

AIRFLOW

#	Requirement	Question	Justification	URL
1	AIRFLOW-1414	["bolke] Any though on that this (the corresponding PR is https://github.com/apache/incubator-airflow/pull/2444) We are trying to switch from Oozie to Airflow at Booking.com and would really need this feature (or maybe there's already a way to achieve this) for us to go forward with the migration.	any thoughts	https://issues.apache.org/jira/browse/AIRFLOW-1414
2	AIRFLOW-1414	["bolke] thanks for your feedback, would you have time to have a quick look at this ticket too: https://issues.apache.org/jira/browse/AIRFLOW-1398?	ask opinion on another issue	https://issues.apache.org/jira/browse/AIRFLOW-1414
3	AIRFLOW-1401	I'm particularly interested in thoughts from ["alexvanboxel"], ["wwlian"], and ["fenglu"]	ask thoughts	https://issues.apache.org/jira/browse/AIRFLOW-1401
4	AIRFLOW-1401	SGTM, maybe we can include other standardization efforts on GCP operators as well (see https://issues.apache.org/jira/browse/AIRFLOW-820)? Let's also make sure to do it in a backwards-compatible way and not break any existing DAGs that depend on an older version of these operators.	ask possible solution	https://issues.apache.org/jira/browse/AIRFLOW-1401
5	AIRFLOW-1398	["bolke] Any though on that this (the corresponding PR is https://github.com/apache/incubator-airflow/pull/2431) We are trying to switch from Oozie to Airflow at Booking.com and would really need this feature (or maybe there's already a way to achieve this) for us to go forward with the migration.	any thoughts?	https://issues.apache.org/jira/browse/AIRFLOW-1398
6	AIRFLOW-1398	["bolke] I implemented your suggestion, would you mind reviewing?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-1398
7	AIRFLOW-1325	To be clear, are you planning to REMOVE support for S3/GCS? We are heavy users of GCS for logging, and this would be a show stopper for us.	ask implementation	https://issues.apache.org/jira/browse/AIRFLOW-1325
8	AIRFLOW-1310	It was exactly the same sort of thinking on my side. Good point about the logs. It was sort of design decision as we capture all logs centrally in ELK so we had no need to stream the logs to airflow as such. It is not hard to add though. Also regarding your comment that it would be more powerful if you can supply a pod definition directly - that's another decision to hide the pod definition as such to make deployment of simple jobs simpler. I have a version where you can give a pod definition instead in case you need to do something complex, but for most jobs we have we just have an image (often with default command) and we don't want to force people to learn yet another library just to run a simple command. Do you mind collaborating to get out something nice and useful to everybody?	ask for collaboration	https://issues.apache.org/jira/browse/AIRFLOW-1310
9	AIRFLOW-1019	["aen], changing this to blocker for 1.8.1. LMK if you don't think we should block the release on it.	Let me know	https://issues.apache.org/jira/browse/AIRFLOW-1019
10	AIRFLOW-983	Was this confirmed yet?	ask status	https://issues.apache.org/jira/browse/AIRFLOW-983
11	AIRFLOW-983	It seems like the commit was reverted but I'm not sure if testing was done: https://github.com/apache/incubator-airflow/pull/2195? If not I can try to take a look soon but just wanted to confirm first ["bolke]	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-983
12	AIRFLOW-979	Please review my Pull Request #2149: https://github.com/apache/incubator-airflow/pull/2149	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-979
13	AIRFLOW-932	verify_integrity in models.py is where you would like to look at. As mentioned it sets a state to REMOVED if a task instance does not have a corresponding task in a DAG. It is there for auditing and versioning reasons. Does it affect running anything? A changed DAG cannot have a dependency on a non-existent task. I.e. please explain why this is a blocker.	ask details	https://issues.apache.org/jira/browse/AIRFLOW-932
14	AIRFLOW-932	Command: <code>airflow backfill testg -t create_and_transfer_roster -s 2017-06-14 -e 2017-06-14</code> DAG: <pre>import airflow from airflow.executors import SequentialExecutor from airflow.models import DAG from airflow.operators import PythonOperator from airflow.operators.subdag_operator import SubDagOperator from airflow.models import DAG from datetime import datetime, timedelta default_args = { 'owner': 'airflow', 'start_date': datetime(2016, 3, 15), } dag = DAG('testg', default_args=default_args) subdag = DAG('testg.create_and_transfer_roster', default_args=default_args) subdag_operator = SubDagOperator(task_id='create_and_transfer_roster', subdag=subdag, executor=SequentialExecutor(), dag=dag) def create_roster(ds, macros, **kwargs): pass PythonOperator(task_id='create_roster', python_callable=create_roster, dag=subdag, provide_context=True).first_task = DummyOperator(task_id='first_task', dag=dag) Expected: first_task TI does not change (stays as none) create_and_transfer_roster TI is run Actual: first_task TI gets its state set to removed create_and_transfer_roster TI is not run, backfill errors with airflow.exceptions.AirflowException: Task first_task not found ["bolke] if you think these are related to your changes (addition of task.REMOVED state/backfill changes) do you mind taking a look? Otherwise let me know and I can do so.</pre>	do you mind taking a look	https://issues.apache.org/jira/browse/AIRFLOW-932
15	AIRFLOW-932	I just ran this on my AIRFLOW-910 branch and I am not seeing this behaviour. I did change your execution date to 2016 as 2017 won't run. <pre>{code} MariaDB [airflow]> select * from task_instance; +-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+ task_id dag_id execution_date start_date end_date duration state try_number hostname unixname +-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+ job_id pool queue priority_weight operator queued_dttm pid +-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+ 2017-03-02 08:58:35.867219 2017-03-02 08:58:46.868470 11.0013 success 1 bolkes-macbook-pro.local bolke 2 NULL default 1 SubDagOperator NULL 46856 create_roster testg.create_and_transfer_roster 2016-06-14 00:00:00.000000 2017-03-02 08:58:42.681159 2017-03-02 08:58:42.702009 0.02085 success 1 bolkes-macbook-pro.local bolke 4 NULL default 1 PythonOperator NULL 46862 +-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+ Can this be related to your local patches?</pre>	ask details about implementation	https://issues.apache.org/jira/browse/AIRFLOW-932
16	AIRFLOW-932	Ah I think I know why it's not reproducing for you. Can you try letting the scheduler run a dagrun, and then backfill the dagrun once it completes (same DAG)? so e.g. let it run for one dagrun via scheduler (might need to unpause it if it's paused) then <code>airflow backfill testg -t create_and_transfer_roster -s 2016-03-15 -e 2016-03-15</code>	task	https://issues.apache.org/jira/browse/AIRFLOW-932
17	AIRFLOW-932	Issue resides in cli.py and will only happen when a specific task is used: <pre>{code} if args.task_regex: dag = dag.sub_dag(task_regex=args.task_regex, include_upstream=not args.ignore_dependencies)</pre> This creates a subset of the tasks from a dag_run, with the same name as the original. Hence it will set a task to removed if you verify the integrity of a dag run, cause that will consider the tree of the original (whole). <code>get_task_instances</code> picks up all instances, including 'removed' ones from the original (whole) dag and this is not filtered in the backfill. Hence the lists mismatch and an AirflowException is thrown. The quick and dirty fix is to not mark removed if run from backfill and filter the list of "get_task_instances". However the functionality if sub_dag is awkward imho - and seems not to be tested, and might need a real fix. What do you think?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-932
18	AIRFLOW-932	PR is available with unit tests ["aen]. Please review	task	https://issues.apache.org/jira/browse/AIRFLOW-932
19	AIRFLOW-931	["bolke] I added a bunch of manual logging to verify this. What seems to happen (in my simple gist example) is DAG starts with 4 tasks. The multiple scheduler processes kick them off (the jobs.py:1095) because at that time they are all runnable (since none are yet running). That models.py:1291 does another check to make sure it is actually runnable (2 of them are, but then due to the dag concurrency of 2, the third and fourth one aren't). The 2 tasks stay marked as queued but never actually get picked up by the scheduler again (unless you restart it). Is there some other state that that models.py:1291 check could move the tasks to if that last minute check found they aren't actually runnable? (I found by making them State.NONE it worked, but that seems hackish as it keeps bouncing back and forth between QUEUED and NONE until it can actually run)	ask details	https://issues.apache.org/jira/browse/AIRFLOW-931
20	AIRFLOW-920	["aen] It does propagate in its current form: if you mark success a "subdag" operator from its parent dag it will mark success all tasks part of that subdag. My PR just verifies the tree of the dag to make sure dag runs are created for any missing subdag dag runs. I consider this to be something of a hack as it is in my opinion the wrong place to do it (the webserver). BTW: Also the unit tests confirm propagation. So are we talking about the same thing?	ask confirmations	https://issues.apache.org/jira/browse/AIRFLOW-920

#	Requirement	Question	Justification	URL
21	AIRFLOW-908	Can i take this feature ?	ask status of who is working	https://issues.apache.org/jira/browse/AIRFLOW-908
22	AIRFLOW-908	[~aoen] Would you give me an instruction ?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-908
23	AIRFLOW-891	Any reason why you are suggesting only adding the day like that? It isn't very intuitive to read and could be misunderstood as some time between 21:29 and 23:29. I would suggest keeping the same time format as we use everywhere else YYYY-MM-DD HH:MM:SS	ask why this solution	https://issues.apache.org/jira/browse/AIRFLOW-891
24	AIRFLOW-890	Are the example endpoints just the two in the docs? Specifically: /api/experimental/dags/<DAG_ID>/tasks/<TASK_ID> returns info for a task (GET). /api/experimental/dags/<DAG_ID>/dag_runs creates a dag_run for a given dag id (POST). Or is there a roadmap somewhere?	ask about docs	https://issues.apache.org/jira/browse/AIRFLOW-890
25	AIRFLOW-846	Hi, I see that 1.8 has finally been tagged, is there any plan to release it to pypi ? How is that automated/processed ? Cheers	ask status	https://issues.apache.org/jira/browse/AIRFLOW-846
26	AIRFLOW-846	BTW [~artwr], would you please be kind enough to point me where this vote / release tracking is taking place ? Looks like I'm lost on JIRA...	ask system	https://issues.apache.org/jira/browse/AIRFLOW-846
27	AIRFLOW-679	I see this has a closing PR. Is this ticket still relevant or can it be closed?	ask status	https://issues.apache.org/jira/browse/AIRFLOW-679
28	AIRFLOW-539	Not sure if you saw the note on the related issue, but perhaps this could be implemented via a kwarg passed through to the API, so that one-off feature changes don't require end-to-end plumbing in the future?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-539
29	AIRFLOW-539	I agree re: kwargs, but I'd rather keep it as a separate thing from this just for simplicity's sake. Could you open a second JIRA to use kwargs for Google hook/operator params?	ask task	https://issues.apache.org/jira/browse/AIRFLOW-539
30	AIRFLOW-539	[~criccomini]: although this ticket is closed (can you add a "gcp" label on this ticket), you're probably an admin. Thanks.	task	https://issues.apache.org/jira/browse/AIRFLOW-539
31	AIRFLOW-513	What's wrong with workaround 2? The point of parallelism is to control how many processes are running. A sensor that's waiting is in a while loop eating processing time, so you don't really want to ignore them for parallelism's sake. Short timeout, with tuned retry delay, and retry counts, will work just fine.	ask details of issue	https://issues.apache.org/jira/browse/AIRFLOW-513
32	AIRFLOW-513	Porting convo from gitter to ticket: {code} [kyuen] option 2 is a valid way of preventing the problem. Just created the issue to see what people think the sensors should do. Didn't want to make a pr before seeing how people feel about how the sensors should behave [George Leslie-Waksman] Sensor doesn't really give the best sense for how they work. I think @mistercrunch was talking about more complex operations in the scheduler as an alternative way of doing things, but I think that's still a work in progress. {code} Would love to hear more from others on how they'd like sensors to behave	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-513
33	AIRFLOW-511	[~criccomini] any thoughts on this?	ask thoughts	https://issues.apache.org/jira/browse/AIRFLOW-511
34	AIRFLOW-479	Can I take this issue?	ask status	https://issues.apache.org/jira/browse/AIRFLOW-479
35	AIRFLOW-461	This is for GCS to BQ operator, right?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-461
36	AIRFLOW-454	Where in the docs would you place this? Submit a PR if you think this is a necessary feature. For example, if someone were to require this link in the natural process of reading through a getting-started guide, I'd say it makes sense. Otherwise, I'd expect folks who are working on PRs to use Google to look into Flask-Admin, SQLAlchemy, or someone other 3rd party framework that Airflow leverages.	ask about docs	https://issues.apache.org/jira/browse/AIRFLOW-454
37	AIRFLOW-445	[~jlowin][~criccomini] What do you think?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-445
38	AIRFLOW-409	[~mtdewulf] I'm considering taking ownership of this issue since you mention on your PR: "Since I've posted this PR, my team has decided that Airflow is not appropriate for our use case, so I probably won't be able to spend any cycles updating it." I assume you're happy for me to assign this to myself and take it over?	ask who is working	https://issues.apache.org/jira/browse/AIRFLOW-409
39	AIRFLOW-382	[~maxime.beauchemin@apache.org], {noformat} While we're [decreasingly] comfortable installing all sorts of packages (R packages, Python libs, related apt packages, ...) on Airflow workers, we clearly don't want any of that on the Yarn boxes directly.{noformat} It could be easier to implement once Docker on YARN infrastructure is complete? YARN-3611, YARN-4793 https://hadoop.apache.org/docs/r2.7.2/hadoop-yarn/hadoop-yarn-site/DockerContainerExecutor.html	possible solution?	https://issues.apache.org/jira/browse/AIRFLOW-382
40	AIRFLOW-374	> The task can monitor the time since last heartbeat and kill itself to prevent such cases. In your example, if the task can't get access to the DB how does it know? Are you planning on keeping some local in-mem variable that tracks the last successful heartbeat write?	ask questions	https://issues.apache.org/jira/browse/AIRFLOW-374
41	AIRFLOW-350	Could you add some more description to this ticket? I don't understand it.	ask description	https://issues.apache.org/jira/browse/AIRFLOW-350
42	AIRFLOW-340	Also, should I track other unreferenced dependencies? I'm sure there's a couple of others lying around. If so, should I open an another ticket/PR to cover this extended work?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-340
43	AIRFLOW-331	Ah, nice catch. Can you send a pull request, please?	ask PR	https://issues.apache.org/jira/browse/AIRFLOW-331
44	AIRFLOW-329	I'm a little concerned about performance. The dag runs page, for example, barely loads for us, since we have 100s of DAGs at a 15m interval. Perhaps recent DAG runs is better?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-329
45	AIRFLOW-325	Isn't this possible with an ExternalTaskSensor(same_dag_id, previous_task_id, allowed_states=[State.SUCCESS, State.FAILED])?	provide workaround as option	https://issues.apache.org/jira/browse/AIRFLOW-325
46	AIRFLOW-312	Can you try running off of master? I'm wondering if my PR in AIRFLOW-200 fixes this.	task	https://issues.apache.org/jira/browse/AIRFLOW-312
47	AIRFLOW-288	Thoughts on support for this? It would be nice to see this at the DAG level so individual DAGs could be triggered in a timezone aware way by the scheduler.	ask thoughts	https://issues.apache.org/jira/browse/AIRFLOW-288
48	AIRFLOW-276	Considering https://issues.apache.org/jira/browse/AIRFLOW-758 , would it not be more straight forward to just revert back to execpv and run gunicorn with the --reload argument? Granted, this functionality is listed under the "Debugging" section of gunicorn's settings documentation, but it might be still be appropriate for this use case. Link to gunicorn docs: http://docs.gunicorn.org/en/19.0/settings.html#reload	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-276
49	AIRFLOW-275	[~bolke], you committed the PR for this, right?	ask status	https://issues.apache.org/jira/browse/AIRFLOW-275
50	AIRFLOW-258	Makes sense to me. Want to send a PR?	ask PR	https://issues.apache.org/jira/browse/AIRFLOW-258
51	AIRFLOW-258	Hey guys - this seems like the closest thing to an issue I'm seeing. I have a custom executor and a custom operator as a plugin. If I specify my custom executor in the config, airflow takes a different path than if I don't specify using the custom executor. If I'm configured with a built in executor like LocalExecutor, the plugin loads fine. If I'm configured to use my custom executor, then it loads the plugins which also loads the modules my custom operator is defined in, which is normal. I'm subclassing BaseSensorOperator and as its importing modules it throws an error about not being able to import BaseOperator from airflow.models from airflow/operators/_init_.py. Long story, any ideas on how to tackle this? Would be glad to work on a PR with a bit of guidance. Thanks!	ask solutions	https://issues.apache.org/jira/browse/AIRFLOW-258
52	AIRFLOW-249	This is really awesome analysis. Can you send a PR? (unless I missed it)	ask PR	https://issues.apache.org/jira/browse/AIRFLOW-249
53	AIRFLOW-247	[~sanand], not sure what the protocol is for creating issues related to something I'd like to implement. Is what I've done here acceptable?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-247
54	AIRFLOW-247	[~rfoetscher], are you still planning to work on this? If not, we can re-assign to [~gjreda].	ask status	https://issues.apache.org/jira/browse/AIRFLOW-247
55	AIRFLOW-247	[~gjreda], want to have a look at the PR?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-247

#	Requirement	Question	Justification	URL
56	AIRFLOW-247	I'm searching for documentation related to how Airflow works with EMR. I'm struggling to find anything here: https://airflow.incubator.apache.org/integration.html#aws My main question is, can Airflow create an EMR cluster and bring it back down like AWS Data Pipeline? Thanks! EDIT: Found some information here: Spark, EMR: - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/example_dags/example_emr_job_flow_automatic_steps.py - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/example_dags/example_emr_job_flow_manual_steps.py - (uses emr hooks, operators) https://docs.google.com/presentation/d/1NG1P86HRIX43qTVucCTOfqIbCvYdOhq_np90VlbVRC/edit#slide=id.gd40e667_1_0 - (uses shells scripts to launch and terminate emr clusters) https://www.agari.com/automated-model-building-emr-spark-airflow/ - (use shell script to spark-submit on a local spark installation) https://blog.insightdatascience.com/scheduling-spark-jobs-with-airflow-4c66f3144660 - (installs spark on each airflow worker node and runs local spark jobs without use of spark submit) https://medium.com/@calvertmg/airflow-integrating-with-apache-spark-50a7704dcebd - (alternative mozilla implementation for emr spark job) https://github.com/mozilla/telemetry-airflow/blob/master/dags/operators/emr_spark_operator.py EMR: - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/hooks/emr_hook.py - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/operators/emr_create_job_flow_operator.py - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/operators/emr_add_steps_operator.py - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/operators/emr_terminate_job_flow_operator.py Spark: - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/hooks/spark_submit_hook.py - https://github.com/apache/incubator-airflow/blob/master/airflow/contrib/operators/spark_submit_operator.py	ask question	https://issues.apache.org/jira/browse/AIRFLOW-247
57	AIRFLOW-245	[~jlowin] Thanks for the feedback. But I don't quite get it. There's 'BaseExecutor' class that all executors derived from, it has a field 'queued_tasks' that's a dictionary of tasks where the key is the task key and the value is a tuple of '(command, priority, queue, task_instance)'. During heartbeat we already [retrieve the task from the db. https://github.com/apache/incubator-airflow/blob/f4c44146cd38afe21c517230bff71d33ab041246/airflow/executors/base_executor.py#L110] We can pass to the 'execute_async' method. I have a question then: I write a custom executor. How can I reuse available Operators and their fields in the executor. Say I write executor for Mesos, I want to have access to Docker operator fields from there and transform them to an HTTP call. What would be the preferable way to do that in Airflow?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-245
58	AIRFLOW-245	The Scheduler is an object that decides what tasks to run; the Executor is an object that decides where to run those tasks. So the Scheduler hands a TaskInstance to the Executor and says "I want to run this task". That's why the Executor has access to the TaskInstance when tasks are being queued. The Executor then extracts the TI's "command" (for example: airflow run my_task) and its sole job is to decide how to run that command. execute_async() is where that decision is made. The important thing is that the thing being run is the string command, not the TI object, and the signature of execute_async should reflect that. The larger question here is why would an operator ever influence the executor? An Operator is a series of instructions that are designed to be run in a self-contained way, no matter what the executor is. An executor is a way of kicking off an Operator. There should be no dependency between them. To be explicit: an Operator must be able to run under any Executor. However, I am not familiar with Mesos so unfortunately I don't know exactly what you're trying to do -- but have you looked at the existing MesosExecutor in airflow.contrib.executors? Is it insufficient for your objectives? If your operator requires special handling, then it sounds like maybe you want a MesosOperator that runs your docker command?	ask why	https://issues.apache.org/jira/browse/AIRFLOW-245
59	AIRFLOW-199	This can be done via the tree view by clicking on the failed task(s), and clicking "clear" with "downstream". Is this what you're looking for? It will allow the DAG to resume.	ask confirmations	https://issues.apache.org/jira/browse/AIRFLOW-199
60	AIRFLOW-177	[~criccomini] I was thinking that would it be helpful to have a cli support for that.. say {{ resume_dag }} which takes a execution_date as input and rerun only failed tasks?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-177
61	AIRFLOW-169	What do you mean by expired schedules? You mean DagRuns that have finished? Which page in the UI are you referring to?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-169
62	AIRFLOW-167	This sounds reasonable to me. Want to send a PR?	ask PR	https://issues.apache.org/jira/browse/AIRFLOW-167
63	AIRFLOW-167	Specifically, the CLI solution sounds reasonable. The REST solution is better, but we haven't yet set up a REST API for Airflow yet. In the mean time, want to send a PR for the CLI {{dag_state}} command?	ask PR	https://issues.apache.org/jira/browse/AIRFLOW-167
64	AIRFLOW-165	I think this is already done, right? https://pythonhosted.org/airflow/concepts.html#task-documentation-notes You can do: <code>t = BashOperator("foo", dag=dag) t.doc_md = ""\n #Title" Here's a [url](www.airbnb.com) ""\n [code]</code>	ask status	https://issues.apache.org/jira/browse/AIRFLOW-165
65	AIRFLOW-161	What context is this under? You mean like you want to use the EmailOperator or SlackOperator to notify people to download a file that's been created as part of the DAG? We do this using XCom+EmailOperator. XCom variables can be accessed via templates. We store the file in a blob store (like S3). The file location is stored in XCom, which the EmailOperator references when it sends the email.	ask details	https://issues.apache.org/jira/browse/AIRFLOW-161
66	AIRFLOW-161	I don't think that we want to embed Quoble logic directly in Airflow. I'm a bit out of my element on the UI-front, though. Perhaps there's a way to achieve this through plugins, or by simply putting the link in the logs?	ask solutions	https://issues.apache.org/jira/browse/AIRFLOW-161
67	AIRFLOW-161	Sure, can you please cc top contributors to PR or here.	ask task	https://issues.apache.org/jira/browse/AIRFLOW-161
68	AIRFLOW-135	I'm +1 on cutting branches for releases (instead of tags). I'm also +1 on deleting dead branches, but could you post a list of the proposed branches to delete so we can allow people time to notify if they have a branch that they don't want deleted?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-135
69	AIRFLOW-129	[~maxime.beauchemin@apache.org], [~aoen], [~artwr], could one of you take a look at this? I think you guys are the big Celery users.	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-129
70	AIRFLOW-128	[~bolke], OK to close this?	ask status	https://issues.apache.org/jira/browse/AIRFLOW-128
71	AIRFLOW-123	What exactly do you define as a timeout? # Task is killed because it's a zombie # A connection to some infrastructure times out # A task runs for longer than it should (1) seems to be the one that Airflow could do right now. (2) is hard--it requires cooperation from the operator to throw some kind of known TimeoutException (unless you want to use Python built-in TimeoutError). (3) I supposed could be handled as a task-level SLA.	ask details	https://issues.apache.org/jira/browse/AIRFLOW-123
72	AIRFLOW-121	Hello. I think you should file a new bug issue with an example DAG that triggers this bug. Please also specify which version you are running.	ask details but this is not quite about this requirement	https://issues.apache.org/jira/browse/AIRFLOW-121
73	AIRFLOW-118	Was there a PR for this?	ask status	https://issues.apache.org/jira/browse/AIRFLOW-118
74	AIRFLOW-115	[~artwr] Are you working on this issue? I'm facing many errors when using the current 'S3Hook' mostly caused by it using 'boto' instead of 'boto3' (e.g. https://github.com/boto/boto/issues/2836)	ask status	https://issues.apache.org/jira/browse/AIRFLOW-115
75	AIRFLOW-113	[~Shenghu], can you: # Test this against master? # Post your DAG (or some equivalent) Python code so we can try and replicate	tasks	https://issues.apache.org/jira/browse/AIRFLOW-113
76	AIRFLOW-113	Also, I wonder if perhaps you want max_active_runs instead of concurrency? From [FAQ](https://pythonhosted.org/airflow/faq.html): {quote} concurrency defines how many running task instances a DAG is allowed to have, beyond which point things get queued {quote} Whereas: {quote} Is the max_active_runs parameter of your DAG reached? max_active_runs defines how many running concurrent instances of a DAG there are allowed to be. {quote}	tasks	https://issues.apache.org/jira/browse/AIRFLOW-113

#	Requirement	Question	Justification	URL
77	AIRFLOW-111	Thanks [~cricomini]. I fixed the description by changing 'jobs' to 'task instances'. Can you please help on review the related PR: https://github.com/apache/incubator-airflow/pull/1579 ?	ask review	https://issues.apache.org/jira/browse/AIRFLOW-111
78	AIRFLOW-101	[~bolke], which PR was this part of?	ask traceability	https://issues.apache.org/jira/browse/AIRFLOW-101
79	AIRFLOW-93	[~underyx] I'm rapidly learning about ExternalTaskSensor but could you not create 6 ETS in your hourly DAG, one for each of the 10 minute tasks (so one with delta of 10 minutes, one with 20 minutes, one with 30 minutes, etc.) and then have your hourly task depend on all of those sensors completing successfully? See for example AIRFLOW-47.	ask detailss	https://issues.apache.org/jira/browse/AIRFLOW-93
80	AIRFLOW-91	ok i pushed new PR https://github.com/apache/incubator-airflow/pull/1497 but i really misunderstand (if i replace 'ssl_certfile': None, section) i got : (color:red) Traceback (most recent call last): File "/usr/local/bin/airflow", line 13, in <module> parser = get_parser() File "/usr/local/lib/python2.7/dist-packages/airflow/bin/cli.py", line 689, in get_parser default=configuration.get('webserver', 'ssl_keyfile'), File "/usr/local/lib/python2.7/dist-packages/airflow/configuration.py", line 520, in get return conf.get(section, key, **kwargs) File "/usr/local/lib/python2.7/dist-packages/airflow/configuration.py", line 428, in get "in config".format(**locals())) airflow.configuration.AirflowConfigException: section/key [webserver/ssl_keyfile] not found in config (color) so please check ones more my changes ? thanks !	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-91
81	AIRFLOW-91	i update related issue, pls take a look ones more ?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-91
82	AIRFLOW-91	hi for all, i already update PR, plz take a look on it? thanks !	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-91
83	AIRFLOW-81	What do you think about writing a new sensor based off of the TimeSensor that blocks during the blackout period?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-81
84	AIRFLOW-79	Is this dead in the water? I've been reading about it in the pull request and on the mailing list and on the Gitter channel...	ask status	https://issues.apache.org/jira/browse/AIRFLOW-79
85	AIRFLOW-78	[~abridgett] Can you tell me more about your expectations of clear?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-78
86	AIRFLOW-78	Assigning this to Norman. We can clear dag runs issued by clear in both the UI and CLI. [~bolke] [~jlowin][~maxime.beauchemin@apache.org] Do you see any issue with this idea?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-78
87	AIRFLOW-78	[~abridgett]: Just to clarify, the issue here is still that the cleared dagruns get rerun without respecting the max_active_runs amount?	ask clarification	https://issues.apache.org/jira/browse/AIRFLOW-78
88	AIRFLOW-78	[~abridgett]Is max_active_runs in this scenario set to 1? Did you set it to depend on upstream? And nothing else is running?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-78
89	AIRFLOW-71	[~asnir], can you describe what this is a bit more?	ask description	https://issues.apache.org/jira/browse/AIRFLOW-71
90	AIRFLOW-71	[~asnir], do you want add config to the DockerOperator, such that it can execute a 'docker login' before it runs the Docker command?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-71
91	AIRFLOW-71	Hi [~cricomini], I'm really needing it, and I think it's simple to do that... I was thinking in: Expose the Client (docker-py.client.Client) that is instantiated inside of DockerOperator.execute() method, changing it to instantiate in __init__(), so we will be able to execute the login passing username, password, and the specific registry too.... eg: d = DockerOperator(task_id='example_task_id', image='example/image-name', force_pull=True, owner='felipetancredo', command="/bin/bash", registry='https://registry.example.com', reauth=True) Can I contribute to docker_operator? Do u see any problem? (sorry, I'm beginner in the open source projects)	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-71
92	AIRFLOW-71	When will [~asnir]'s PR be approved? It's been 8 months.	ask status	https://issues.apache.org/jira/browse/AIRFLOW-71
93	AIRFLOW-58	[~maxime.beauchemin@apache.org], any reason why this was closed?	ask why closed	https://issues.apache.org/jira/browse/AIRFLOW-58
94	AIRFLOW-58	[~sanand], did you mean to put a real external URL in? You just put the JIRA in. :)	ask details	https://issues.apache.org/jira/browse/AIRFLOW-58
95	AIRFLOW-55	Overall, this looks good to me. Given that it's in contrib, and basically can't be supported by the committers (unless one of us wants to use it?), I'm OK without tests. [~sanand], what's your take?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-55
96	AIRFLOW-54	for this feature as well. cc @jlowin (the original creator of xcom) [~amustafa] Would you like to take this on? The one case you need to think about is subdags. I think the relationships above translate across subdags, but please do test with subdags as well. Also, please change 'upstream' to 'parent'	ask if you want	https://issues.apache.org/jira/browse/AIRFLOW-54
97	AIRFLOW-54	[~amustafa] Any PRs that we can look at?	ask status	https://issues.apache.org/jira/browse/AIRFLOW-54
98	AIRFLOW-52	[~cricomini] Can we do an 'airbnb' release for now then? I.e. not use the apache brand? Cause this release is long overdue and 1.7.0 is a bit buggy. Chasing licenses will cost at least another week or even more (with broken PRs to add)	ask status	https://issues.apache.org/jira/browse/AIRFLOW-52
99	AIRFLOW-52	[~jghoman]/[~hitesh]/[~cnauroth] what's the best way to handle this? We basically had an RC in flight for like a month before Apache.	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-52
100	AIRFLOW-52	So, the release will be coordinated (and built from) github and github resources? I think Chris is right, but I'd be concerned about running the vote or whatever process was being used for the release on ASF resources. I'd want to avoid the appearance of trying to make an end-run around asf procedure (which I know isn't the case). I'm honestly not sure though; may be worth pingging general@incubator.	ask about release	https://issues.apache.org/jira/browse/AIRFLOW-52
101	AIRFLOW-52	[~aoen], are you planning to burn in [~jlowin]'s fix before release?	ask release statys	https://issues.apache.org/jira/browse/AIRFLOW-52
102	AIRFLOW-52	This one right https://github.com/apache/incubator-airflow/pull/1473 ? In this case yes.	ask assurance	https://issues.apache.org/jira/browse/AIRFLOW-52
103	AIRFLOW-52	Should this be closed?	ask status	https://issues.apache.org/jira/browse/AIRFLOW-52
104	AIRFLOW-51	Shouldn't we support both, rather than deleting the 'cpus' field? As I read it, it sounds like Docker supports both CPU count and cpu_shares percentage.	ask details	https://issues.apache.org/jira/browse/AIRFLOW-51
105	AIRFLOW-48	It's likely not implemented because no one is using it. Can you provide an example of your uri? Is it something like *postgres://localhost:2345/your_db?sslmode=require*? What version of postgresql are you running? I'm not sure how well this parameter is supported for versions of postgresql before 9.2. Another alternative is to take control of this by implementing this logic in a PythonOperator's python callable. That could read this param and use psycopg2 to connect to your db. Of course, then you would not be using the Connection or Hook, but, you would be able to achieve your end goal. In the meantime, feel free to submit a PR for review -- we'd be happy to accept it.	ask example	https://issues.apache.org/jira/browse/AIRFLOW-48
106	AIRFLOW-45	What do you think about defining a heuristic that variables would be hidden from the UI if they contain the text: hidden, password, secret	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-45
107	AIRFLOW-45	{quote} I think we should have a configuration variable to hide/expose the encrypted values in the UI for both Variables and Connections. It should likely affect all Variables and/or Connections in a given Airflow installation and not be DAG specific. {quote} This seems OK to me. So, current proposal is: # Add a config param to airflow.cfg called hide_encrypted_ui_fields # Default hide_encrypted_ui_fields to true The hide_encrypted_ui_fields param would hide 'passwords' and 'extras' in the hooks view, as well as the 'value' of variables in the variables view. Does that sound OK?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-45
108	AIRFLOW-45	Sounds good. [~cheny258], would you like to try this?	ask assign	https://issues.apache.org/jira/browse/AIRFLOW-45

#	Requirement	Question	Justification	URL
109	AIRFLOW-45	[~cheny258] As part of this Variables work, there are currently 2 views. One is the edit view, the other a list view of all variables. Currently, the list view only shows the variables names, not the values. i'd like to display the values if we the configuration param proposed above is set to False. Can you add this feature? Currently, one must click edit to view a variable value.	ask status/doability	https://issues.apache.org/jira/browse/AIRFLOW-45
110	AIRFLOW-45	[~sanand] will do that for the list view of variables. Also, for conn edit page, {{extra}} and {{password}} are already in a hidden_fields list. Do we want to make any changes in edit or list view for that?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-45
111	AIRFLOW-45	Could we not handle this with a regexp of fields that should be hidden rather than a boolean?	ask solution	https://issues.apache.org/jira/browse/AIRFLOW-45
112	AIRFLOW-45	{quote} I'd have to go change the config whenever I do this, which is kind of annoying {quote} (Yes, quoting myself :). Perhaps not. As long as the regex was something like '.*secret.*', and I stuck to that as a naming convention, I'd be fine. Still, to limit burden on [~cheny258], I'd like to keep the regex to a separate ticket. I'll open it up after this one is closed off. [~withnale]/[~sanand], does that sounds OK?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-45
113	AIRFLOW-39	Can you provide sample code?	ask example	https://issues.apache.org/jira/browse/AIRFLOW-39
114	AIRFLOW-39	Oh, so not version 1.7.0? What does 'git log head' look like for you?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-39
115	AIRFLOW-35	[~sanand], any particular reason you closed this?	ask why status	https://issues.apache.org/jira/browse/AIRFLOW-35
116	AIRFLOW-31	I'm not sure where the best place to comment is... Wherever you're comfortable? Let me know if you discover anything -- in my experience getting objects to behave like modules is tricky (supporting full {{import}} syntax etc). Maybe some of what I did for plugins would be reusable?	LMK	https://issues.apache.org/jira/browse/AIRFLOW-31
117	AIRFLOW-30	Initially, this seemed fine to me, but it actually sounds a little tricky to me. The concern I have is what happens if the DAG fails half way through and the transaction is reverted? In such a case, Airflow will show prior operators as having successfully run, but their state mutation will not have taken place. This effectively means that the entire DAG needs to be re-run, even if some of the operators show a successfully executed view.	ask what happens in use case	https://issues.apache.org/jira/browse/AIRFLOW-30
118	AIRFLOW-30	In your use cases, are the source and destination DBs part of the same DB cluster?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-30
119	AIRFLOW-29	One thing to consider: what happens if the scheduler machine is overloaded, and imports are taking a long time because of that? (similar to a GC pause in Java).	ask use case	https://issues.apache.org/jira/browse/AIRFLOW-29
120	AIRFLOW-29	[~underyx], re (1), I wasn't suggesting GC would be an issue--was just saying it's a similar problem to timeouts in Java when GC occurs. Anyway, I agree, probably not worth worrying about. What default did you have in mind? 5s?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-29
121	AIRFLOW-29	[~underyx] I'm not connecting the configuration suggestion to how such a change would resolve your problem. Do you mind providing a little more context on impact of such a change? If the import is actually only taking 1 second on your box and if you are not hitting a timeout, then why does the larger configured value have an impact on you?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-29
122	AIRFLOW-28	Okay, but what's preventing someone else from implementing this? How come the issue was closed as won't fix?	ask why staus	https://issues.apache.org/jira/browse/AIRFLOW-28
123	AIRFLOW-23	Hey [~alexvanboxel], checking in to see how this and AIRFLOW-24 are coming.	ask status	https://issues.apache.org/jira/browse/AIRFLOW-23
124	AIRFLOW-22	Can public buckets/BQ datasets (thinking about GDELTT) be used without a project?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-22
125	AIRFLOW-22	What's GDELTT? In the case of GCS, yes I believe we can use public buckets. I'm uncomfortable baking in someone's own buckets, though. It'd have to be public buckets that Google, themselves, provide. For BQ, unfortunately not. The problem with BQ is that it requires a project to run its queries on. Even if the datasets are public, you still need to have access to a whole project, and be a viewer for it, in order to execute any queries against the data.	ask details	https://issues.apache.org/jira/browse/AIRFLOW-22
126	AIRFLOW-21	Do you need some help? I could do it.	ask status/help	https://issues.apache.org/jira/browse/AIRFLOW-21
127	AIRFLOW-21	[~alexvanboxel], that would be awesome. Can you wait until AIRFLOW-16 goes in, though, to avoid merge conflicts? Should be today.	task-ish	https://issues.apache.org/jira/browse/AIRFLOW-21
128	AIRFLOW-20	The DagRun.previous handling is used to connect previous related DagRuns in order to make depends_on_past work with a moving interval or an automatically aligned start_date. Imagine the following Use Case 1 Dag1 has an assigned start_date "2016-04-24 00:00:00" with a cron schedule of "10 1 * * *", ie run every day at 1.10am. Depends_on_past is true. The current scheduler kicks off a run "2016-04-24 00:00:00" and then stops as the previous interval of "2016-04-24 01:10:00" is "2016-04-23 01:10:00". This cannot be found as it has not taken place and thus the whole thing grinds to a halt and the TaskInstances refuse to run. What the user expects here is that the first run is "2016-04-24 01:10:00", ie start_date + interval, unless the start_date is on the interval, ie. start_date is first interval. This is what I address by start_date normalization in the PR. However, the second issue then kicks in as the "previous" run can still not be found. Use Case 2 Dag2 has an assigned start_date "2016-04-24 01:10:00" with a cron schedule of "10 1 * * *". Depends_on_past is true. The scheduler happily goes through a few runs, but then the dag is updated and the schedule adjusted. Because the previous interval cannot be found by the TaskInstance (emphasis), tasks get stuck again requiring an manual override. What the user expects here is that the scheduler is smart enough to figure out that we are still running the same dag and that it needs to look up the previous run for that dag and make sure dependencies are met with that previous dagrun in mind. I don't think those two use cases are edge cases, considering the amount of questions we get on these subjects. To resolve the remaining issues (aside from start_date normalization) I first made DagRun aware of its predecessors. Then I strengthened the relationship between TaskInstances and DagRuns slightly, by optionally including a dagrun_id in the TaskInstance. Now a TaskInstance can lookup its predecessors in the previous DagRun and know for sure that it is either the first run or it has a predecessor somewhere in time instead of guessing. What I am unsure of is what to consider is the unit of work: 1) is a TaskInstance dependent on its past predecessor ignoring the outcome of the DagRun? (previous TaskInstance can be successful, but previous DagRun as a whole can fail) or 2) is it dependent on the outcome of the DagRun? 3) Can it be both? In case of 1 and 3 my logic needs to be updated slightly, but that should not be too much of a big deal. However I have difficulty imagining why you want to do that.	ask solutions/details	https://issues.apache.org/jira/browse/AIRFLOW-20
129	AIRFLOW-20	Max I understand your concerns, I'm working with Jeremiah to see if I can remove the need for the dag run id in taskinstances. I would be happy to as it would simplify the change. Rigorous testing needs to take place before making it part of a release, current unit tests do not cover enough. Documentation is forthcoming but I will only assemble it when all assumptions have been worked out and discussions finished. Please also note that we (Jeremiah and I) are seeing this as the first step towards his bigger pr. Sort of paving the way. If you can regard it in that context and not just "making it easier for schedules that move". In that regard the more fundamental questions I have also asked on the list have not been answered yet: what do we consider the unit of work? Is it a DagRun or is it a taskinstance? We talk about dags, but we can run tasks without a DagRun. I consider taskinstances part of a DagRun. They should not be able to look beyond the borders of the containing DagRun. If they want to do so they should query the DagRun for this. If you do this backfills and so on become much easier to handle.	ask details	https://issues.apache.org/jira/browse/AIRFLOW-20
130	AIRFLOW-20	I took the plunge and made the "dag_run_id" field for TaskInstances NOT NULL. This allowed me to catch places where TaskInstances are manipulated and to fix these. So indeed there are some strange places like the PythonOperator and views.py that create TaskInstances. That is really a refactor task towards the future (eg. dao principle would be smart maybe) The good news is these have all been fixed. All tests pass, including some additional ones I added. Now I can dive in to some other cases. Like what to do when a backfill has an execution_date that is equal to a previous scheduled run or backfill? 1. Do we update the tasks and run new ones and skip ones that were successful? If so do we connect old ones to the backfill run or do we leave them as is (ie connected to the old run). Think lineage here. How do we handle dependencies if the old tasks stay connected to their previous run? 2. Do we refuse to run? That would disallow you to "add tasks" in the past. I tend to lean towards no 1 and connect all tasks to the backfill and setting the state of the "old" dagrun to "SUPERSEDED" or "OVERRIDDEN". The same flexibility as today would exist but an audit trail is preserved.	ask details	https://issues.apache.org/jira/browse/AIRFLOW-20

#	Requirement	Question	Justification	URL
131	AIRFLOW-17	Got it. Is there any way to fix it?	ask solutions	https://issues.apache.org/jira/browse/AIRFLOW-17
132	AIRFLOW-16	Pull request up at: https://github.com/airbnb/airflow/pull/1452 This PR should also fix: https://github.com/airbnb/airflow/issues/1354 [~jlowin]/[~alexvanboxel], I could use a review from you.	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-16
133	AIRFLOW-16	> Got it working with the new hooks Did you find any bugs, or started up without incident? > GCP operations and hooks be prefixed with gcp ... Could be a ticket, I'll pick it up (after this is merged). Sounds good. Please do! Go ahead and open the JIRA.	ask details about run	https://issues.apache.org/jira/browse/AIRFLOW-16
134	AIRFLOW-15	Pull request: https://github.com/airbnb/airflow/pull/1448 [~jlowin], could use a review.	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-15
135	AIRFLOW-15	[~criccomini] is this issue resolved?	ask status	https://issues.apache.org/jira/browse/AIRFLOW-15
136	AIRFLOW-14	Overall, major +1 on this. One thing that I think is glossed over is what happens when a DagRunJob dies ({kill -9}), for example)? You mention heartbeats, but I don't see a mention of who is going to clean up/take over the execution of the wedged DagRuns when their job dies. I see this problem a lot with Airflow right now (DagRuns stuck in 'running' forever), so it'd be good to think through.	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-14
137	AIRFLOW-14	If I'm not mistaken, I think this work should solidify running multiple schedulers as a first-class citizen, right? We are not going to run Celery for distributed execution, so having a trustworthy way to run more than one scheduler is useful. /cc [~sanand]	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-14
138	AIRFLOW-14	How does a DRJ grab a Dag to execute? It seems to me like there's a race condition where two DRJs could simultaneously see that a DAG needs to be run, and both could insert entries into the dag_run table simultaneously. Is this being enforced by a primary key/unique constraint in the DB?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-14
139	AIRFLOW-14	> DRJ needs to make sure to refresh from the db to check for a lock immediately before running it Isn't there a race condition where two DRJs check the lock in the DB before running, both see it's unset, and both set it to themselves? Even if they then check to verify that they are the owners of the lock, it's not trustworthy. The only way I can see this working is if a real read/write lock is obtained. Is SQLAlchemy doing this?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-14
140	AIRFLOW-14	The good news is that lock_id is already set to the job's id, so I just need to add the additional check that it matches (right now it just assumes if it was able to lock, then it owns the lock, which is not as strict). But I'm still curious about your point: > Even if they then check to verify that they are the owners of the lock, it's not trustworthy. Is that not still an issue?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-14
141	AIRFLOW-14	To clarify, would different DagRuns be evaluated in the same process, or in different processes? In the current work-in-progress branch, I noticed that task scheduling is done by each dag run. However, with the current approach, pool limits may not be correctly respected?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-14
142	AIRFLOW-14	While I understand that for being truly distributed it would be great to have a distributed scheduler, are not not looking for a distributed executor instead [~criccomini]? Now the scheduler and the executor are coupled, but that does not need to be and does not even make sense imho.	ask details	https://issues.apache.org/jira/browse/AIRFLOW-14
143	AIRFLOW-14	{quote} Now the scheduler and the executor are coupled {quote} Can you expand on this a bit? I know [~sanand] is running multiple LocalExecutor schedulers. We were planning to do this as well.	ask details	https://issues.apache.org/jira/browse/AIRFLOW-14
144	AIRFLOW-14	Now the scheduler is, in case of using the LocalExecutor, tied to the executor: ie. both are same process. This is not really required and maybe not even the best way forward. A scheduler should, imho, only schedule and an executor should execute. With celery you already get a kind of decoupling. The celery executor puts tasks on a queue and then celery workers pick these up. It is not so difficult to imagine that the local executor can do the same and pick up tasks for a queue that is either something like rabbitmq or just the db. In other words the scheduler should queue tasks in a certain order and the executor should pick these up from anywhere. ie. producer/consumer like. Does this make sense?	ask opinion	https://issues.apache.org/jira/browse/AIRFLOW-14
145	AIRFLOW-14	Yea, I agree with this. I guess the two things that I'm worried about are: # Will this mean we can't safely run more than one LocalExecutor? # Will this break the CeleryExecutor?	ask details	https://issues.apache.org/jira/browse/AIRFLOW-14

ANY23

#	Requirement	Question	Evaluation	URL
1	ANY23-297	GitHub user lewismc opened a pull request: https://github.com/apache/any23/pull/30 ANY23-297 Any23 doesn't build under JDK1.8 This issue addresses https://issues.apache.org/jira/browse/ANY23-297 @ansell can you please check how this affects your open pull request? Do you have any objections to merging this into master branch? Thanks You can merge this pull request into a Git repository by running: \$ git pull https://github.com/apache/any23/pull/30 Alternatively you can review and apply these changes as the patch at: https://github.com/apache/any23/pull/30.patch To close this pull request, make a commit to your master/trunk branch with (at least) the following in the commit message: This closes #30 ---- commit bb55685859a26b36ca4a9893bb93aa9eb7687b8c Author: Lewis John McGibbney <lewis.mcibbney@gmail.com> Date: 2016-11-25T21:33:57Z ANY23-297 Any23 doesn't build under JDK1.8 ----	ask opinion	https://issues.apache.org/jira/browse/ANY23-297
2	ANY23-297	GitHub user lewismc commented on the issue: https://github.com/apache/any23/pull/30 Hi @ansell any objections to merging into master branch? Seems the like the semargl issues are progressing which is great, however I don't think that should block us here in Any23. I'll commit by CoB today unless there are objections. Thanks	ask opinion	https://issues.apache.org/jira/browse/ANY23-297
3	ANY23-291	I agree [~thomas.francart], are you working off of master branch? If so, are you able to provide us with a patch and a test case if possible?	ask for patch	https://issues.apache.org/jira/browse/ANY23-291
4	ANY23-280	GitHub user lewismc commented on the pull request: https://github.com/apache/any23/pull/24#issuecomment-212761670 Anyone get a chance to take a look at this? This is THE critical issue for Any23 to address right now IMHO.	ask opinion	https://issues.apache.org/jira/browse/ANY23-280
5	ANY23-280	GitHub user ansell commented on the issue: https://github.com/apache/any23/pull/24 There are a large number of whitespace modifications to change from tabs to 2-space indentation. Is 2-space indentation what Any23 is aiming for, given that most java code is either tab or 4-space indentation.	ask opinion about spacing	https://issues.apache.org/jira/browse/ANY23-280
6	ANY23-276	GitHub user ansell commented on the issue: https://github.com/apache/any23/pull/27 I rebased this to fix the conflicts. There are some test failures, some of which I @ignore'd while other I left so far. Could you have a look in to see if you can figure out why the tests that have been ignored or still failing are broken?	ask opinion	https://issues.apache.org/jira/browse/ANY23-276
7	ANY23-276	GitHub user lewismc commented on the issue: https://github.com/apache/any23/pull/27 Hi @ansell whats the current state of this patch? Thanks for merging in the PR.	ask status	https://issues.apache.org/jira/browse/ANY23-276
8	ANY23-276	GitHub user lewismc commented on the issue: https://github.com/apache/any23/pull/27 Hi @ansell this is starting to look REALLY good. I've pulled your suggestions. I did notice that the appassembler is still generated in the core/target/appassembler. I think that this can be removed based upon the suggestion for CLI to be modularized as you've implemented. Is there anything else you want to add to this patch? If not then I suggest we merge into master.	ask status (in the form of patch needs more or not)	https://issues.apache.org/jira/browse/ANY23-276
9	ANY23-259	[~p_ansell] is this the issue which is leading to the unstable builds?	ask result of issue.	https://issues.apache.org/jira/browse/ANY23-259
10	ANY23-259	bq. There also appear to be a changed (lower, unless the junit assertEquals was applied backwards) number of extractors than expected. Hi [~p_ansell] what do you mean by this :(	ask meaning	https://issues.apache.org/jira/browse/ANY23-259
11	ANY23-257	I added implementations of Sesame RDFParser's for common OWL formats, including Manchester Syntax, Functional Syntax (using the obo to owl algorithm that is in owlapi), and OWL/XML to owlapi-rio:4. See the list at: https://github.com/owlls/owlapi/blob/version4/rio/src/main/resources/META-INF/services/org.openrdf.rio.RDFParserFactory As OWL isn't an input format, per se, which particular format were you wanting to target?	need description of targer	https://issues.apache.org/jira/browse/ANY23-257
12	ANY23-250	GitHub user lewismc commented on the pull request: https://github.com/apache/any23/pull/13#issuecomment-76448755 all tests pass locally. Would be appreciated if someone could confirm.	ask opinon	https://issues.apache.org/jira/browse/ANY23-250
13	ANY23-247	GitHub user ansell commented on a diff in the pull request: https://github.com/apache/any23/pull/17#discussion_r27437088 --- Diff: core/src/main/java/org/apache/any23/validator/rule/MissingItemscopeAttributeValueRule.java --- @@ -0,0 +1,52 @@ +/* + * Licensed to the Apache Software Foundation (ASF) under one or more + * contributor license agreements. See the NOTICE file distributed with + * this work for additional information regarding copyright ownership. + * The ASF licenses this file to You under the Apache License, Version 2.0 + * (the "License"); you may not use this file except in compliance with + * the License. You may obtain a copy of the License at + * http://www.apache.org/licenses/LICENSE-2.0 + * + * Unless required by applicable law or agreed to in writing, software + * distributed under the License is distributed on an "AS IS" BASIS, + * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. + * See the License for the specific language governing permissions and + * limitations under the License. + */ + package org.apache.any23.validator.rule; + import org.apache.any23.validator.DOMDocument; +import org.apache.any23.validator.Fix; +import org.apache.any23.validator.Rule; +import org.apache.any23.validator.RuleContext; + +/** + * This fixes missing attribute values for the 'itemscope' attribute, + * which was be associated with <div> nodes. + * Typically when such a snippet of XHTML is fed through the + * {@link org.apache.any23.extractor.rdfa.RDFA11Extractor}, and + * subsequently to Sesame's {@link org.semarglproject.sesame.rdf.rdfa.SesameRDFaParser}, + * it will result in the following behavior. + * <pre> + * {@code + * [Fatal Error] :23:15: Attribute name "itemscope" associated with an element type "div" must be followed by the '=' character. + * } + * </pre> + * This Fix is an effort to mitigate against that happening. + * + */ +public class MissingItemscopeAttributeValueRule implements Fix { --- End diff -- How is this class recognised or instantiated? META-INF/services/ or another method?	ask decription	https://issues.apache.org/jira/browse/ANY23-247
14	ANY23-247	Any ideas about this [~p_ansell]. Did you manage to debug where and when rules and/or fixes are invoked and when they are applied?	ask status	https://issues.apache.org/jira/browse/ANY23-247
15	ANY23-247	GitHub user lewismc commented on the pull request: https://github.com/apache/any23/pull/17#issuecomment-201537723 hi @ansell OK I've added in the correct rule and fix as well as a test to verify that empty Itemscope values are identified and fixed. Whilst debugging this however the core issue persists. Reasoning for this is that ""RDFA11Extractor extends BaseRDFExtractor"" which inherits the [parser function inputStream parameter](https://github.com/apache/any23/blob/63ba2fc82966cc056a2e475af849154d0dfdcf93/core/src/main/java/org/apache/any23/extractor/rdf/BaseRDFExtractor.java#L105). This input stream is not the 'fixed' steam but the raw document. The only way I can think around this is for us to + refactor the [RDFA1.1Extractor](https://github.com/apache/any23/blob/master/core/src/main/java/org/apache/any23/extractor/rdfa/RDFA11Extractor.java) such that it extends [TagSoupDomExtractor](https://github.com/apache/any23/blob/63ba2fc82966cc056a2e475af849154d0dfdcf93/api/src/main/java/org/apache/any23/extractor/Extractor.java#L60) as oppose to (eventually) the [ContentExtractor](https://github.com/apache/any23/blob/63ba2fc82966cc056a2e475af849154d0dfdcf93/api/src/main/java/org/apache/any23/extractor/Extractor.java#L44), or + undertake a mass refactoring which essentially removes the [ContentExtractor](https://github.com/apache/any23/blob/63ba2fc82966cc056a2e475af849154d0dfdcf93/api/src/main/java/org/apache/any23/extractor/Extractor.java#L44) altogether... this would provide us with a much more flexible and adaptable extraction framework IMHO. What do you think?	ask opinion	https://issues.apache.org/jira/browse/ANY23-247

#	Requirement	Question	Evaluation	URL
16	ANY23-247	Github user lewismc commented on the pull request: https://github.com/apache/any23/pull/17#issuecomment-201550510 I agree. Jumping through this in the debugged made me think the same. I think it is different if Any23 is to be a PURE implementation... But that is clearly not the case. Any23 fits in best when it can be used to extract semantics from any old crap input that it is fed. Parsers and extractors *should not* fail when there is a piece of crap input HTML. Currently, that's exactly what happens and it is extremely limiting. I would like to propose that this PR is committed to master as is, we then open a brand new issue which acts exactly your comments refactoring out content extractor and reusing the input stream which has been fixed, etc. Any thoughts Peter? Thanks fr quick response. On Friday, March 25, 2016, Peter Ansell <notifications@github.com> wrote: > The system does seem a little too complex for our purposes and isn't > usable because of that. > > Removing generics would be the first step IMO as there are too many > rawtypes definitions which indicate generics are being used badly. > > ContentExtractor may be able to be completely removed instead of being > refitted into the process after that and the parser should always be set to > parse as far as practical for our purposes. > > It is a little strange that there isn't a buffered, markable, InputStream > provided for all of the steps to reuse as necessary rather than pushing a > raw InputStream or other source into different extractors. > > — > You are receiving this because you authored the thread. > Reply to this email directly or view it on GitHub > < https://github.com/apache/any23/pull/17#issuecomment-201545776 > > -- *Lewis*	ask opinion	https://issues.apache.org/jira/browse/ANY23-247
17	ANY23-247	Github user lewismc commented on the pull request: https://github.com/apache/any23/pull/17#issuecomment-202702530 @ansell any further comments here? I will try to get to work on the larger issue this week.	ask opinion	https://issues.apache.org/jira/browse/ANY23-247
18	ANY23-238	Github user lewismc commented on the pull request: https://github.com/apache/any23/pull/9#issuecomment-55443359 Hi @Timpy thank you so much fo this PR. I absolutely agree with you that this is a problem as I do not like the look of the BNode hash either, it is useless for many purposes. For the record, can you please post here what the new outout will look like. Thank you p.s. I posted https://issues.apache.org/jira/browse/ANY23-238 so that we can add this to our release report once we release Any23 1.0	ask output	https://issues.apache.org/jira/browse/ANY23-238
19	ANY23-226	[~p_ansell], can you detail the current status of this over on jsonld-java? I feel that it would be best to bake in support in jsonld-java library which we wrap in Any23. Thanks in advance.	ask status	https://issues.apache.org/jira/browse/ANY23-226
20	ANY23-226	Github user lewismc commented on the pull request: https://github.com/apache/any23/pull/16#issuecomment-84384643 Ok Peter thank you for looking. This is great. I have not seen the test failures. Can you please tell me if it is in Any23 or in jsonld-Java? We could upgrade the Jsonld-Java implementation as well. To the 0.5.1 release On Saturday, March 21, 2015, Peter Ansell <notifications@github.com> wrote: > The main bug was that the entire script node was being sent to > JSONLD-Java, and not just its content. > > However, I also made a few other changes while doing that testing. > > It turned out that the jsonld was invalid, but somehow the exception when > parses fail was changed to be silently swallowed, so the only indication > was that the count was 0. I turned on the exception propagation again (no > reason it should be swallowed outside of temporary testing). > > However, in addition to the 4 tests currently failing on the core tests, > there are now other tests failing due to an inability to parse " >" > > — > Reply to this email directly or view it on GitHub > < https://github.com/apache/any23/pull/16#issuecomment-84257478 >. > -- *Lewis*	ask status	https://issues.apache.org/jira/browse/ANY23-226
21	ANY23-207	Hi [~cuibonobo] I am VERY sorry for getting back to you so late. Right now we have a Microformats2 GSoC project underway so we can definitely sync. Are you up for that?	ask if he wants to participate	https://issues.apache.org/jira/browse/ANY23-207
22	ANY23-207	Hi Folks, there is a PR for this now available at https://github.com/apache/any23/pull/18 Would be great if you guys could help us peer review the patch. Thanks.	ask for peer review	https://issues.apache.org/jira/browse/ANY23-207
23	ANY23-202	Do you want to do Tomcat monitoring? It would mean less code for us and more reuse of existing monitoring. We can automate this as well along with the nightly pull script. wdyt? I've classified this for 1.1 as it is certainly non-blocking as the WebService at any23.org will be updated nightly there is no need for us to hold a 1.0 release for this.	ask opinion	https://issues.apache.org/jira/browse/ANY23-202
24	ANY23-202	Does Apache have a privacy policy that details this collection of information by Google Analytics?	ask for details about implementation	https://issues.apache.org/jira/browse/ANY23-202
25	ANY23-202	[~p_ansell], I honestly do not know! I am not really too happy about Google Analytics being part of Any23 source or binary release. It means that everyone who deploys an Any23 server with the Google Analytics code embedded can be tracked! I am not really too happy with that. [~michele.mostarda], did you consider adding Tomcat logging metrics instead of Google ones? This would mean that it is down to whoever consumes Any23 to do their own log analysis, not for us to do it for them.	ask for details about implementation	https://issues.apache.org/jira/browse/ANY23-202
26	ANY23-202	Hi guys, I used Google Analytics because as you can easily verify other Apache projects already embed it within their landing pages (see for example http://hadoop.apache.org/). Any metrics provided by Tomcat is not comparable with the informations obtained by Google Analytics like geographic distribution and disambiguation of accesses. The tracker is embedded only in the demo front-end service page and anybody can customize (change the tracking ID) or remove it. Suggestions to make it less invasive: we could perhaps add the tracking ID within the POM so that it is more visible and left it empty in the distributions? we could add a disclaimer in the README? Suggestions? In general I suggest to keep a way to fine track user activity because we need to quantify the usage, the community and the distribution to better improve support.	ask opinion	https://issues.apache.org/jira/browse/ANY23-202
27	ANY23-202	The fact that some other projects have it doesn't absolve the privacy issue. There are many questions that should have been asked before the tracker was added. Who is using the information other than yourself and Lewis now or in the future? Was this a request from your employer or just from personal interest? Why didn't you add a privacy policy when you added the tracker? Why doesn't the entire project team have access to the information? I am not sure what you mean by adding information to the POM, but this is certainly something that is never done by any other projects.	ask for details about implementation	https://issues.apache.org/jira/browse/ANY23-202
28	ANY23-198	Hi Norbert. Are you using the Any23 source code? Are you able to contribute a patch for this feature? It would be a nice contribution.	ask for details	https://issues.apache.org/jira/browse/ANY23-198

#	Requirement	Question	Evaluation	URL
29	ANY23-189	what is the plan in this ticket? to integrate MO in any23 in some way?	ask for details about issue	https://issues.apache.org/jira/browse/ANY23-189
30	ANY23-185	hi [~niala12], nice patch :) Does it comply with the following ruleset "...Note that the global name has a specific semantic in the <meta> element and that the itemprop must not be set when one of the name, http-equiv or charset is already used."?	ask for details about implementation	https://issues.apache.org/jira/browse/ANY23-185
31	ANY23-185	Hi Lewis John McGibbney, One thing I have notice in the spec[1] itemprop is stated as "unimplemented", and I cant really find a reason behind this labeling. WDYT ? ----- if the itemprop is set, a user-defined metadata, transparent for the user-agent as the semantics of the metadata is user-specific. Living Standard Unimplemented ----- Also for the cases other cases name and http-quiv and charset (attribute type), this labeling has not been used and more information is given for these case where as for the itemprop case no any other information or sample case is provided. (see examples in the bottom of the the spec) So my patch is written considering these 3 cases except itemprop. I am more than happy to include this case of itemprop if more information is available. :) [1] https://developer.mozilla.org/en-US/docs/Web/HTML/Element/meta Regards Nisala	ask opinion	https://issues.apache.org/jira/browse/ANY23-185
32	ANY23-177	Hey [~lanthaler], do you have any interest in working on this issue? If so I can assign it to you.	interest in working?	https://issues.apache.org/jira/browse/ANY23-177
33	ANY23-177	[~lewismc], it's been a couple of years since I programmed in Java the last time and I'm not familiar with the code of Any23 either. [~p_ansell], could you look into this please? I think you are quite familiar with both code bases and already have a quite good idea of what would be needed. I don't how much work it would be, but it would perhaps also be interesting to extract JSON-LD embedded in HTML. See http://www.w3.org/TR/json-ld/#embedding-json-ld-in-html-documents Thanks!	look into this?	https://issues.apache.org/jira/browse/ANY23-177
34	ANY23-177	[~p_ansell] integrating this into the extractor/rdf package was pretty straight forward. Is there anything else you would want to see implemented here?	ask for details	https://issues.apache.org/jira/browse/ANY23-177
35	ANY23-155	Hi Debmalya, can you submit a patch we can apply against trunk? Thanks for opening this issue.	ask for patch	https://issues.apache.org/jira/browse/ANY23-155
36	ANY23-155	Hi, Thanks for your mail. I will feel happy to apply the patch and work for Any23. I am new to this project. Let me know how to apply patch. Is there any link to follow? Please give me some time. Is there any date line for this? Best Regards, Debmalya Jash On Tue, Apr 2, 2013 at 3:47 AM, Lewis John McGibbney (JIRA) <jira@apache.org	ask timeline	https://issues.apache.org/jira/browse/ANY23-155
37	ANY23-155	The only test that is failing for me right now in the core module is testBlankNodesViaURL, so the RDFSchemaUtilsTest is working for me. Could you post your java and maven versions (mvn -version and java -version) here	ask for source code	https://issues.apache.org/jira/browse/ANY23-155
38	ANY23-155	Could you also post the output of: core/target/surefire-reports/org.apache.any23.util.DiscoveryUtilsTest.txt After you revert your changes to RDFSchemaUtilsTest and post the output of: core/target/surefire-reports/org.apache.any23.vocab.RDFSchemaUtilsTest.txt	ask for output	https://issues.apache.org/jira/browse/ANY23-155
39	ANY23-155	So RDFSchemaUtilsTest is working now for you after you reverted your changes to it?	ask for status on bug	https://issues.apache.org/jira/browse/ANY23-155
40	ANY23-147	http://purl.org/dc/terms/ doesn't look to be deprecated. Are you thinking of http://purl.org/dc/elements/1.1/ ?	ask for clarification	https://issues.apache.org/jira/browse/ANY23-147
41	ANY23-147	Well its a redirect right? Which URI so we use as the NS?	ask details	https://issues.apache.org/jira/browse/ANY23-147
42	ANY23-147	I am not sure what you mean. Everyone I know is still using http://purl.org/dc/terms/ for current applications, although older applications may still support http://purl.org/dc/elements/1.1/ . The DCTERMS vocabulary in Any23 has a small subset of the http://purl.org/dc/terms/ properties, but it could definitely be improved, particularly to include the classes. Most people refer to the dc/elements/1.1 namespace using the prefix "dc", with the newer dc/terms either being "dcterms" or "dct". Is that part of the confusion?	refinement to avoid confusion	https://issues.apache.org/jira/browse/ANY23-147
43	ANY23-144	What do you suggest then Peter? I am happy to go with a positive consensus on this one.	ask opinion	https://issues.apache.org/jira/browse/ANY23-144
44	ANY23-144	{bq}My preference is for all uppercase vocabulary class names if the commonly used prefix is an acronym, and camel-case where it is not, however, in some cases that distinction is not necessarily clear.{bq} OK so in an attempt to move towards a more appropriate (partial) solution... the DOAP vocab would remain named as DOAP and LKIF-Core would be changed to CamelCase e.g. LKIFCoreAction, etc... Can you confirm? It appears that this was not so much of a pressing issue before Simo pointed it out and you commented. This being said I appreciate that and I think it is important to get it right :) {bq}To clear up the naming I would prefer if we avoided the getInstance() pattern.... like access to the properties without storing them in a local private field to get around the getInstance() call.{bq} Excuse me if I am wrong, but I see this as a fundamental different issue from renaming of classes. Can we address this more directly in a related but separate issue?	ask for another issue	https://issues.apache.org/jira/browse/ANY23-144
45	ANY23-144	patch for trunk which addresses comments and concerns. If someone can review we can sort this one out. Thank you	ask for review	https://issues.apache.org/jira/browse/ANY23-144
46	ANY23-144	I can't seem to apply either of the patches using "patch -p0 < ANY23-144.patch" or similar for ANY23-144-v2.patch. Do you need to apply both of the patches to get it to work, because that wasn't working for me either but I may just be doing something wrong? What command have you been using to apply the patches?	how to use implemetation	https://issues.apache.org/jira/browse/ANY23-144
47	ANY23-141	Expected release date of Sesame 2.7.0 has been set for about 3 weeks time. Is it worth waiting until then to release 0.8.0 or should we still focus on 0.9.0?	ask for target release	https://issues.apache.org/jira/browse/ANY23-141
48	ANY23-141	Sesame-2.7.0 is now available, and the Git repository is now writable. What should we do about this issue?	ask opinion	https://issues.apache.org/jira/browse/ANY23-141
49	ANY23-138	Hi [~simone.tripodi], can you provide some guidance as to this one? Do you have any 'typical' steps or guidance which we could work on? Thanks	ask opinion	https://issues.apache.org/jira/browse/ANY23-138
50	ANY23-137	Hi Lev. You've marked this issue as being related to a number of other issues. Can you please provide info on how many of these have been addressed within the scope of your proposal? Thank you.	ask refinement on connected issues	https://issues.apache.org/jira/browse/ANY23-137
51	ANY23-137	I think it would be great to push on with this RDFa proposal then make an interim release for Any23. Is anyone interested in working to get this done? This issue (with a few others) would IMO justify pushing an RC for 0.9.0. Any thought please?	ask for thoughts	https://issues.apache.org/jira/browse/ANY23-137

#	Requirement	Question	Evaluation	URL
52	ANY23-137	Is there still any glue code that needs to be written to achieve this? or is ready for testing? Are there any instructions somewhere on how to build any23 with this patch so I can test it?	ask status	https://issues.apache.org/jira/browse/ANY23-137
53	ANY23-137	What's the status of this issue to merge semargl into Any23? Is there anything I can test to make progress?	ask issue	https://issues.apache.org/jira/browse/ANY23-137
54	ANY23-137	Hi! I've prepared pull request to fix RDFa related tests. See https://github.com/apache/any23/pull/2 . Are there any other things to do with this branch?	Ask versions	https://issues.apache.org/jira/browse/ANY23-137
55	ANY23-137	Success. I tested https://github.com/apache/any23/pull/2 on ubuntu 13.04 and it worked with java 1.7.0_51 and mvn 3.0.4! Not sure why it's not working on Mac OS X, do I need java 7?	ask about testing requirements	https://issues.apache.org/jira/browse/ANY23-137
56	ANY23-137	Thanks, Stephane! Completely missed that RDFa was used as a part of extraction process in other tests. I've added related fixes. Brief description. *ServletTest* Old RDFa implementation produces {{<issue level="Warning" row="14" col="5">Error while processing node /HTML(1)/HEAD(1)/META(9) : 'Cannot map prefix 'fb'</issue>}} while {{<fb:app_id>}} is completely valid predicate which shouldn't be resolved against fb: prefix. *Any23Test* *RoverTest* Changed RDFXMLWriter to NTriplesWriter in some tests to improve precision (they basically check line count). Changed expected triples count. It was reduced in most cases, because old RDFa parsed produced a lot of invalid triples like: {quote} <http://host.com/service> <http://host.com/serviceexternal> <http://host.com/service/ambiente/> . <http://host.com/service> <http://host.com/serviceexternal> <http://host.com/service/salute/> . <http://host.com/service> <http://host.com/serviceexternal> <http://host.com/service/legalita/> . <http://host.com/service> <http://host.com/serviceexternal> <http://www.ansamed.info/> . <http://host.com/service> <http://host.com/serviceexternal> <http://host.com/service/web/notizie/regioni/lazio/provinciadirioma/> . {quote} Fixed markup in {{test-resources/src/test/resources/html/rdfa/ansa_2010-02-26_12645863.html}} to conform declared XHTML 1.0 Strict. Fixed RDFa markup in {{test-resources/src/test/resources/html/encoding-test.html}} otherwise it shouldn't produce any triples. Disabled second part of {{Any23Test.testExtractionParameters}}. Should it do anything after RDFa parser replacement? Also, ExtractionException thrown from BaseRDFExtractor is escalated in test suite. It leads to some failed tests in Any23Test. What's the correct behaviour for ANY23 parser in case it gets SAXException?	ask for refinement	https://issues.apache.org/jira/browse/ANY23-137
57	ANY23-134	Will you have time to work on this before the 0.8.0 release or is it okay to postpone it to 0.9.0?	ask timeline	https://issues.apache.org/jira/browse/ANY23-134
58	ANY23-127	[~lewismc] I'm failing to understand the point of this story. The purpose of the Semantic Web is that parsing tools like any23 should be vocabulary agnostic. Baking support for a whole load of vocab is by design out dated as soon as it ships, given that vocabularies tend to evolve (e.g. schema.org). Maybe if you explain what you are trying to achieve I will be better able to help...	ask for description	https://issues.apache.org/jira/browse/ANY23-127
59	ANY23-127	.bq The purpose of the Semantic Web is that parsing tools like any23 should be vocabulary agnostic. Yes of course. I take your point. However if we are not namespace (or more appropriately vocabulary) aware then do we just _expect_ to recognize some namespace declarations? .bq ...is by design out dated as soon as it ships, given that vocabularies tend to evolve (e.g. schema.org). I need to disagree here. I would agree if you had added "given that _some_ vocabularies..." Some vocab's do evolve you are right. You've picked probably the most dynamic of all as your example however this ignores the point that _some_ are not as dynamic. For example, I added some Legal vocab's for a specific project I was working on. These may or may not be useful to other people. I also know that the vocabulary has NOT changed, therefore they are as useful now as they were when I wrote the implementation. If you navigate to the BBC vocabularies you will also see that these are well established and are not as dynamic as you may indicate. Furthermore, by implementing other developers' vocabularies we are directly inviting those same developers to use Any23 as a tool for extracting triples from their markup. Personally I've reached out the DOAP and LKIF communities to state that we are now namespace aware. This is the wider goal. I hope this provides some justification behind my reasoning on this one.	ask details	https://issues.apache.org/jira/browse/ANY23-127
60	ANY23-121	Has this got anything to do with Any23?	ask for details	https://issues.apache.org/jira/browse/ANY23-121
61	ANY23-110	I honestly didn't realise that this was committed. Have you verified it works Peter? I will try tomorrow	ask for test	https://issues.apache.org/jira/browse/ANY23-110
62	ANY23-105	One fix for this may be to create a Sesame Sail implementation that can talk to TDB. Has anyone looked into this before?	ask expert opinion	https://issues.apache.org/jira/browse/ANY23-105
63	ANY23-105	[~andy.seaborne] [~rvesse] [~p_ansell] what do you think is the best implementation here? I am _thinking_ towards a more abstract interface which wouldn't just restrict us to TDB. Is there anything out there which I could use?	ask opinion	https://issues.apache.org/jira/browse/ANY23-105
64	ANY23-103	So this has been addressed right? It doesn't quite look like what you implemented in your commit @ r1371537 is what you initially proposed, but all seems well anyway.	ask status	https://issues.apache.org/jira/browse/ANY23-103
65	ANY23-83	Having had a browse of both o.a.a.vocab.RDFSchemaUtils, and o.a.a.rdf.RDFUtils respectively, I understand where you are coming from. I also notice that you propose a more generic factory type implementation similar to Tika which (based on some criteria) guesses the specific parser which should be used per situation. What other mimetypes would we encounter other than the ones mentioned in both of the above classes as specified within the hardcoded RDFParsers? I am interested to find out what you think. Thank you	ask opinion	https://issues.apache.org/jira/browse/ANY23-83
66	ANY23-83	Hi Peter, {bq} I am not familiar enough with the rest of the codebase to know why the shortlist of supported parsers is in place {bq} Although I have tried to get to grips with as much of the codebase as I have a use-case for, I am sorry as I cannot provide insight into the original design considerations which you refer to... maybe one of the other guys could chime in here and clear the picture for us? I must admit though, I am in agreement with your overall analysis and think that the more extensible we can make the Any23 libraries the better. There was some discussion however over in ANY23-19 where the vision was to have an abstraction layer between Any23 and the underlying RDF API's. In my opinion, this would provide far greater potential for us to make Any23 a more flexible library... wdyt?	ask opinion	https://issues.apache.org/jira/browse/ANY23-83

#	Requirement	Question	Evaluation	URL
67	ANY23-83	I like these two patches and based on the conversation on any23-dev give the +1. Marking for 0.7.0 and happy to commit on behalf of Peter unless there are objections?	ask if req done? ask if objections to commit	https://issues.apache.org/jira/browse/ANY23-83
68	ANY23-83	I rebased my git repository onto a copy of the current trunk to check how many patches I still had in this area and noticed that RDFSchemaUtils still hardcodes the supported formats. I may not be seeing the significance of VocabularyFormat currently, but I assume that it could be generalised to any RDFFormat. Is it related to the ability to type in a human readable alias for the format in VocabPrinter? If so, have you tried using RDFFormat.valueOf(String) in the converter inside VocabPrinter which matches against registered formats using: format.getName().equalsIgnoreCase(formatName) These names may not exactly match the historical names for each format that were previously used, but they are extensible at runtime so it makes RDFSchemaUtils easier to maintain. This may also help with ANY23-3 as N3 would be available, although it would just be the Turtle subset in practice as we only hold RDF triples in memory.	ask dependancies?	https://issues.apache.org/jira/browse/ANY23-83
69	ANY23-75	Setting for 0.7.0-incubating. Is it possible for you to explain a bit about the patch and the underlying reason as to why the existing parser implementation seems to clog up? Also this is really trivial but can you please have a look at the coding format if it differs or not. My initial thoughts of the patch are great, it's a nice one to have caught, but some additional explanation would really help us out. Thank you very much. Lewis	ask for clarification about patch	https://issues.apache.org/jira/browse/ANY23-75
70	ANY23-70	Yep, unfortunately, as Lewis put it they are "horrible for searching". I usually end up doing searches elsewhere. Is there a way to search on the Apache mail archive?	ask archive search	https://issues.apache.org/jira/browse/ANY23-70
71	ANY23-69	Does this duplicate ANY23-7 ?	ask opinion	https://issues.apache.org/jira/browse/ANY23-69
72	ANY23-69	Yeah OK thanks for clarification. The W3C guys that chimed in on the thread on any23-dev, stated that they would like us to produce this by the end of the month so we have plenty of time to cook something up. I had a look into the RDFa 1.1. spec, and subsequently into the EARL spec... unsurprisingly it's left me a bit puzzled as to how to pull this together so I was going to start a thread on our dev list. However this is as good a place as any I suppose. Have you any initial thoughts Michele?	ask opinions	https://issues.apache.org/jira/browse/ANY23-69
73	ANY23-65	Well, last night I found that no, it isn't working flawlessly. I was just looking at the whole thing again to trace what is going wrong. I was about to send my XSLT file to the Sindice-dev mailing list, to bridge the period between now and the moment Sindice starts using Any23 0.7.0, when I thought of an edge case to test against. My stylesheet doesn't handle it well, and neither does my build from the SVN code. The case: when a prefix attribute contains "rnews:http://www.iptc.org/std/rNews/1.0: foaf:http://xmlns.com/foaf/0.1/dbpedia:http://dbpedia.org/resource/" (no space between prefix name and prefix URI, but at least one prefix URI ending with a colon), my stylesheet wrongly assumes there are spaces between prefix name and prefix URI, because it tests whether the attribute contains ": ". The RDFa11Parser outputs warnings that it cannot map the foaf prefix when I test it on my 0.7.0 build. It is an edge test case because it is invalid content for a prefix attribute, but since I saw that Any23 is accepting / testing against no-space prefix definitions and I think namespace URIs with a colon are valid (e.g. <http://dbpedia.org/resource/Category:>), I figured it made an interesting test case. I sent the file anyway, with this test result too :) Another issue I have is that the RDFa11Parser doesn't infer the right triples from <link> elements with a @rel in the <head> section. I believe the @rel values "icon", "stylesheet", "bookmark" etc are to be treated specially. Sindice (Any23 0.6.1) produces URIs for these like <http://www.w3.org/1999/xhtml/vocab#icon> from my blog post. When I extract from my blog post locally, I get properties like <http://ben.companjen.name/2011/08/het-gezin-timmer-de-bruijn-in-amsterdam/icon>. The parser (also) complains my HTML doesn't declare 'xmlns="http://www.w3.org/1999/xhtml"'". I couldn't find the part of the HTML or RDFa specifications that says it should do so - as far as I can tell it's not necessary in HTML5. Looking at the code, I see some references to RDFa 1.0 in the comments in the processDocument method. This method seems to be the source for the complaint. Maybe the problem and complaint (well, warning actually) are linked? And maybe the incorrect handling of the @rel values is also linked to "// TODO: introduce support for RDFa profiles. (http://www.w3.org/TR/rdfa-core/#s_profiles)"? BTW, these profiles don't exist (anymore) in RDFa Core. I hope someone can shed more light on this, I'm too confused from reading all the RDFa related docs and drafts right now ;)	ask opinion	https://issues.apache.org/jira/browse/ANY23-65
74	ANY23-61	The repo directory is a maven repository. Are the artifacts in this directory in another maven repository? We are currently deploying snapshots to the Apache Snapshots repository and new releases to the Apache staging server, based on the definitions in the apache-10.pom parent. http://svn.apache.org/repos/asf/maven/pom/tags/apache-10/pom.xml	ask for details	https://issues.apache.org/jira/browse/ANY23-61
75	ANY23-50	i had a look to the patch and it looks very nice. If the META-INF data is built at compile time and it is not necessary to package the plugins to annotate them I don't have any other concern about applying this simplification. How does it work /does it work the hot-loading with ServiceLoader?	ask details about implementation	https://issues.apache.org/jira/browse/ANY23-50
76	ANY23-50	Hola! META-INF/services metadata are always generated at compile time, indeed plugins are annotated with [@MetalnService http://weblogs.java.net/blog/kohsuke/archive/2009/03/my_project_of_t.html] that is part of a provided-scope Annotation Processor wrote by Kohsuke (the Hudson/Jenkin's father) In Any23 the ServiceLoader is used inside the Any23PluginManager, relying on the already managed dynamic ClassLoader, so when adding new jars/resources at runtime to the ClassLoader, SL queries it. I think that the same strategy can be applied to Extractors' factories as well, WDYTY? Anyway, please double-ack if you agree I can commit the modification! TIA!	ask opinion	https://issues.apache.org/jira/browse/ANY23-50

#	Requirement	Question	Evaluation	URL
77	ANY23-50	I think this should definitely be committed? Once question, is there anything specific about the Any23 plugin system which make this issue only applicable to Any23? What I mean, is would it be possible to apply this issue in a project elsewhere which relies heavily on plugins? (e.g. Nutch)	ask details about propect	https://issues.apache.org/jira/browse/ANY23-50
78	ANY23-48	OK so Committed @ revision 1303963 in site I also committed some changes to trunk but I can't remember which revision. If someone could please review where we are with this one, unless someone can point me to further improvement (which obviously I'm very open to) then I'd like to close this one off so we can focus on some of the other issues which lie between us and the 0.7.0-incubating release. Thanks	ask status	https://issues.apache.org/jira/browse/ANY23-48
79	ANY23-46	Thanks Andy. Taking this forward, obviously we can copy in Any23-dev@ anyway, but do we wish to combine these resources requests to infra@, e.g. is there any way that we can share resources to limit the infra@ involvement if you were wanting to bring the SPARQL stuff to ASF as well?	ask for sharing resouces	https://issues.apache.org/jira/browse/ANY23-46
80	ANY23-38	@Lewis, you mean org.apache.any23.util.LogUtils or what? That one is used to activate / deactivate the verbose mode of CLI tools by changing the log level programmatically.	ask for clarification	https://issues.apache.org/jira/browse/ANY23-38
81	ANY23-38	Upon closer inspection looks like your correct Michele. What about we re-factor this issue to address Simone's opinions? I'm a +1 for Slf4j	ask for refinement	https://issues.apache.org/jira/browse/ANY23-38
82	ANY23-36	I genuinely have not got a clue what has to be done to resolve this one. I've tried using the maven changes plugin but getting nowhere. Anyone else got a scooby?	Ask opinion. I think scooby = clue	https://issues.apache.org/jira/browse/ANY23-36
83	ANY23-36	Isn't this issue related to Any23 history munging?	ask for clarification	https://issues.apache.org/jira/browse/ANY23-36
84	ANY23-33	I attempted to replace the classes with public implementations, but the once found within the classpath are marked as @deprecated. Ideas? {code} -import com.sun.org.apache.xml.internal.serialize.OutputFormat; -import com.sun.org.apache.xml.internal.serialize.XMLSerializer; +import org.apache.xml.serialize.OutputFormat; +import org.apache.xml.serialize.XMLSerializer; {code}	ask for ideas	https://issues.apache.org/jira/browse/ANY23-33
85	ANY23-33	Are you getting this from your compiler output? I've not tried to implement yet. We don't explicitly depend upon any xerces related stuff in the parent pom.xml, so I am curious where the imports are being pulled from for you to get this output.		https://issues.apache.org/jira/browse/ANY23-33
86	ANY23-32	Thanks Mic for reporting! I need indeed doing more work on the service, my bad that I didn't notice! there is a specialized README in the produces distribution package that explains how to install&run, can you double chek please? TIA!	ask for double check	https://issues.apache.org/jira/browse/ANY23-32
87	ANY23-26	Hi Peter. Regarding your edited comment above. Have you started/finished work on the Extractor factory patterns? I would like to get this issue integrated into trunk... therefore I need to test it. If the answer to the above is yes, are you able to produce a patch which I can apply to the SVN trunk HEAD? Thank you	ask status	https://issues.apache.org/jira/browse/ANY23-26
88	ANY23-26	Has anyone else had the opportunity to test Peters patches? There are some trivial issues which we have discussed off list but generally speaking the spi and tika upgrades look dynamite to me. I also confirm that all tests pass locally for these patches.	ask if patches were tested	https://issues.apache.org/jira/browse/ANY23-26
89	ANY23-21	Is there an automated tool for doing this?	CONTEXT DEPENDANT. ask for resources	https://issues.apache.org/jira/browse/ANY23-21
90	ANY23-19	Hey Paolo, with regards to Chris' recent comments on the dev list and the stub report he filed, how did you get on wrt moving away from the packages you've highlighted above and also with progressing with more collaboration between ASF RDF-type projects? I'm using this issue as a basis to kick the conervation off and hopefully others will join in.	ask expert opinion	https://issues.apache.org/jira/browse/ANY23-19
91	ANY23-19	Hi Max, the thing is: adding yet another dependency is going to make Any23 bigger and more complex. Can something simple be done without adding yet another dependency? As I said, an example of a small RDF distiller which does this is java-rdfa (https://github.com/shellac/java-rdfa). Third party RDF APIs become only an optional dependency and people are free to choose which one they want/need simply adding that to their classpath (in a very similar way to what SLF4j does for logging frameworks). The only thing tricky is "validation", but Any23 could chose not to validate URLs/IRIs, etc. and simply relay on the third party RDF APIs to do that.	ask for alternate solution	https://issues.apache.org/jira/browse/ANY23-19
92	ANY23-19	Hi Paolo, The example library that is continually being referred to here, java-rdfa, abstracts away from clerezza, sesame and jena interfaces by representing everything as strings. See https://github.com/shellac/java-rdfa/blob/master/core/src/main/java/net/rootdev/javardfa/StatementSink.java The java-rdfa library is also not being developed or used actively, so it isn't the best example of a successful library that doesn't use a type-safe RDF Statement/Value API. I worked on java-rdfa a little but there is no push behind it. Even its author refers to it as "The cruftest RDFa parser in the world" Two of the three goals of Any23, command line utility, and web service, are completely ambivalent to the technology being used internally, as long as it is high quality, and the Sesame libraries are very high quality in my opinion and experience. The only goal that would be affected would be the use of Any23 as a library. How often is Any23 currently being used as a library? Could another project easily implement the same functionality using Jena with less effort than it would take to either create a custom string based solution, or move to another framework that may have just as many or more dependencies? In terms of the use of Any23 as a library, is there anything about the Sesame Model hierarchy (Value/Resource/Literal/BNode/URI) that would be better represented using a custom solution? As one example, I have been working with OWLAPI recently and its RDF handling is shocking, it merges URIs with Blank Nodes to form what it refers to as IRIs. It has a custom internal solution that only recognises two types of triples, those with an IRI in the object position (where IRI is not type-safely defined between URI and BlankNode) and those with Literals in the object position. I can't imagine Any23 going down this route, but it is the worst case scenario if the API is converted without a reason. In the simplest scenario, it may be possible to reuse the Sesame Model hierarchy to produce Values that work across all three libraries, using custom ValueImpl etc., implementations that actually implement the relevant interfaces from other libraries, along a custom ValueFactory to produce these multi-library-compatible Values (custom ValueFactories can be plugged into any Rio Parser using Rio.getParser(RDFFormat, ValueFactory), a functionality which I haven't seen in other libraries). In terms of the actual packages that are currently used, there are four basic packages sesame-model, sesame-rio-api, sesame-repository-api, sesame-sail-api, sesame-sail-memory. These base libraries are small dependencies. One other dependency is some small utilities, sesame-util that are used by sesame-model and other sesame libraries. 82K - sesame-model-2.6.5.jar 36K - sesame-repository-api-2.6.5.jar 22K - sesame-rio-api-2.6.5.jar 56K - sesame-sail-api-2.6.5.jar 54K - sesame-sail-memory-2.6.5.jar 53K - sesame-util-2.6.5.jar The value Impl classes should not be directly referenced. They should be accessed using a ValueFactory and used as their Interfaces. This doesn't change any of the libraries that are used, but it is better practice. The other libraries that pull in the Rio parsers can be linked in dynamically without compiling in the dependency, so the use of Any23 as a library would enable people to pull them in as needed. See Rio.getParser(RDFFormat, ValueFactory) and Rio.getWriter(RDFFormat) methods. It would be valuable if Any23 could dynamically pull in all of its parsers and writers using the Rio.* static methods. Then it could be used with the absolute minimum number of parsers and writers for the current user. 4.6K - sesame-rio-n3-2.6.5.jar 14K - sesame-rio-ntriples-2.6.5.jar 33K - sesame-rio-rdfxml-2.6.5.jar 17K - sesame-rio-turtle-2.6.5.jar Switching to another library may cause the bloat that you say you do not want. For example, Jena and its immediate dependencies is quite large compared to the modular sesame jar files, and that doesn't include the SPARQL parsing libraries from ARQ, as indeed the sesame libraries quote above do not include the sparql libraries. 1.7M - jena-core-2.7.0-incubating.jar 151K - jena-iri-0.9.0-incubating.jar 3.1M - icu4j-3.4.4.jar 1.4M - xercesImpl-2.10.0.jar	ask dependencies	https://issues.apache.org/jira/browse/ANY23-19

#	Requirement	Question	Evaluation	URL
93	ANY23-18	Ciao Michele, what are the problems you found with java-rdfa? There is an open issue for java-rdfa to add RDFa 1.1 support (https://github.com/shellac/java-rdfa/issues/24). In any case, good to see the XSLT replaced by a new programmatic RDFa extractor.	ask for details	https://issues.apache.org/jira/browse/ANY23-18
94	ANY23-11	I've updated the w3c page above, but can't seem to create an account for the RDFizers!!! Is there anything else we wish to add to this issue? I think it goes to say that at any given opportunity make people aware of Any23.	any other improvements?	https://issues.apache.org/jira/browse/ANY23-11
95	ANY23-4	I think this certainly duplicates ANY23-69 ?	ask opinion about related issues	https://issues.apache.org/jira/browse/ANY23-4
96	ANY23-3	Is this not addressed?	ask status?	https://issues.apache.org/jira/browse/ANY23-3
97	ANY23-3	Are there any full N3 implementations in Java? Sesame and Jena both parse the RDF subset, but not the full N3 grammar. In practice, the RDF subset amounts to Turtle. Is Any23 interested in the non-RDF components of N3?	ask for details about current implemation and future work	https://issues.apache.org/jira/browse/ANY23-3
98	ANY23-2	@Peter, Are you interested in committing this in line with thr new SPI Factory implementation you wrote?	ask details about implemetation?	https://issues.apache.org/jira/browse/ANY23-2

DASHBUILDER

#	Requirement	Question	Justification	URL
1	DASHBUILDE-237	Regarding individual chart filters (see example => !chartfilter.png thumbnail!) I think both "master" toolbar and individual chart filters can co-exist. A toolbar at the top of the dashboard screen could be implemented as another displayer component which allows for filtering on any of the properties of a given data set. Dashboard designer users can decide which approach to take. For toolbar, a toolbar filter component is dropped at the top of the page and filtering is disabled for every single chart. [~uxdlc] [~mstevens1] wdyt?	ask opinion	https://issues.jboss.org/browse/DASHBUILDE-237
2	DASHBUILDE-237	[~davidredhat] I see these filter tags used mostly with big data, faceted/complex filtering. I believe the use case there is that there is so much criteria (and objects) that you're filtering for you really need this kind of visual summary. But for individual chart filters I'm not sure the use case exists, perhaps a slightly more complex filter control would be enough, like a multi-select dropdown: http://dotansimha.github.io/angularjs-dropdown-multiselect/docs/#/main Wdyt?	ask opinion	https://issues.jboss.org/browse/DASHBUILDE-237
3	DASHBUILDE-237	There are some cases where individual chart filters control might make sense. See attached exampleB & exampleC. These are simple scenarios where we do not require a full featured filter toolbar as they are quite simple uses cases. A filter toolbar is more suitable for full dashboard scenarios, see exampleA. Anyway, users can always choose whether to use individual chart filters or a global toolbar. [~uxdlc] We do already have selector dropdowns for single column filtering, unlike the filter toolbar which is multi-column. [~mstevens1] 1) As per each icon's tooltip: "Export to CSV", "Export to Excel", "Show/hide columns" 2) +1. The string is a hard coded literal 3) What are you trying to build? Can you elaborate on this please?	ask refinement	https://issues.jboss.org/browse/DASHBUILDE-237
4	DASHBUILDE-237	I seem to be losing the thread here a bit. A quick conversation might help. Some stray questions (none of which are directly related to the project Metrics page anymore)... 1) A filter at the top of the screen (a global filter) - does that data column need to be present in every chart on the screen? I can only globally filter by _time_, for example, if every chart has a time value. 2) When filtering one chart on the screen (local filter), should you also filter other charts that have the same data column? I filter one chart by "David" and the other charts with people update? 3) If answering yes to #2, should you then show the "tags" on the other charts to make it clear that the chart is filtered? 4) Have I gathered correctly that Dashbuilder uses NVD3? 5) Does the user get to choose whether the chart displays a legend? And, further, whether that legend is clickable / filters the chart? 6) Are all of these things customizable by the user anyway and my questions are moot?	ask questions about use of filters	https://issues.jboss.org/browse/DASHBUILDE-237
5	DASHBUILDE-237	Hello [~davidredhat] do you have time to chat on Thursday or Friday? [~uxdlc] Thank you, Matthew	need to talk offline about requirement	https://issues.jboss.org/browse/DASHBUILDE-237
6	DASHBUILDE-237	The following screenshot shows the toolbar implemented in RedHat's BPMS => !bpms_toolbar.png thumbnail! The first dropdown ("Id") shows a set of predefined columns that support text based search. The rest of the dropdowns ("State", "Errors", "Name" and "Last update") are extra filters for specific columns. They have been expanded in the toolbar to make it easier to operate with them as well as to provide custom filter controls depending on the column type. Notice, for example, the "Last update" column is an advanced control for selecting a date range. In order to support the features above the following requirements must be also considered: 1.- The toolbar requires a list of columns to display in the generic column dropdown 2.- It also requires a sublist of the columns to be expanded in the toolbar 3.- Each expanded column may optionally have a divider element. This enables the concept of column groups in the expanded section The main goal is to provide a generic toolbar able to support different toolbar configurations, including the one used in BPMS. Regarding the discussion in previous comments about having individual chart filters. The best option is to remove them and rely only in either the toolbar component or the existing selectors components (dropdown, label set and slider). Whenever a filter is required in a dashboard, a filtering component needs to be dropped into the page. Charts will no longer provide filtering functionality. That will not only reduce the codebase but it will also keeps the filtering approach, Patternfly and the UXD on the same page. [~uxdlc] [~mstevens1] wdyt?	ask opinion	https://issues.jboss.org/browse/DASHBUILDE-237
7	DASHBUILDE-232	Hi Team, it would be a great help if a workaround is provided to have default ascending sort of category. Currently, there is no sorting and category axis appears in random order,	asks for workaround	https://issues.jboss.org/browse/DASHBUILDE-232
8	DASHBUILDE-230	This is the old implementation => !^dashbuilder_filter_labels.mp4 This is how it looks the new version => !^dashbuilder_filter_labels_new.mp4 [~uxdlc] wdyt? :-)	ask opinion	https://issues.jboss.org/browse/DASHBUILDE-230
9	DASHBUILDE-230	[~davidredhat] I sent my comments to a design related jira, next time will just post them here directly. The visual styling of the labels themselves looks more like PF, but some of the interaction aspects aren't aligned. * For future iterations/enhancements: The PF filters are typically used in conjunction with a global [toolbar control http://www.patternfly.org/pattern-library/forms-and-controls/toolbar/]. It's presented at the top of the screen (or content view) to indicate that everything on that screen is being filtered. The filter labels are shown there to help users to know at a glance which filters have been applied. I haven't seen an example where they are applied out of that context. * Usability concerns with the current dashboard presentation are: _Because the labels are visually associated with one widget, an assumption will be made that the filters have only be applied to that widget. _Users need to hunt around the screen to find the labels to figure out which filters have been applied, whereas in the PF toolbar design they are all presented in one location at the top of the screen. To align the interaction with PF (in addition to the styling) could we look into the options for picking up the toolbar? Of it that can't be used for some reason, perhaps we could look into some alternative design solutions. Let me know what you think, thanks!	ask for opinion	https://issues.jboss.org/browse/DASHBUILDE-230
10	DASHBUILDE-229	Initial implementation done. Preview video => !^dashbuilder_perspective_component.mp4 As the video shows, there are issues with the layout rows. The logic behind the calculation of the row height seems not to be working very well. [~ederign] any ideas?	Ask for ideas	https://issues.jboss.org/browse/DASHBUILDE-229
11	DASHBUILDE-228	[~davidredhat] or [~porcellini], can you guys please put this on Sprint5?	ask timeline	https://issues.jboss.org/browse/DASHBUILDE-228
12	DASHBUILDE-226	This issue is a bit tricky. The purpose of having a fixed width for a label set is mostly to expand its size (see !selectorlabelswidth.png thumbnail!). Even in this example there might be cases where one item's name is still too large. Also, it does not make sense to set a very small width since the item's name won't fit in the button. So, what's the solution? I think a good one could be as follows: 1.- Expand If the total width is larger than the selector width then all the items shall be displayed without cuts and the remaining width should be split and added up evenly to every item. The total width is the sum of all items width whereas an item's width is calculated as a function based on its number of chars. For example, if we have the following items: "Red", "Green" and "Blue", the total length is 12 chars and the total width 120px (provided we set 1 char=10px). If, for instance, the selector width is set to 210 then the width of each item will be: 210-120=90px excess width => we split evenly 90/3=30px => Red=3chars*10+30=60px, Green=5chars*10+30=80px, Blue=4chars*10+30=70px As a result, every button is expanded 30px. 2.- Reduce If the width set is smaller than the total width then we need to reduce each item so that they fit into the button reduced space. Given the example above and a selector width=90. then 90-120=-30px to reduce overall, -10px each item => Red=30-10=20px, Green=50-10=40px, Blue=40-10=30px. The buttons will show as "Re", "Gree", "Blu" with their corresponding tooltip showing the full name. [~jhrcek] wdyt?	ask opinion	https://issues.jboss.org/browse/DASHBUILDE-226
13	DASHBUILDE-218	[~uxdlc], [~jhrcek], the current implementation is ready to push. I've implemented two alternatives: 1.- Item's menu attached to the type icon on the left => !^dashbuilder_navitem_actions2.mp4 2.- Item's menu separated into a new icon on the right => !^dashbuilder_navitem_actions.mp4 Option #2 makes more explicit the presence of a menu since the icon shows up only when the mouse is over the item. Which do you like the most?	ask opinion	https://issues.jboss.org/browse/DASHBUILDE-218
14	DASHBUILDE-218	[~jhrcek][~davidredhat] Cool! And I like Option #2 better as well. Couple of ideas: - I noticed that the cog icon turns blue on hover. For consistency, could we have all of the action icons (on the right) turn blue on hover? - Except, wdyt of having the delete (trash can) icon turn red on hover to indicate that it's a "destructive" action? For example: !Screen Shot 2017-05-24 at 10.40.20 AM.png thumbnail!	ask opinion	https://issues.jboss.org/browse/DASHBUILDE-218
15	DASHBUILDE-93	Hello [~davidredhat], I'm looking at latest kie-wb and dashbuilder showcase webapps, but I don't see any effect of this setting (nor in the chart preview in dataset editor, nor in the created perspectives). For which charts is this setting expected to work? For every that has it in the Display tab? I tried with Bar/Area/Line charts and I didn't see any effect on the axis labels. Could you please verify that it works?	need description to test	https://issues.jboss.org/browse/DASHBUILDE-93
16	DASHBUILDE-70	I cannot reproduce it yet. I have tried it using the dashbuilder's showcase and the whole kie-wb, both webapps tested in Chrome, Firefox and Safari (on mac), and the text alignment is fine... any ideas?	can't reproduce. ask for ideas	https://issues.jboss.org/browse/DASHBUILDE-70
17	DASHBUILDE-28	Ok Roger. Talking about Patternfly , could you please share a short info on this initiative? I saw the webpage (https://www.patternfly.org/), but I have no idea what's the connection to our projects. Does it mean using some additional library of widgets or something, or is it just a set of design recommendations to make the apps more user friendly / unified?	ask for more info	https://issues.jboss.org/browse/DASHBUILDE-28
18	DASHBUILDE-10	Questions about this issue: - The project already depends on resteasy (2.3.7). It is a dependency of Uberfire. Should I keep this version? - Is there a code formatter, to decrease merge needs? - I think that the rest-api should reside on module 'dashbuilder-webapp'. Is it ok? - Any consideration about packages structure/names?	need more info	https://issues.jboss.org/browse/DASHBUILDE-10

DROOLS

#	Requirement	Question	Notes	URL
1	DROOLS-1652	What should the URL look like in your opinion [~tirelli]? From a RESTful point of view, we're creating a new kie-container resource, so in theory it should be a PUT (create resource) on URL"/server/container/{id}" that accepts a DMN file. However, that URL already exist. We could provide support for different request bodies in that operation (e.g. a 'KieContainerResource` and a `DmnKieContainerResource`), but that could become a bit nasty ... We can also use a POST on something like "/server/containers/{id}/resources" and have that create a container resource under "/server/containers/{id}" (the reason I don't like a PUT in this case is because we create the resource (container) under a different URL than our request URL). The reason for the name "resources" (plural) is that we could provide multiple resources in the body (e.g. multiple DMN files). Second, apart from the DMN XML file (or files), do we also want to be able to post more information in the body (e.g. like additional information we pass in `KieContainerResource`)? And finally, should this be part of the DMN extension or KIE-Server Commons? The argument for DMN would be that otherwise we would expect Commons to have DMN knowledge. On the other hand, it's just deploying a KIE Containerner based on a resource file. The fact that that's a DMN model does not really matter (in theory this could be any resource file). Personally KIE-Server Commons makes most sense.	Ask opinion of solution, description	https://issues.jboss.org/browse/DROOLS-1652
2	DROOLS-1652	I don't think we should make this for jBPM as by nature it is stateful so once you create such resource and bounce the server it's gone while the data in db remains - kind of orphans. That in general applies to stateless/in memory stuff too, how would you deal with server restarts? Or is this intended for quick checks of resources? My vote is to include it in DMN extension or Drools if you would like to support other rule related resource types.	Ask solution	https://issues.jboss.org/browse/DROOLS-1652
3	DROOLS-1652	I see two approaches to this: - if we consider it as demo only I would opt for having this to be as separate extension - so it can be deployed outside of the official distribution that way we avoid potential issues/bug reports that it won't survive server restarts - make it all the way down to survive the restarts by saving given kjar into maven repo locally and then simply register it as any other container. Not sure what is preferred as I don't know the origin of this request....	Compare solution	https://issues.jboss.org/browse/DROOLS-1652
4	DROOLS-1577	[~rmorrisse] , for sequential execution, have you tried using ruleflow-group instead? I am closing this ticket as I believe ruleflow-group provides you the functionality you are looking for, and for the future we are working on the concept of rule units for orchestration. Feel free to add more information/justification in case you believe the existing features do not meet your requirements and we can re-evaluate.	Ask rmorrisse try sth	https://issues.jboss.org/browse/DROOLS-1577
5	DROOLS-1561	Can you please check the container in the kie-server if you haven't enabled the kie-scanner by mistake? There is a configuration there to run the kie-scanner continuously or to deactivate it.	Ask code review	https://issues.jboss.org/browse/DROOLS-1561
6	DROOLS-1561	[~mfusco] can you please quickly fix this in the docs? Use this ticket for it.	Ask mfusco fix issue	https://issues.jboss.org/browse/DROOLS-1561
7	DROOLS-1444	Implementing AutoCloseable is only useful if you're planning to use the AutoCloseable resource in a try-with-resources block, but to be honest I don't think it will be common at all to use a KieSession in that way. Do you have a use case where you may need (or want) to do that? If so can you please paste the code snippet where you tried to do this, so I could have a better idea regarding the possible implications?	Ask use case	https://issues.jboss.org/browse/DROOLS-1444
8	DROOLS-1444	In that example you're showing how to use try-with-resource with the StringWriter. That makes perfectly sense to me but it is not what I was asking. I was asking for a practical example using the KieSession because in my opinion in that case it is not that useful and using a try-with-resource is more verbose and cumbersome than simply adding a dispose. I cannot add this to drools 6.x because it has to stay Java 6 compatible, but I could do this on master (drools 7) that is Java 8 based. However I really don't see the usefulness in doing so. Moreover your scenario seems to fit far better a StatelessKieSession which doesn't require any dispose. Did you give it a try?	Provide solution and ask some one to try it	https://issues.jboss.org/browse/DROOLS-1444
9	DROOLS-1443	What you're suggesting sounds interesting but it is not entirely clear to me. In particular I don't understand how you could implement meaningful hashCode and equals methods on a ClassLoader. If you already implemented this could you maybe send a pull request with what you did or at least paste it somewhere? Thank you.	Ask PR	https://issues.jboss.org/browse/DROOLS-1443
10	DROOLS-1419	[~geoffrey] it seems this JIRA might need an update to resolved?	Ask update issue	https://issues.jboss.org/browse/DROOLS-1419
11	DROOLS-1398	[~Rikkola] sorry, I am confused. Does this change need to be done in the workbench or in core?	Ask which project this issue belongs to	https://issues.jboss.org/browse/DROOLS-1398
12	DROOLS-1387	I think I can confirm this. I created the "same" benchmarks in kie-benchmarks using JMH (I used DRLs and fact inserts from [~tkobayashi]'s reproducer). I put it in my separate branch, which you can find here [1]. Please *DO NOT MERGE* it into upstream, because it contains changes in pom.xml's and other benchmarks that contain references to APIs available only in 7.x. I will create separate PR for upstream/master when needed. [~tkobayashi] could you please check my benchmarks if I migrated them correctly? My results when I ran the benchmarks locally on my computer are here [2] . [1] https://github.com/baldirir/kie-benchmarks/tree/DROOLS-1387 [2] http://pastebin.com/dsAMv5x3	Ask code review	https://issues.jboss.org/browse/DROOLS-1387
13	DROOLS-1387	[~mfusco], Sorry about not having distilled the test case enough. I'll make sure to be careful about eval use. Btw, I have removed eval from the test case rules (Rule1.drl and Rule2.drl) and attached (brms-perf-comarison-02.zip). The result still tells that ReteOO is 30% faster than PHREAK in average. Am I still overlooking some eval use in the rules? Or could we say "eval" and "accumulate" and "exists" are all slow in PHREAK?	Ask describe of req	https://issues.jboss.org/browse/DROOLS-1387
14	DROOLS-1387	Hi [~tkobayashi]. The resulting diff can be also influenced by thread contention. I see that you have some sync points in the reproducer. Could you please try to run benchmarks from my reproducer? If you use IDEA or Eclipse, there is JMH plugin for these IDEs. If you install it, you can run the benchmarks similarly as JUnit tests. If you have any questions about this, please let me know. You can find me also on IRC on #droolsdev channel. I will also try to rerun the benchmarks. We can compare the numbers then.	Question reproducer	https://issues.jboss.org/browse/DROOLS-1387
15	DROOLS-1385	Friends, Please comment your ideas	Ask solutions	https://issues.jboss.org/browse/DROOLS-1385

#	Requirement	Question	Notes	URL
16	DROOLS-1378	[~tkobayashi] I don't understand why an user may want to put a text into a cell with Numeric type. Is there a valid use for this? In general we should stop trying to parse any sort of garbage thrown in by users. This sort of guesses and type coercions causes many more damages than benefit on the long term and once they are in you cannot get rid of them anymore, otherwise you'll break compatibility.	Ask use case of issue	https://issues.jboss.org/browse/DROOLS-1378
17	DROOLS-1355	[~porcelli] can you please take a look at this and check how that relates to the work you are doing?	Ask some one take look at issue	https://issues.jboss.org/browse/DROOLS-1355
18	DROOLS-1345	This may be harder to remove that expected. It seems some people are still using it. That's fine, we just need a sources for the jsr94 jar. [~mfusco] or [~mark.proctor] do you know if those sources are available somewhere? (I was not able to find anything googling for a while).	Ask extenal resouces	https://issues.jboss.org/browse/DROOLS-1345
19	DROOLS-1344	[~kdwivedi], I have talked to Mario about this issue (docs jira linked). We do not have any Aries BP docs in our current guides (upstream nor downstream). If this issue blocks our release, as suggested by [~rsynek], this should be addressed (and incorporated in the downstream docs) by 6.4 release. I am not sure whether Mario has the time to finish this jira on time, maybe prioritization is in order? WDYT?	Ask some one's schedule	https://issues.jboss.org/browse/DROOLS-1344
20	DROOLS-1319	Hi Edson, Could you please share me some links / sample configs so that i can try out the same in my project ? Thanks, Kavi sri K K	Ask project setting document	https://issues.jboss.org/browse/DROOLS-1319
21	DROOLS-1270	Can you please clarify your request and provide an example?	Ask more details of req	https://issues.jboss.org/browse/DROOLS-1270
22	DROOLS-1250	I fixed the issue by adding <code>*dialect "mvel" *</code> at the two rules in Sample.drl. Is it a bug in the drools wizard that dialect is missing in the examples? or It supposes to work without dialect? Thanks.	Ask type of issue (bug or feature)	https://issues.jboss.org/browse/DROOLS-1250
23	DROOLS-1250	I am facing the same issue. But the code runs fine despite compile time error. Also the error message makes no sense. Why does it say <code>com.sample.DroolsTest.Message</code> resolves to a package when it clearly isn't.	Question why issue marked as fix when it is not fix	https://issues.jboss.org/browse/DROOLS-1250
24	DROOLS-1250	Hi Tianwei Chen, just followed your comment and still got the same error. Did you still use droolsjbpm tools 6.4 plugin? I have to add dialect mvel to make it work. Is it an issue of plugin or runtime?	Ask version used by issue reporter	https://issues.jboss.org/browse/DROOLS-1250
25	DROOLS-1225	Hi [~kuniyasu], if I'm understanding correctly what you're asking here is adding those new SequentialMaxPriority and SequentialMinPriority optional fields in the "RuleSet" section of the XLS for each sheet and setting the salience of the generated rules accordingly. Can you confirm?	Ask clarification of req	https://issues.jboss.org/browse/DROOLS-1225
26	DROOLS-1221	Hi Youcef, Thanks for the slide set about the Italian HACEP project. It makes it much more clear. I intend to port the Italian HACEP solution to the Amazon cloud replacing Infinispan, ActiveMQ, Camel Route, etc. with Amazon equivalents. The grouping is done on Tenant (aka customer environment). The domain is Business Event Monitoring in the cloud. Are you going to use Apache Geode in your project as a replacement for Infinispan? If so, why?	Ask which server container reporter use and why	https://issues.jboss.org/browse/DROOLS-1221
27	DROOLS-1168	ok, but what can be the workaround for the problem? How can I force the execution of update (final FactHandle handle, BitMask mask, Class <?> ModifiedClass). Since you are in the development you could give me some guidance? Best Regards	Ask solution	https://issues.jboss.org/browse/DROOLS-1168
28	DROOLS-1168	Did you try what I suggested in my last comment? Something like <code>{code} DroolsUtil.completeEntity(\$entity, fields, \$untMis); \$entity.setUntMisVen(\$entity.getUntMisVen()); update(\$entity) {code}</code>	Ask to try suggestion solution	https://issues.jboss.org/browse/DROOLS-1168
29	DROOLS-1168	This is a workaround finded by me rule "completeUntMisVen" when <code>eval(checkRulesEnabled == 0) \$entity : MtbLisvData(untMisVen == null && codMart != null) \$untMis: java.lang.String() from SalesRules.completeUntMisLisv(conn, \$entity) then //modify (\$entity) { setUntMisVen(\$untMis) } String[] fields = new String[]{"untMisVen"}; DroolsUtil.completeEntity(\$entity, fields, \$untMis); long l = 434576465786346l; //random long LongBitMask ll = new LongBitMask(l); update(\$entity, ll, \$entity.getClass()); //update(\$entity) end It is a correct solution? In this way bitmask is not 0 and next rule fire correctly</code>	Ask review solution	https://issues.jboss.org/browse/DROOLS-1168
30	DROOLS-1168	Solution proposed in last post is not possible because <code>\$entity : EntityBase()</code> from <code>\$entityFind \$entity is EntityBase</code> and <code>setCodDiviCont()</code> , <code>setCambioDiviCont()</code> methods are in one of class listed (JtbRLavr, JtbRLavr, GtbSpes, ...) In compile phase drools throw error because not find methods in entitybase. You confirm this behavior in 6.4.0?	Ask when issue will fix	https://issues.jboss.org/browse/DROOLS-1168
31	DROOLS-1168	I would like a solution, I can invoke the mask calculation in some way?	Ask solution	https://issues.jboss.org/browse/DROOLS-1168
32	DROOLS-1168	May I understand why you're insisting using property reactivity for this use case? Which advantages do you see?	Ask retional/use case of req	https://issues.jboss.org/browse/DROOLS-1168
33	DROOLS-1168	I'm sorry, but property reactivity doesn't fit your use case. Have you tried to disable it and see what happens? I don't see why it shouldn't work.	Ask try suggestion solution	https://issues.jboss.org/browse/DROOLS-1168
34	DROOLS-1168	No, because I want <code>PropertyReactive</code> . With no property reactive rules have cascade firing?	Show uncertainty	https://issues.jboss.org/browse/DROOLS-1168
35	DROOLS-1168	And how to Drools know that property was changed?	Ask detail of existing code	https://issues.jboss.org/browse/DROOLS-1168
36	DROOLS-1168	Ok, it is just for efficiency. But you consider my request for enhancement next version? Best Regards	Ask put issue to furture work	https://issues.jboss.org/browse/DROOLS-1168
37	DROOLS-1137	Note that the RPM work is working around this, so RPM's shouldn't really impact the decision here. I do think it's an improvement nonetheless. In this proposal (let's call it proposal A), what are the changes? - Will it remove 1 POM (because kie-parent-with-dependencies and kie-platform-bom will become kie-third-party-bom)? - Child poms don't change, only top level multiproject parent poms.	Ask changes	https://issues.jboss.org/browse/DROOLS-1137
38	DROOLS-1071	Probably a null dereferencing inside an mvel constraint in the DRL, but the drools-compiler should mention the right line number in the stacktrace. [~idoflax] Have you tried commenting out some of the rules, to identify which of the DRL rules causes this exception?	Ask try to find which code cause problem	https://issues.jboss.org/browse/DROOLS-1071
39	DROOLS-1071	Can you specify which DRL line exactly caused it? We want to improve the error message in a future version, so it mentions that exact DRL line.	Ask where cause problem	https://issues.jboss.org/browse/DROOLS-1071

#	Requirement	Question	Notes	URL
40	DROOLS-1053	this has been recently added so you can set system property to enable text based date format for JSON marshaller: https://github.com/droolsjbpm/droolsjbpm-integration/blob/master/kie-server-parent/kie-server-api/src/main/java/org/kie/server/api/marshalling/json/JSONMarshaller.java#L66 though it does use certain data format pattern: https://github.com/droolsjbpm/droolsjbpm-integration/blob/master/kie-server-parent/kie-server-api/src/main/java/org/kie/server/api/marshalling/json/JSONMarshaller.java#L73 would this be sufficient or shall we have another system property to control actual date format pattern?	Ask completeness of solution	https://issues.jboss.org/browse/DROOLS-1053
41	DROOLS-1052	[~psiroky] only for master or for other branches too ?	Ask completeness of solution	https://issues.jboss.org/browse/DROOLS-1052
42	DROOLS-1052	Should I only look in the java files or in the drl files too? What should I do in this case: <code>/** * Finds currently available server ports. * * @author Apache Directory Project * @see IANA.org */</code> There are a bunch of @authors in test java files, like i.e. <code>public class StandaloneDocBuilder { public static void main(String[] args) throws FileNotFoundException, ParseException { String drl = ""; drl += "# important information\n"; drl += "# about this package\n"; drl += "# it contains some rules\n"; drl += "package org.kie.test\n"; drl += "global java.util.List list\n"; drl += "# Really important information about this rule\n"; drl += "# Another line because one was not enough\n"; drl += "#\n"; drl += "# @author: trikkola\n";</code> Should this be removed too?	Ask completeness of solution	https://issues.jboss.org/browse/DROOLS-1052
43	DROOLS-1008	I am not sure of what you're asking here. I guess you're talking about the maven default behaviour of only fetching SNAPSHOTS once per day, but if this is the case I think you can configure maven to enforce of checking SNAPSHOTS at each request. Please let me know if this is case or try to better clarify this ticket.	Ask clarify req	https://issues.jboss.org/browse/DROOLS-1008
44	DROOLS-1008	[~mfusco] [~andrew.collins] Any status on this. I am able to duplicate this error and have to clear out my local .m2 before the snapshot gets updated in business-central on a deployment. I have the update policy set to always in my settings.xml so I am pretty sure everything is configured correctly My scenario is slightly different but basically the same build and deploy kjar to remote nexus - using a separte settings.xml for this build and a separte local reposiotry as to not conflict with what BC is using deploy GAV to BC The deployment is still the old snapshot	Ask status	https://issues.jboss.org/browse/DROOLS-1008
45	DROOLS-972	Can you please send a PR adding a new example in drools-examples-cdi so I could start working from there?	Ask PR	https://issues.jboss.org/browse/DROOLS-972
46	DROOLS-963	This should be fixed for 6.4.0.Final, does it sound possible?	Ask fix issue in a specific release version	https://issues.jboss.org/browse/DROOLS-963
47	DROOLS-905	Thanks for the pointers [~martenscs]. Is there any other prefix we need to worry about? Or does <code>startsWith("bundle")</code> cover all the needed cases?	Ask completeness of solution	https://issues.jboss.org/browse/DROOLS-905
48	DROOLS-901	any news here [~mfusco]?	Ask status of issue	https://issues.jboss.org/browse/DROOLS-901
49	DROOLS-823	Hi Geoffrey, I don't know what should be the expected behaviour? i.e. I run in drools <code>jdeps -jdkinternals ./drools-verifier/src/test/resources/org/drools/verifier/model.jar</code> only for test purposes. It says: <code>Picked up JAVA_TOOL_OPTIONS: -Dfile.encoding=UTF-8 ./drools-verifier/src/test/resources/org/drools/verifier/model.jar -> /usr/java/jdk1.8.0_20/jre/lib/rt.jar</code> Hope thta this is OK. When an error is found, what would be the behaviour?	Ask correct reaction when error happened	https://issues.jboss.org/browse/DROOLS-823
50	DROOLS-823	I think 2) is ok and 1) is ignorable. 3) is definitely a true negative, something that we need to have an inventory of and see if we can reduce that inventory to 0. But those are <code>sun.misc.Unsafe</code> usages from our depended libraries (<code>com.google.common</code> in this case), not from ours ourself. Are there any from us directly (<code>org.drools/org.kie/org.optaplanner/org.jbpm</code>)? We should handle those first.	Ask whether any error caused by project itself	https://issues.jboss.org/browse/DROOLS-823
51	DROOLS-821	Last week [~geOffrey] added new feature for <code>optaplanner-engine</code> (see https://github.com/droolsjbpm/droolsjbpm-integration/commit/5221c0731ec6008e434dd93b9a7c31f514211592). I think it is the time to rename the osgi related modules and use "kie" instead of "drools" whenever possible. So we would have <code>kie-osgi</code> , <code>kie-osgi-integration</code> , <code>kie-karaf-features</code> , etc. [~mfusco], [~tirelli] what do you think?	Ask opinion on rename files	https://issues.jboss.org/browse/DROOLS-821
52	DROOLS-760	There were some changes related to formatting done by [~kgaevski]. I believe the formatters are now up-to-date. [~kgaevski] can you please confirm that?	Ask calrify issue	https://issues.jboss.org/browse/DROOLS-760
53	DROOLS-760	My understanding was that this JIRA is mainly about not having up-to-date formatters for IDEs, which I believe got mostly fixed? We can keep this open until <u>everything</u> gets fixed, but that may very well never happen.	Ask whether issue is fixed	https://issues.jboss.org/browse/DROOLS-760
54	DROOLS-729	Can you come onto irc and speak to Mario (mfusco). This will make it much easier to discuss all the related issues and for us to figure out how to integrate your changes. http://drools.org/community/chat.html	Ask some one talk to Mario	https://issues.jboss.org/browse/DROOLS-729
55	DROOLS-686	Petr in that SNAPSHOT setup, after building successfully once, did you try pluggin out your network cable and building again? The second time it must be able build offline (and that originally didn't work with <code>pressgang</code> 3).	Ask specific actions of user	https://issues.jboss.org/browse/DROOLS-686
56	DROOLS-660	<code>public Class<?> fastFindClass(final String name) { lock.lock(); Class<?> cls = findLoadedClass(name); if (cls == null) { final byte[] clazzBytes = this.store.read(convertClassToResourcePath(name)); if (clazzBytes != null) { String pkgName = name.substring(0, name.lastIndexOf('.')); if (getPackage(pkgName) == null) { definePackage(pkgName, "", "", "", "", "", null); } cls = defineClass(name, clazzBytes, 0, clazzBytes.length, PROTECTION_DOMAIN); } if (cls != null) { resolveClass(cls); } } lock.unlock(); return cls; }</code> i that ok? i find the lock has been removed in 6.2.0.Final, is that means 5.6.0Final has dead lock problem ? can i use 5.6.0 in my project?	Ask usability of a specific version	https://issues.jboss.org/browse/DROOLS-660

#	Requirement	Question	Notes	URL
57	DROOLS-660	Are you reporting a deadlock problem ??? If so, please change this ticket to bug - blocker AND provide a reproducer. Reading your comment, this is not a bug report, not a feature request, but a question. In this case, please close this ticket and ask your question here https://groups.google.com/forum/#!forum/drools-development (By the way, the classloader architecture was completely redesigned for 6.x, so you can't just compare 5.6 and 6.x. Is there any reason why you would prefer an old, no longer supported version for your project? This would be a discussion for here: https://groups.google.com/forum/#!forum/drools-usage) Thanks	Ask type of issue	https://issues.jboss.org/browse/DROOLS-660
58	DROOLS-660	Could you try using 6.x? The 5.6 series is a community version and no longer maintained, so even with a reproducer it may take a long time to fix it and you would have to build a SNAPSHOT version anyway. If the issue was still present in 6.x, things would be very different Davide	Ask user to try a specific version	https://issues.jboss.org/browse/DROOLS-660
59	DROOLS-656	AFAIK this was done? I am not sure about this. Geoffrey, If you know that this werd done please resolve and close this Jira.	Question whether issue is done	https://issues.jboss.org/browse/DROOLS-656
60	DROOLS-656	mbiarnes, would you mind creating a new repo under droolsjbpm, push some stuff to it and then force push some more stuff to it? That will answer the question above. Note: We can't on existing repo's without potentially breaking them which is not an option.	Ask new PR	https://issues.jboss.org/browse/DROOLS-656
61	DROOLS-652	{quote} What about allowing this, instead? declare event Ev // declaration here end define Ev as // LHS-like definition here end {quote} Seems to me that you are proposing the same thing that is already in the document, just using different keywords?	Ask review of solution	https://issues.jboss.org/browse/DROOLS-652
62	DROOLS-651	The documentation uses the term <code>_binding variable_</code> 8 times, so bindings seems to be variables too (but of course special ones). I understand that these bindings are not yet cached because of constraint idempotence. However, in my view the implication should be, that especially a non changing result <code>+can+</code> be cached! Rather than a <code>_good practice_</code> , the documentation says: {quote} "Property accessors <code>+must+</code> not change the state of the object in a way that may effect the rules. Remember that the rule engine effectively <code>+caches+</code> the results of its matching in between invocations to make it faster." {quote} So if kind of functional immutability is required and used inside the engine, I would indeed think about caching method calls (with variables) in constraints, as this saves potentially massive processor time and would be much more intuitive. I don't understand your sentence {quote} "If you can't or desire not to guarantee this property, Drools won't enforce it caching the first value returned." {quote} Is it possible, to enforce Drools to cache variable bindings? Drools shouldn't cache values, that change. Let me say that I'm still intrigued about the huge, powerful Drools framework and that we get this as Open Source including this helpful support. Feel free to change the issue-type to a feature request, as the behavior is not strictly incorrect but optimizable.	Ask posibility of a solution	https://issues.jboss.org/browse/DROOLS-651
63	DROOLS-599	Hello, just a quick question, is it really required to keep the header-comment-license to the generated files of this archetype, please? I've noticed for example [this in commit 7172a50] https://github.com/droolsjbpm/droolsjbpm-knowledge/commits/master/kie-archetypes/kie-drools-archetype/src/main/resources/archetype-resources/src/main/resources/rules.drl by [~mbiarnes] the net result is that an user would generate artifact from maven and then erase the header-comment-license as the user starts to write own rules in the generated drl file, because the rules are the user's own. Similar analogy would apply to the other files generated by this artifact, as generated files are indeed meant to be modified to adapt to user's need, and the reason why I didn't put this header in the original PR. I can understand it is not a big deal to manually erase this header in all generated files where required, but if somebody is creating a lot of self-contained rule packages on a daily basis, it has an impact trust me =) Just open question, thank you, Ciao!	Ask required element in code	https://issues.jboss.org/browse/DROOLS-599
64	DROOLS-599	Hi [~mbiarnes] , in reference to my previous comment above, and now discussing with [~mfusco] , we would like to highlight the availability of the Kie Archetype to end users. However, [since this archetype was first proposed] https://github.com/droolsjbpm/droolsjbpm-knowledge/pull/74/files#diff-ae073dbf7d2d212e8ef79caab795cb87 , now [License file-headers are attached to all files] https://github.com/droolsjbpm/droolsjbpm-knowledge/commit/7172a50c4989800a8cda029b1e5efaec9c1d7dcb#diff-42a73ec3fc7b3eda342f8c8ad2bdddec31 ; with the actual consequence that once you use the archetype, as a user, your files are containing the Red Hat file-header license; and I believe this could be very confusing for the end-users, and also I believe is wrong. My proposal is to remove License file-headers, at least from all files which are serving as a template-generation. I have two question, please? # Do you agree with me reverting your change? # Did you notice the License file-headers are missing thanks to some automatic tool? Is there a way to manage exclusion, either in the tool or by adding a special header to the file itself? Thank you in advance for your feedback.	Ask mbiarness look at specific code file	https://issues.jboss.org/browse/DROOLS-599
65	DROOLS-566	Thanks, what about the other way around, eg DROOLS SNAPSHOT with mvel 2.1.9 - does that work?	Ask issue working environment	https://issues.jboss.org/browse/DROOLS-566
66	DROOLS-552	I double checked that all the tests in your pull request run successfully and merged it. I am not sure of what we could do to better support other encodings. Any suggestion?	Ask completeness of solution	https://issues.jboss.org/browse/DROOLS-552
67	DROOLS-529	Could you please submit the rules you used to benchmark the engine? Thanks!	Ask more code	https://issues.jboss.org/browse/DROOLS-529
68	DROOLS-529	We'd like to try and fix this for 6.1, which is due within 2 weeks. But to do so we need some benchmarks, can you upload a zip for us?	Ask source code	https://issues.jboss.org/browse/DROOLS-529
69	DROOLS-507	Is there any update on this? i'm facing the same issue..	Ask status	https://issues.jboss.org/browse/DROOLS-507
70	DROOLS-507	The FireUntilHalt command is marked as unpersistable (for obvious reasons - it is usually executed in a dedicated thread), which causes the exception even before the command is attempted. Could you elaborate more on your use cases? In particular, what would you expect to happen if a running session was persisted and then restored?	Ask why/use case of issue	https://issues.jboss.org/browse/DROOLS-507
71	DROOLS-492	[~cmoulliard] Great! If the exception isn't being thrown anymore, it sounds like this issue (the exception, old class, etc) has been fixed. Does it make sense to assign the issue to you and close it?	Ask some one take look at this issue	https://issues.jboss.org/browse/DROOLS-492
72	DROOLS-483	Thanks for reporting this, can you submit a PR? It will be easier to discuss the details Thanks! Davide	Ask PR	https://issues.jboss.org/browse/DROOLS-483

#	Requirement	Question	Notes	URL
73	DROOLS-483	Hi there. It's been more than 3 months and neither this issue nor the pull request has seen any kind of response, so I was wondering if something important is missing (other than time of the Drools team side). Should I continue rebasing the modification on master from time to time?	Ask status	https://issues.jboss.org/browse/DROOLS-483
74	DROOLS-470	This issue still hurts. Any chance we can get it for 6.4?	Ask possibility of fix it in specific version	https://issues.jboss.org/browse/DROOLS-470
75	DROOLS-470	Still hurts. Since this is on the drools engine (not drools workbench), [~mfusco] any chance you want to take a look at it?	Ask some one take look at it	https://issues.jboss.org/browse/DROOLS-470
76	DROOLS-424	The configuration part of this would be easy. e.g. in the xsd: <xsd:attribute name="type" use="required" type="knowledgeSessionTypeEnum"/> could be replaced with: <xsd:attribute name="type" use="optional" type="knowledgeSessionTypeEnum"/> <xsd:attribute name="scope" use="optional" type="knowledgeSessionScopeEnum"/> <xsd:simpleType name="knowledgeSessionScopeEnum"> <xsd:restriction base="xsd:string"> <xsd:enumeration value="request" /> <xsd:enumeration value="session" /> <xsd:enumeration value="application" /> </xsd:restriction> </xsd:simpleType> and then in KnowledgeSessionDefinitionParser, scope="request" would be equivalent to type="stateless", and scope="application" would be equivalent to type="stateful". The new option would be scope="session". However I'm not sure when the StatefulKnowledgeSession would actually be created as it requires access to the HTTP session to store it for future requests. Does the HTTP request get passed all the way through from Camel to the Drools classes? I tried tracing it through but lost it somewhere in the depths of Camel.	Ask extenal source	https://issues.jboss.org/browse/DROOLS-424
77	DROOLS-365	I see what you are trying to do, but I am not sure that such an attribute would be feasible, or even solve your problem. What happens if a rule inserts (and modifies, and retracts...) more than one object? And what happens if the same object is insert by more than one rule? Have you tried using the EQUALITY assertion behavior and/or the @propertyReactive annotation? The former considers equivalent (i.e. equals()) object to be the "same" and thus do not activate rules multiple times, regardless of the rule that generated the objects. The latter allows you to control how rules fire in case of updates. Best Davide	Ask requirement clarification	https://issues.jboss.org/browse/DROOLS-365
78	DROOLS-365	I don't see the problem with modifications to multiple facts: you could condition on the NOT of the logical AND of all facts. Or the logical OR; I think both approaches are defensible. As for your second example, what's the issue with multiple rules asserting the same fact? In any case, I think the biggest challenge to implementation is that the rule consequence is a snippet of Java, and thus can do absolutely anything. For example, it could conditionally assert a fact. How do you modify the rule then? Also conditionally? Based on runtime input? It seems really challenging. The example that actually drove me to file this bug looks something like this: {noformat} rule "1" when accumulate (Foo(\$type : type, \$value : value, \$total : sum(\$value)) then insert(new Bar(\$type, \$total)); end rule "handle all Bars of type first" when \$bar : Bar(type == 'first', \$total : total) ... then retract(\$bar); end rule "handle all Bars of type second" when \$bar : Bar(type == 'second', \$total : total) ... then retract(\$bar); end rule "generic rule to handle all leftover Bars" salience -1 // let custom Bar rules fire first when \$bar : Bar(\$total : total) ... then retract(\$bar); end {noformat} The flow here is to convert the sum of all Foes (of a given type) into a Bar. Subsequently, there's a rule to handle each type of Bar, with a generic rule to handle Bars without a custom rule. To get that fallthrough behavior I need: # Salience -1 on the generic rule, so it fires last. # Retraction of each Bar as it's processed. However, the retraction causes the summation rule to reactivate, unless I add an explicit {{not Bar(type == \$type)}} condition to the rule. Granted, in this case, I suppose the feature request mentioned here wouldn't work, because it'd add {{not Bar(type == \$type, total == \$total)}}, which would allow the rule to reactivate. For this example, I don't think equality-based assertion will solve my problem; fundamentally the retraction of a Bar causes us to reevaluate the accumulate rule. Likewise, I don't think @propertyReactive will help either; the issue here isn't with modifying a single property, but with removing an entire fact.	Ask requirement clarification	https://issues.jboss.org/browse/DROOLS-365
79	DROOLS-355	Hi, I'm Henry from Red Hat IoT engineering team. We are working on putting Drools on IoT gateway devices. I'd like to use kie-server-client and install it into Eclipse Kura server (which is really OSGi container based on Eclipse Equinox). I'm running into the same problem as drools-core is a dependency of kie-server-client. Do you have an estimate of when this issue could be resolved? Cheers!	Ask estimation of when issue can be fixed	https://issues.jboss.org/browse/DROOLS-355
80	DROOLS-355	Hello [~hkonsek]. Would it possible to attach a reproducer which shows the error you are encountering? It would help us a lot. Thanks!	Ask code	https://issues.jboss.org/browse/DROOLS-355
81	DROOLS-355	I did test and merge your pull request. And that part works. However I am now stuck at the next bundle: {code} org.osgi.framework.BundleException: The bundle "org.kie.remote.client_6.4.0.201604141724 [26]" could not be resolved. Reason: Missing Constraint: Import-Package: org.kie.remote.services.ws.common; version="0.0.0" {code} However the JAR {{kie-remote-ws-common-6.4.0.Final.jar}} is not an OSGi bundle. I updated the repository, so you can test if you like. Should I open a new issue?	Ask extent to new issue	https://issues.jboss.org/browse/DROOLS-355
82	DROOLS-355	And which is not available on Equinox ;-)	Ask extent to new issue	https://issues.jboss.org/browse/DROOLS-355
83	DROOLS-337	Ok, you are right for the backward compatibility. ==> ena"B"led is fine. What should i do for the next step ? I didn't understood. Should i select under the workflow button, the "Link Pull Request" choice ?	Ask correct operation	https://issues.jboss.org/browse/DROOLS-337
84	DROOLS-315	Edson - what state is the patch in - has it made it in - will it make it?	Ask status	https://issues.jboss.org/browse/DROOLS-315
85	DROOLS-315	Hey Edson, What's the status of this issue? Do we have sumInt(), sumLong(), sumBigDecimal() available in Drools? Or better yet, support for overloaded methods support, so sum(int), sum(long), sum(BigDecimal) just do the right thing (instead of using sum(double)).	Ask status	https://issues.jboss.org/browse/DROOLS-315

#	Requirement	Question	Notes	URL
86	DROOLS-300	Issue #2 In our discussions, we have remarked that drools is not real-time, so the scheduled time is a lower bound for the actual activation: i.e., expect the rule to activate any time AFTER the scheduled time. Should we have an activation-specific deadline? E.g., activate with a period of 1 day, but cancel any activation that can't be fired within 1 hour since the activation. If you miss the 1hr-window, just skip and go for the next day	Ask when to fix issue	https://issues.jboss.org/browse/DROOLS-300
87	DROOLS-299	Do you mean I need to add the org.drools:drools-jboss-integration dependency to the project itself? If I do this it adds over 80mb of dependencies to the deployment.	Ask clarification of req	https://issues.jboss.org/browse/DROOLS-299
88	DROOLS-297	[~mbiarnes] On the ip-bom list, Keith and Gary we're wondering how to do this. How did you fix this for us? We 'll want to extract that to the ip-bom, so the other projects benefit from it too.	Ask solution or code	https://issues.jboss.org/browse/DROOLS-297
89	DROOLS-297	I added a profile "nonsnapshot" (locally) on kie-parent. When I do a mvn clean install -Dfull -DskipTest -Dnonsnapshot in any project, when there is a SNAPSHOT, the build fails. OK. I don't see a time-saving here, since during a release, after the project version update I do a find in path for -SNAPSHOT, to find -SNAPSHOTS (if there are any). If we then execute after version update -Dnonsnapshot, this will be like a second check if there is still a SNAPSHOT. I don't know if it is worth to add this profile to kie-parent pom?	Ask benefit and cost	https://issues.jboss.org/browse/DROOLS-297
90	DROOLS-297	Does this need to be an extra profile? Can't the enforcer rule be smart enough to do the following "if the version of this module is not a SNAPSHOT, then it should have no SNAPSHOT dependencies"?	Ask completeness	https://issues.jboss.org/browse/DROOLS-297
91	DROOLS-290	Constraints are handled differently in 5.5 than 5.0. In later versions, you can have "free-form" constraints as opposed to simpler <Field OP Value>. This is not a bug. Are you manipulating the Descr tree yourself? Those classes have never been part of the public API, as far as I know	Ask operation when find the issue	https://issues.jboss.org/browse/DROOLS-290
92	DROOLS-166	With reference to: https://groups.google.com/d/msg/drools-usage/fbF8KhOPCKM/D9pCQvLdfIEJ Feedback by Edson Tirelli: {quote} As we moved to a minimal public API, the getConfiguration() method was moved to the internal API, but I think this can be reconsidered. May I ask you please to open a JIRA requesting this method be moved to the public API? Meanwhile, as a workaround, you will need to cast it to the internal API: InternalKnowledgeBase.getConfiguration(). {quote} As I understand what should be the proper solution based on Edson's feedback, I also was able to discover this bug on JIRA; so I'm going to re-use this record instead of creating a new one. I do confirm that the workaround in the meantime works for me, thank you very much Edson. @Petr, for KieSession you can do .getSessionConfiguration() of the public API interface, works for me at least.	Ask extend to new issue	https://issues.jboss.org/browse/DROOLS-166
93	DROOLS-166	[~tari_manga] It seems you can ask for a value for a particular KieSession option, but you can get a list of the options. Or maybe I am overlooking something? Basically I would like to get the list of options (and their values) for current KieSession. I know it could be done by asking for a value of all the available options, but that is a bit cumbersome. However, it would work, so thanks pointing this out!	Ask solution	https://issues.jboss.org/browse/DROOLS-166
94	DROOLS-147	Hi, thanks for the reply. Couple of things i want to point out, 1) For the attached code which i had sent, i had already added the jBPM 5.4 dependency. 2) In the attached code i have commented the part wherein i am using knowledge builders, knowledgepackages and stateless knowledge sessions. The code works fine for the same scenario when i use the above mentioned classes. The problem is with the usage of PackageBuilders and that too in a peculiar way in which the rule flow is called. Initially we are calling a rule and from the rule we are invoking the rule flow using the drools.getWorkingMemory().startProcess(<rule flow id>) Kindly refer the drl and rf files attached. In the above scenario the first rule is executed successfully but the execution of 2nd rule fails (ie the call to rule flow). But the exact case works fine with the knowledge builders. Is it possible to provide a workaround / patch for the case to work with packagebuilders?	Ask possibility of solution	https://issues.jboss.org/browse/DROOLS-147
95	DROOLS-113	Is some work being done on this feature?	Ask status	https://issues.jboss.org/browse/DROOLS-113
96	DROOLS-49	So what do you suggest to do ?	Ask solution	https://issues.jboss.org/browse/DROOLS-49

IMMUTANT

#	Requirement	Question	Justification	URL
1	IMMUTANT-628	[~gantengx] - I have concerns about implementing this - all that will happen is that the call to {{stop}} will return immediately, it won't actually kill any long running jobs (unless you are killing an uberjar/{{lein run}} process with {{^CJ}}, etc, since that will kill all of the scheduler threads). How are you stopping your system currently? And are you running this as an uberjar or inside WildFly?	ask questions about implementation details	https://issues.jboss.org/browse/IMMUTANT-628
2	IMMUTANT-628	[~crawley] Ah so there's no way to actually force killing the jobs? Currently I'm running this as an uberjar using init script. When stopping the system it's basically calling {{kill}} command	ask questions about implementation details	https://issues.jboss.org/browse/IMMUTANT-628
3	IMMUTANT-580	This has been implemented, and is available in incremental build #682 and newer. Can you give it a try and confirm it works for you (see http://immutant.org/builds/2x/ for instructions on using incremental builds)? Relevant commits: https://github.com/projectodd/wunderboss/commit/0e6f802debb429d5ffb1537ecffae561fa72fc65 https://github.com/immutant/immutant/commit/be8e1009f51047dae432b7f5136a867871e23418	ask for feedback about implementation. can you give it a try.	https://issues.jboss.org/browse/IMMUTANT-580
4	IMMUTANT-568	Replying to the mailing list thread [1] here: > I think this option should be separate from `singleton`, despite the > intuitive implication, and treat it as a local-only option. I can > imagine cases where you want a job running on only one machine, but > concurrently there. [~crawley] I'm not sure I understand correctly. Do you mean that setting `allow-concurrent-exec?` to false would allow the job to execute `in parallel` on another machine but not the same machine? Intuitively I would expect the option to behave differently by default: if set the job should not execute concurrently cluster wide, thus implying `singleton` true. It's then up to the user to override by setting `singleton` to true explicitly. ...just my two cents. [1] http://immutant.markmail.org/thread/n3wnofa6dg4dcgtg	ask clarification about requirement'	https://issues.jboss.org/browse/IMMUTANT-568
5	IMMUTANT-568	This is now available in [incremental build][http://immutant.org/builds/2x/] 591 and newer. Can you give that a try and confirm it works for you?	ask for feedback about implementation. can you give it a try.	https://issues.jboss.org/browse/IMMUTANT-568
6	IMMUTANT-562	Ah, that's a good idea. Would you mind sharing an example of how you are doing that?	ask for details about implementation	https://issues.jboss.org/browse/IMMUTANT-562
7	IMMUTANT-456	Immutant uses a custom configurator itself, to return a singleton Endpoint, a feature also described at that link. So if we took in a custom configurator, we would need to override its getEndpointInstance as we do now. That might lead to surprising behavior. What if we just go ahead and make Immutant's configurator put the HttpSession and cookies in the user properties?	ask if solution is viable	https://issues.jboss.org/browse/IMMUTANT-456
8	IMMUTANT-456	Another thought: how do you feel about passing a function instead of a custom configurator? Its signature would match that of the Configurator's modifyHandshake method. So we'd introduce another option key, .handshake, which would be a fn that should expect to be passed a ServletEndpointConfig, a HandshakeRequest and a HandshakeResponse. It's more limiting, in that we wouldn't expose the other Configurator methods, but possibly less surprising, and simpler to create a function than an object, I think.	ask opinion on another solution	https://issues.jboss.org/browse/IMMUTANT-456
9	IMMUTANT-453	I've fixed Immutant to honor META-INF in resources (see IMMUTANT-455), so you no longer need to put he META-INF at the root, and it should work the same for .ima's or on-disk deployments. This will be in the latest incrementals (build #1206 or newer, #1206 is building now) - can you give it a try and confirm?	ask for feedback about implementation. can you give it a try.	https://issues.jboss.org/browse/IMMUTANT-453
10	IMMUTANT-436	Alex: The latest incremental build implements IMMUTANT-437 (which will load log4j.properties from resources/, or anywhere else on the classpath components that are within the app root) and IMMUTANT-443 (which restores log messages from Immutant namespaces when you use a custom logging configuration). Can you try the latest incremental build and make sure it works for you? You can get it via `lein immutant install LATEST`. I plan to get a 1.1.2 release out this week with these changes.	ask for feedback about implementation. can you give it a try.	https://issues.jboss.org/browse/IMMUTANT-436
11	IMMUTANT-436	Hmm, actually, since we now detect the presence of c.t.logging for IMMUTANT-443, we could force c.t.logging to use the log4j-factory. The question is: is that a bad idea? I suspect that most c.t.l users don't care which logging provider is being used unless they are providing custom logging configuration.	ask opinion on another solution	https://issues.jboss.org/browse/IMMUTANT-436
12	IMMUTANT-415	I just pushed (via http://github.com/immutant/immutant/commit/047042ac0) a function that exposes this setting: {{{immutant.jobs/set-scheduler-options {:thread-count 10}}}}. Can you see if this works for you? It will be available in incremental build #1166 (building now) or newer. Note that this function must be called before any jobs are scheduled, else it will throw an IllegalStateException.	ask for feedback about implementation. can you give it a try.	https://issues.jboss.org/browse/IMMUTANT-415
13	IMMUTANT-410	Kanwei: I tried to recreate this issue with https://github.com/tobias/immutant-410/blob/master/src/core.clj , but it's working for me. Can you share something that doesn't work, or guide me to make changes to my test app? The "Handler returned nil" error means that the ring handler fn returned nil instead of a response map. That error message could be more helpful. I'll take a look at that.	ask for more details to reproduce	https://issues.jboss.org/browse/IMMUTANT-410
14	IMMUTANT-405	Ian: Do you have any details about what header-based sessions might look like? Is it just a matter of storing a session-id in a header? How is that different than storing the session-id in a cookie header?	ask for more details to reproduce	https://issues.jboss.org/browse/IMMUTANT-405
15	IMMUTANT-390	Why is this ticket got closed? What does it mean "out of date"? Today I redeployed a simplest hello-world application war (built in clojure) on a standalone wildfly 8 (no clusters) and I was getting 404 during 10 seconds. If you think it's an upstream issue, could you link to it here?	ask ABOUT status of issue	https://issues.jboss.org/browse/IMMUTANT-390
16	IMMUTANT-359	Have we marked them as deprecated somehow? A grep of the codebase says "nope".	ask details about current implementation	https://issues.jboss.org/browse/IMMUTANT-359
17	IMMUTANT-334	What would the workaround be for this? I added {{{lein-profiles [:default]}}} to the {{{immutant}}} descriptor in my {{{project.clj}}}, but e.g. {{{target/classes}}} still isn't on the classpath. (Though, {{{target/user+dev+base/classes}}} or whatever it is isn't being created anymore...)	ask workaround	https://issues.jboss.org/browse/IMMUTANT-334
18	IMMUTANT-334	You shouldn't have seen {{{target/user+dev+base/classes}}} being generated I don't believe - were you? The issue is that since Immutant doesn't apply the default profile internally when resolving the project, but instead a subset, it expects the classes to be in {{{target/user+dev+base/classes}}}, where they actually are. I think the fix is to have Immutant always add {{{target/classes}}} to the effective classpath to cover this case, and the case where users are using a version of lein older than what Immutant uses internally.	ask details	https://issues.jboss.org/browse/IMMUTANT-334
19	IMMUTANT-334	The latest incrementals are using 2.3.2, and I modified an integ to confirm that {{{target/classes}}} is on the ECP when no profiles are applied: https://github.com/immutant/immutant/commit/6f6be5cc6567f677b5cc90a99d2a0c3639159682 . Chas: can you confirm?	ask expert opinion	https://issues.jboss.org/browse/IMMUTANT-334
20	IMMUTANT-334	Mariano: can you confirm as well? (I forgot that you were the original reporter).	ask reporter to confirm	https://issues.jboss.org/browse/IMMUTANT-334
21	IMMUTANT-329	I opted to just add the com.sun.net.httpserver.* packages instead of trying to determine a list of all 'public' com.sun packages, since the workaround is simple if more are discovered to be missing. This should be fixed by http://github.com/immutant/immutant/commit/87013ab53 and https://github.com/projectodd/wildfly/compare/56f59b98def1441727d9e797a7db57b65183724c...bd2a5fc891417c355690fe0e436b6a73b542a46 . Note that these changes will only affect the slim build of Immutant. The full build will not have these classes available, since we don't modify that build at all. Chas: can you confirm? These changes are already in the latest incremental builds.	ask expert opinion	https://issues.jboss.org/browse/IMMUTANT-329
22	IMMUTANT-326	Paul, forgive my ignorance, can you give some more details around this notion of a passive listener, please? Are multiple clients supposed to see the same message, but only one ack's it? This sounds more like a topic to me. I'm not opposed to expanding the API, but it'd be nice if it had more application than just testing production traffic without affecting it (which seems kinda risky anyway). I'm not crazy about moving from jms to core, mainly because we may be switching to activemq in the distant future, and i'd like to support both it and hornetq. I would like to make sure that whatever core HornetQ abstractions you need (managers, factories, servers, etc) are within arm's reach of you and convenient, of course.	ask for more detail	https://issues.jboss.org/browse/IMMUTANT-326
23	IMMUTANT-326	So will it help to expose the ACK mode? It's not clear to me how setting the mode to CLIENT will achieve your goal, though. I don't think the broker will deliver a message to the "real" consumer just because the passive one doesn't ack it. There's no guarantee the passive consumer won't get the same message redelivered multiple times before the real consumer gets it once (though that's unlikely). Makes sense that you can't turn your production queues into topics just to get monitoring, and I hate for you to introduce another topic just for that. You can list the queue contents via the mgmt/jmx interfaces, but it's difficult to map them to actual consumers, if that's the level of granularity you need.	ask for clarification	https://issues.jboss.org/browse/IMMUTANT-326

#	Requirement	Question	Justification	URL
24	IMMUTANT-317	I looked into what it would take to do this at deploy-time inside Immutant, and it's currently not straightforward, because: 1) we explicitly strip any Immutant deps from the project before passing them to {{resolve-dependencies}}. If we wanted to get the pedantic output for those deps, we would have to call {{resolve-dependencies}} twice, which is time consuming 2) most apps won't have any explicit Immutant deps, so we would need to add them before passing the project to {{resolve-dependencies}}. That would be possible for versioned releases, but wouldn't work for incremental builds, since the jar artifacts won't be available. 3) {{leinigen.core.classpath}} currently memoizes the pedantic output, which means an app that triggered warnings would not do so on a redeploy, nor would a second app that happens to have the same dependencies. Putting this check in the plugin may be a better option, though we would still have the same issue with incremental builds. Maybe we could just support this for versioned builds? I'll give it some more thought.	ask opinion	https://issues.jboss.org/browse/IMMUTANT-317
25	IMMUTANT-287	Is this still true (from web.org)? bq. NOTE: As of this writing, Pedestal route tables are not context-aware, so unless you deploy them to the root context, you must hard-code the context path in your routes.	ask about details	https://issues.jboss.org/browse/IMMUTANT-287
26	IMMUTANT-287	cache.org: the default locking mode is nil, but the description states: "One of :optimistic or :pessimistic". so what does nil mean?	ask for clarification	https://issues.jboss.org/browse/IMMUTANT-287
27	IMMUTANT-287	jboss.org: do we need to update the logging section?	ask for clarification	https://issues.jboss.org/browse/IMMUTANT-287
28	IMMUTANT-219	Should this be resolved as 'Done'? We didn't actually add scheduled delivery to publish.	ask ABOUT ticket status	https://issues.jboss.org/browse/IMMUTANT-219
29	IMMUTANT-195	Pragmatically, I'm guessing that filling a bug with tomcat is not likely going to be profitable? (the code is very clearly wrong by design, to deal with a flaw in another part of the system, which is undoubtably even harder to fix) Not that I have any strong desire to do this, but I don't suppose it's possible to easily replace CoyoteAdapter with subclass that works?	ask for details	https://issues.jboss.org/browse/IMMUTANT-195
30	IMMUTANT-195	I've turned on ALLOW_BACKSLASH by default (https://github.com/immutant/immutant/commit/218b85bd4af809e22eb0d44136b03799459208fb). Can you confirm that the latest incremental works for you? If so, I'll resolve this issue.	ask for feedback about implementation. can you give it a try.	https://issues.jboss.org/browse/IMMUTANT-195
31	IMMUTANT-182	It might be worthwhile to do this by implementing one proper Chef/Pallet/Ansible/Puppet script; this is both practical and documents best practices for configuring a cloud node to be part of a cluster, would include init scripts for common init systems, also EC2-specific stuff for S3PING or AWSPIING, logging, debugging, remote nREPL connectivity from a development system, etc. Ideally, if Ansible (my favorite), also shows how to do cluster-wide updates on EC2 through adding domain mode in Immutant or more likely a script to do rolling upgrades across the cluster. Other questions: - Is upgrading an app on an AMI better / worse than just doing a DNS swap across an entire cluster? - Maybe there is a good balance between power cycling an app, immutant, or an entire AMI/Node?	ask opinion about solutions	https://issues.jboss.org/browse/IMMUTANT-182
32	IMMUTANT-107	Why? If we include deps in archives by default, why would we resolve them by default at deployment? It's too much of a pain to manually create the deployment descriptor just to turn :resolve-dependencies to false.	ask motivation so a better description	https://issues.jboss.org/browse/IMMUTANT-107
33	IMMUTANT-96	Another reason to consider this change is that using '/' in queue/topic names will make it difficult to use HornetQ REST interface to JMS. See Section 3.1 of http://docs.jboss.org/resteasy/hornetq-rest/1.0-beta-1/userguide/html_single/index.html#d0e162 . The URI for a topic/queue resourc would be /queues/jms.queue.{immutant_queue_name} Using '/' clearly doesn't play well in this scheme. Is there a strong reason to continue with '/', and not simply use '.' or some other separator?	ask opinion	https://issues.jboss.org/browse/IMMUTANT-96
34	IMMUTANT-96	Right now, I believe the only naming requirement we have is that queues must start with /queue, and topics must start with /topic. We do that so we can recognize the type based on a naming convention. I'm fine with relaxing that, maybe to #^A.?queue". This would make queues starting with "/queue", ".queue", "queue", "xqueue" all legal. How does that sound?	ask opinion	https://issues.jboss.org/browse/IMMUTANT-96
35	IMMUTANT-84	I pushed an update to lein-immutant today (0.8.2) that should resolve this. Can you try it and report back that it indeed works for you?	ask for feedback about implementation. can you give it a try.	https://issues.jboss.org/browse/IMMUTANT-84
36	IMMUTANT-78	Is this one still relevant?	ask opinion on relevance	https://issues.jboss.org/browse/IMMUTANT-78
37	IMMUTANT-55	What does this even mean?	ask for details	https://issues.jboss.org/browse/IMMUTANT-55
38	IMMUTANT-27	Jim: Isn't this kind of done? Or are is this jira waiting on more formal Immutant support on OS?	ask status	https://issues.jboss.org/browse/IMMUTANT-27
39	IMMUTANT-24	We already had this discussion in TB about 'at' API design. So I wrote a draft in Clojure. https://gist.github.com/184f7b2eb203df87727a Lemme know if I got the idea.	ask opinions	https://issues.jboss.org/browse/IMMUTANT-24
40	IMMUTANT-21	Hey Jimbotron - is there anything else we need to do here?	ask status	https://issues.jboss.org/browse/IMMUTANT-21
41	IMMUTANT-21	Hey Jimbotron - is there anything else we need to do here?	ask status	https://issues.jboss.org/browse/IMMUTANT-21

JBTM

#	Requirement	Question	Justification	URL
1	JBTM-2905	This class was added to demonstrate how to support new log record types in the tooling. The bug is that it was included in the default list of types. What is your plan to fix this, are you going to remove it from the default list?	Ask resolution	https://issues.iboss.org/browse/JBTM-2905
2	JBTM-2905	[~mmusgrov] ok, that could be my wrong in understanding the issue here. Without the header reader I got troubles to get working ExpiredAssumedCompleteScanner. But I could take it from a wrong end. In that case, this change would be superfluous and the fix should go in opposite way - removing from the default list and leaving header reader in the test package. Can you please elaborate a bit on what is the correct way of fixing it? Thanks in advance!	Ask resolution	https://issues.iboss.org/browse/JBTM-2905
3	JBTM-2905	I may have misunderstood something. I thought this was required to fix a JBEAP issue? Straight up reverting isn't right because the ObjStoreBrowser (in main) shouldn't depend on a test class. But we could make it so its not in the default list. [~mmusgrov] what do you think?	Ask req clarification	https://issues.iboss.org/browse/JBTM-2905
4	JBTM-2835	What does the error look like with your patch? What is the effect of the NPE?	Ask req clarification	https://issues.iboss.org/browse/JBTM-2835
5	JBTM-2822	Hi [~tomjenkinson], during the check ups JBEAP issues I got to this issue and when applying the same test as I have for JBTM-2636 I found out that there is a branch in code where deferred exception is not added - that's when heuristic happens. Is it intentional? This kind of update fixes the issue that I refer to: https://github.com/ochaloup/narayana/commit/cb2ab4252a6c2e1106751d28103daefea42e86f6	Ask code difference between branches	https://issues.iboss.org/browse/JBTM-2822
6	JBTM-2822	Sorry Ondra - I don't quite follow. Which branch does not have this commit in?	Ask code difference between branches	https://issues.iboss.org/browse/JBTM-2822
7	JBTM-2822	Ah yes, that sounds useful. Perhaps you can add a unit test to ExceptionDeferrerTests and we can look to merge it? This is the type of test to add: https://github.com/jbosstm/narayana/pull/1116/files I think you should raise a new issue though because it will be confusing to track the change against the resolved one.	Ask code	https://issues.iboss.org/browse/JBTM-2822
8	JBTM-2733	Hi [~tomjenkinson], I was wondering on your comment https://issues.iboss.org/browse/JBTM-2733?focusedCommentId=13284121#comment-13284121 and now I found some time to go through the XA spec for it. I wasn't able to find reason behind why the recover sequence should be {{xa_recover(START), xa_commit/rb, xa_recover(END)}} and can't be e.g. like {{xa_recover(START END), xa_commit/rb}}. Would you be so kind and elaborate a bit on this? Or can you point me to spec/article where I got information how this is expected to be used? It's quite possible that I just haven't read/searched properly. I'm just curious. Thanks in advance.	Ask req clarification	https://issues.iboss.org/browse/JBTM-2733
9	JBTM-2686	Where is the link to the PR?	Ask code (PR)	https://issues.iboss.org/browse/JBTM-2686
10	JBTM-2660	Good work Amos, I wonder if we can upgrade the SPI org.jboss.spec.javax.resource:jboss-connector-api_1.5_spec:jar:1.0.0.Final dependency to a later one that includes JTA 1.2?	Ask project dependency	https://issues.iboss.org/browse/JBTM-2660
11	JBTM-2658	Hi Alexis, I think you might be in the wrong place, this is the transaction manager component not the JMS component? Maybe you could talk to the Artemis team? Unless I misunderstood but in which case I recommend a discussion here first: https://developer.jboss.org/en/jbosstm/content?filterID=contentstatus%5bpublished%5d~objecttype~objecttype%5bthread%5d	Suggest them report this issue to another project	https://issues.iboss.org/browse/JBTM-2658
12	JBTM-2634	Can you ask for feedback on the forum now it seems complete.	Suggest the discussion in forum	https://issues.iboss.org/browse/JBTM-2634
13	JBTM-2617	Feel free to discuss in the forum if you would like some more feedback - not sure which class you are referring to.	Ask req clarification	https://issues.iboss.org/browse/JBTM-2617
14	JBTM-2606	Can we get a link to your branch here please?	Ask code	https://issues.iboss.org/browse/JBTM-2606
15	JBTM-2605	Is this done?	Ask status	https://issues.iboss.org/browse/JBTM-2605
16	JBTM-2599	[~stalep] I have no problem with doing a release with this in but I think Narayana will use the one in WFLY anyway so you could just upgrade the version of the journal in there?	Ask upgrate version information	https://issues.iboss.org/browse/JBTM-2599
17	JBTM-2599	[~lbarreiro] - understood. Will do the release as the PR passed. It will be tomorrow AM (UK) though. I don't think characterising this as a bug (in Narayana) is correct so changing the issue type. Also, please can you add a description? Probably just replicate your comment https://issues.jboss.org/browse/JBTM-2599?focusedCommentId=13145088&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-13145088	Ask req clarification	https://issues.iboss.org/browse/JBTM-2599
18	JBTM-2593	OK, so rather than remove this and revert back to what we had for 20+ years, I've made it a trace. There may have been a reason why it was needed, e.g., debugging? Would be good for someone to link any related issues here for the future.	Ask retional/use case of req	https://issues.iboss.org/browse/JBTM-2593
19	JBTM-2564	So do I need to add those to all the quickstart poms?	Ask completeness of solution	https://issues.iboss.org/browse/JBTM-2564
20	JBTM-2564	Because if so we need to do it in a profile at a guess. Unless those arguments work on an IPv6 box. Do we know which quickstarts need those? [~hhovsepy] - can you raise a PR if you have the fix on your box so we can check it works on our CI cluster too please?	Ask code (PR)	https://issues.iboss.org/browse/JBTM-2564
21	JBTM-2564	Hi [~hhovsepy] - I thought Mike said you had tested a fix locally? If so, please can you commit it and PR it and I can amend it as necessary? Alternatively just paste in the surefire config you tested as a comment and I will try it here.	Ask code (PR)	https://issues.iboss.org/browse/JBTM-2564

#	Requirement	Question	Justification	URL
22	JBTM-2564	http://albany.eng.hst.ams2.redhat.com/job/narayana-quickstarts-ipv6/ ?? I didn't even think we had dualstack machines in ams2? If so we should move the dualstack job onto ams2 cluster instead of ncl one.	Ask req clarification	https://issues.jboss.org/browse/JBTM-2564
23	JBTM-2564	Did you need any fixes to Narayana to make it work on IPv6 because I think [~hhovsepy] is reporting they don't work for him? I think the reason they would work on IPv6+IPv4 is arq will be fine to run with its IPv4 args. I therefore don't think we can mark it as resolved because the QS are running on IPv6 with IPv4 too so that is difference to the summary+description of this issue.	Ask necessary of req	https://issues.jboss.org/browse/JBTM-2564
24	JBTM-2564	As I understood, this task was created to add new maven profile into several quickstarts to pass IPv6 parameters (mentioned in earlier comments), right? Is it already done?	Ask req clarification & status	https://issues.jboss.org/browse/JBTM-2564
25	JBTM-2553	Any more details on this?	Ask req clarification	https://issues.jboss.org/browse/JBTM-2553
26	JBTM-2553	I did some debug with the CrashRecovery2 Run with the following command {code} ./build.sh -f ArjunaJTA/jta/pom.xml clean test -Dtest=CrashRecovery2 -Dorg.jboss.byteman.debug -Dorg.jboss.byteman.verbose -PcodeCoverage {code} check the output of the test and it looks like the rule does not get triggered as it can not find the local binding var "\$this" The fail2pc rule is {code} RULE Fail 2PC CLASS com.arjuna.ats.arjuna.coordinator.BasicAction METHOD phase2Commit AT ENTRY BIND NOTHING IF readCounter("phase2commit") == 0 DO debug("Called"), incrementCounter("phase2commit"); com.arjuna.ats.arjuna.coordinator.ActionManager.manager().remove(\$this.get_uid()); RETURN ENDRULE {code} The failure is {code} start:4 triggerPos:0 end:535 (I add this debug message in agent/src/main/java/org/jboss/byteman/agent/adapters/RuleCheckMethodAdapter.java) RuleCheckMethodAdapter.checkBindings : invalid local variable binding \$this checking method phase2Commit(Z)V {code} if we run without the jacoco, the start and end is {code} start:0 triggerPos:0 end:327 {code} [~adinn] how do you think about this ?	Ask req clarification	https://issues.jboss.org/browse/JBTM-2553
27	JBTM-2553	Hi Amos, thanks for reporting this. I don't know why you are seeing that error. It's very odd because \$this is the only local variable reference that we can identify when there is no local variable table. Could you try changing the rule to use \$0 instead of \$this and see if that makes it work.	Suggestion	https://issues.jboss.org/browse/JBTM-2553
28	JBTM-2553	Thanks Adinn, I met this error when running the codeCoverage with the jacoco. The \$0 works. So it looks like the issue is related to the script but not the jacoco ?	Ask req clarification	https://issues.jboss.org/browse/JBTM-2553
29	JBTM-2553	Doh! I see Amos is way ahead of me here (nice work, Amos :-). The trace statement Amos inserted shows that the local variable checking code has found an entry in the local var table for \$this but one that starts at bytecode offset 4 and ends at offset 535. The injection code is trying to inject the rule at offset 0 i.e. at entry to the code. This is after transformation by Jacoco. With no transformation the local var this is found starting at offset 0. I have not yet looked at the Jacoco code but it seems pretty obvious what is happening. When Jacoco inserts profiling instructions into the method it updates the start and end offsets for entries in the local variable table to take account of the new code. One of those profiling instruction sequences will be injected at the start of the method and Jacoco appears to be shifting the local var start offset for this so that it comes into scope after the profiling code rather than at bytecode 0. No doubt it is also shifting the start position for method parameter entries. This is at best unhelpful; actually, I think it is an error because this and the parameters are /always/ in scope from the start of the method even if code is inserted before the start. I thought that ASM fixed this as a special case. Which version of Jacoco is being used here? The latest Jacoco version uses ASM 5 which should deal with this problem. Have you tried updating to a later Jacoco? If that doesn't work I'll take this up with the Jacoco developers.	Ask req clarification	https://issues.jboss.org/browse/JBTM-2553
30	JBTM-2553	Thanks for the detailed explanation Andrew - it certainly sounds a complex issue. Do you know of any other coverage frameworks that Byteman might interoperate with today? If not we can just disable the byteman tests under coverage testing. If you do get a Jacoco issue reference please can you let us know so we can reference it here? Thanks again!	Ask external resource & RIM	https://issues.jboss.org/browse/JBTM-2553
31	JBTM-2532	Did you put back a commit for the unit tests for now?	Ask code (testing)	https://issues.jboss.org/browse/JBTM-2532
32	JBTM-2522	I meant https://github.com/jbosstm/performance/tree/master/comparison , but I'm not sure it's going to work. Is it possible to use JMH to test Arquillian tests? Or applications in WildFly deployed manually?	Ask feasibility of solution	https://issues.jboss.org/browse/JBTM-2522
33	JBTM-2496	Thanks Amos - what is missing to make it work with WildFly 10, please?	Ask reason of req(problem)	https://issues.jboss.org/browse/JBTM-2496
34	JBTM-2495	The failures related to the upgrading to the activemq artimes in WFLY 10. Also in the Step 5, you should run with standalone-blacktie.xml {code} copy docs/examples/configs/standalone-blacktie.xml wildfly<version>standalone/configurations bin/standalone.sh -c standalone-blacktie.xml {code} I'm sorry that I can not find in the QS to say use WFLY 9 but I will update the README.md to say it should integrate with the master from WFLY. can you show me where you find the information to run with WFLY 9 ?	Ask external resource	https://issues.jboss.org/browse/JBTM-2495
35	JBTM-2490	Hi Hayk, Would you have chance to work on this at all? You have marked it as a blocker so trying to work out the priority as you didn't mark a fix version. Thanks, Tom	Ask status	https://issues.jboss.org/browse/JBTM-2490
36	JBTM-2481	Hi [~gytis], this has been Coding in Progress for a while. Is it worth marking "Work Stopped" as I don't think you are currently coding on this one?	Ask status	https://issues.jboss.org/browse/JBTM-2481
37	JBTM-2481	Are we not planning to make the release completely automatic and incorporating it into our CI?	Ask req clarification	https://issues.jboss.org/browse/JBTM-2481

#	Requirement	Question	Justification	URL
38	JBTM-2464	Can this be closed?	Ask status	https://issues.iboss.org/browse/JBTM-2464
39	JBTM-2458	Is this done?	Ask status	https://issues.iboss.org/browse/JBTM-2458
40	JBTM-2458	What was the resolution?	Ask solution	https://issues.iboss.org/browse/JBTM-2458
41	JBTM-2440	Do you know why these only seemed to start failing when we upgraded to JDK8?	Ask reason of req (problem)	https://issues.iboss.org/browse/JBTM-2440
42	JBTM-2440	Hi Mike, Am I right in thinking that this should be fixed, but we still have JBTM-2454? If so it is worth closing this one. Thanks, Tom	Ask status	https://issues.iboss.org/browse/JBTM-2440
43	JBTM-2432	Is there anything in Narayana that requires Java8 though? If not then why not still allow that at build time?	Compare platform	https://issues.iboss.org/browse/JBTM-2432
44	JBTM-2432	That's WildFly and I do understand why they're moving to 8. However, if someone wants to build a component project using 7 then why wouldn't we allow that within that project until/unless we add features only available in 8?	Suggestion	https://issues.iboss.org/browse/JBTM-2432
45	JBTM-2429	I am not sure how easy this will be but is it possible to leave it as a system property where we can enable it? e.g. -DenableDebugger	Ask feasibility of solution	https://issues.iboss.org/browse/JBTM-2429
46	JBTM-2414	Yesterday I pushed the fix that does not upgrade arquillian so we are good to go on our CI builds. Do you think the jira still warrants the blocker status? Mike -- Michael Musgrove Red Hat UK Ltd, Transactions Team e: mmusgrov@redhat.com t: +44 191 243 0870 Registered in England and Wales under Company Registration No. 03798903 Directors: Michael Cunningham (US), Charles Peters (US), Matt Parson (US), Michael O'Neill(Ireland)	Ask RIM	https://issues.iboss.org/browse/JBTM-2414
47	JBTM-2412	Thanks Tom, Which component would that be. Tom	Ask component	https://issues.iboss.org/browse/JBTM-2412
48	JBTM-2396	I see the test passing now - thanks for the quick update :) re the LockException misconception I had, I now understand why throwing an exception is wrong here, though I wonder if it is worth trace logging something to help clarify if a TX is rolling back due to OCC conflict?	Ask benefit and cost of req	https://issues.iboss.org/browse/JBTM-2396
49	JBTM-2393	[~tomjenkinson], do we have analytics id etc.?	Ask code	https://issues.iboss.org/browse/JBTM-2393
50	JBTM-2390	I am surprised you think its an anomaly. As an example, when you turn volatile store on our performance is something much more than a 15 fold improvement irc? Some suggestions: 1. Are you using HQ store (can you use it in this config)? 2. Have you tried enabling the volatile store? In terms of notx vs emptytx I expect that is all down to the creation of the TX endpoints. You could profile the server to get more details there. I imagine your participants are no-ops so in the no tx case you have 2 no-op calls vs the empty tx case where you have 3 no-ops (I will class commit() here as a no-op as I assume it really is very little) and 1 op (the creation of the tx endpoints). Seems reasonable to me. In the emptytx case, can you drop the commit calls (without it falling over) to see the cost of the commit vs the create?	Suggestion of solution	https://issues.iboss.org/browse/JBTM-2390
51	JBTM-2379	Is it OK to raise this one now that 2377 is merged thanks?	Ask RIM	https://issues.iboss.org/browse/JBTM-2379
52	JBTM-2379	Great, please can you let me know when the PR is merged so I can start the 5.1.1 release?	Ask code (PR)	https://issues.iboss.org/browse/JBTM-2379
53	JBTM-2375	I think a link to github is probably OK. Maybe do it by tag?	Discuss solution	https://issues.iboss.org/browse/JBTM-2375
54	JBTM-2366	EDIT: The first link will need to be viewed using: https://drive.draw.io/?#G0B9c6Ahw-picmQkIPZxp5M19nalk as the text does not seem to be visible on the Google Docs preview ----- Gytis has a draft, I have shared it publicly but its still very much in development: https://drive.google.com/file/d/0B9c6Ahw-picmQkIPZxp5M19nalk/view?usp=sharing For comparison, this is the existing version: http://narayana.io/images/narayana-architecture-diagram.jpg I think we need to align the cosmetics with other JBoss.org projects. Are isometric boxes still the most graphically clear?	Ask external dependency	https://issues.iboss.org/browse/JBTM-2366
55	JBTM-2358	We should try to get the best JDBC results we can so should try to use a database that is either on the same machine ideally. Does haverstraw host a postgres?	Ask external dependency	https://issues.iboss.org/browse/JBTM-2358
56	JBTM-2357	Hi Toby, we can't change it to look for jbossts-properties.xml solely, the reason we can't remove the check for arjuna-properties.xml is due to backwards compatibility. The work in JBTM-482 was about changing the file name inside JBoss really which was completed. As such, I see this as an enhancement. Would allowing a system property to override the default name of the file work for you?	Discuss solution	https://issues.iboss.org/browse/JBTM-2357
57	JBTM-2345	Would we need a variant of WFLY-1119 (Assign an unique NodeID automatically) to avoid conflicts (ie would different images share the same JDBC store). Also it is possible to point the tooling (ObjStoreBrowser) at a filesystem to manage the logs although it is unsupported (normally it is done via a JMX browser connected to the TM process). Do we need similar functionality for accessing the JDBC store directly?	Ask new req	https://issues.iboss.org/browse/JBTM-2345
58	JBTM-2343	How is this different to the current class that allows thread-to-transaction association and disassociation?	Ask difference of existing and commit code	https://issues.iboss.org/browse/JBTM-2343
59	JBTM-2334	Does that address the same use case though? If not then we should re-examine what we're doing.	Ask completeness of solution	https://issues.iboss.org/browse/JBTM-2334
60	JBTM-2318	Agreed. Maybe create a separate JIRA for the tooling?	Ask new req	https://issues.iboss.org/browse/JBTM-2318

#	Requirement	Question	Justification	URL
61	JBTM-2318	Yeah, I haven't upgraded the logging for the component. The only thing I can think is that maybe the default used to be WARN or something? From looking at the log statements there looks to be three statements, its the way the log messages are formatted that is causing each of the statements to bridge two lines. The logging format looks to be the format generated by JDK logging rather than log4j so I guess this is in an IDE? {code} 2014-12-18 16:30:49,755 [main] INFO (ats.arjuna :88) - ARJUNA012163: Starting service com.arjuna.ats.arjuna.recovery.ActionStatusService on port 51214 2014-12-18 16:30:49,765 [main] INFO (ats.arjuna :327) - ARJUNA012337: TransactionStatusManagerItem host: 127.0.0.1 port: 51214 2014-12-18 16:30:49,765 [main] INFO (ats.arjuna :138) - ARJUNA012170: TransactionStatusManager started on port 51214 and host 127.0.0.1 with service com.arjuna.ats.arjuna.recovery.ActionStatusService {code} I am happy for those statements to move down to debug but the reason they are info is the default option seems to be to choose a random port so it might be the case then that people would all enable debug just to get at that. I can see why with multiple processes on the same machine why electing a port dynamically has some advantages but I know that WFLY at least configure a default of 4713. (4712 for recovery port). Why don't we use those defaults and only then log the statement at INFO if the user changes the port to 0?	Ask req clarification	https://issues.jboss.org/browse/JBTM-2318
62	JBTM-2318	Can we also add the following to the list of messages (relating to recovery) that should be demoted to debug level: {code} ARJUNA012324: Start RecoveryActivators ARJUNA012343: RecoveryManagerImple is ready. Socket listener is turned off. ARJUNA012296: ExpiredEntryMonitor running at {0} ARJUNA012332: Failed to establish connection to server ARJUNA012304: Removing old transaction status manager item {0} {code}	Ask req clarification	https://issues.jboss.org/browse/JBTM-2318
63	JBTM-2299	It looks like we don't need it as we have removed the TAO dependency in JBTM-1665 and used the http rest methods to access the TM [~tomjenkinson], Does it make sense ?	Ask req clarification	https://issues.jboss.org/browse/JBTM-2299
64	JBTM-2296	Here are the updated name server and transaction service images: https://github.com/Gytis/narayana-docker/tree/master/jts I've added extra configurations and a couple of unit tests. And here is the quickstart: https://github.com/Gytis/quickstart/tree/master-JBTM-2296/jts-docker All the code is the same. However, readme was extended to explain how use different types of object stores. Any ideas?	Ask review solution	https://issues.jboss.org/browse/JBTM-2296
65	JBTM-2296	[~tomjenkinson], you had a go with it some time ago. Do you think it's usable enough to be merged into our dockerfiles repo?	Ask usability of req	https://issues.jboss.org/browse/JBTM-2296
66	JBTM-2286	Two questions: 1) is this being driven by any customer/community request? 2) presumably the javadocs generated will always depend upon the build you perform, e.g., JTS doesn't get covered if I'm only building core?	Ask req clarification	https://issues.jboss.org/browse/JBTM-2286
67	JBTM-2280	Suggestimate suggests that this is similar to JBTM-2225, which appears to be true, no?	Ask RIM	https://issues.jboss.org/browse/JBTM-2280
68	JBTM-2280	[~mmusgrove]: would this be enough to get rid of JNDI for us entirely then?	Ask completeness of solution	https://issues.jboss.org/browse/JBTM-2280
69	JBTM-2280	So this was closed and then fixed?	Ask status	https://issues.jboss.org/browse/JBTM-2280
70	JBTM-2279	Wouldn't it have been easier/better/less invasive to have hidden this within a call to something like getCheckedActionFactory()? _checkedAction = _checkedActionFactory.getCheckedAction(get_uid(), type()); + _checkedAction = allowCheckedActionFactoryToBeOverridden ? arjPropertyManager + .getCoordinatorEnvironmentBean().getCheckedActionFactory() + .getCheckedAction(get_uid(), type()) + : defaultCheckedActionFactory.getCheckedAction(get_uid(), + type());	Suggestion	https://issues.jboss.org/browse/JBTM-2279
71	JBTM-2270	JTS participant records are not in-lined with the transaction log so they appear in a separate part of the log store. The tooling only shows participants for JTA records which do have their participants inlined. Ondrej, is the Fix version field (6.0.0) ok? I just noticed that Gytis is down as the assignee so I will let him update that field.	Ask fixed version	https://issues.jboss.org/browse/JBTM-2270
72	JBTM-2264	Tom, great write up. I suggest changing this issue from a bug to something else (feature request?)	Ask change type of issue	https://issues.jboss.org/browse/JBTM-2264
73	JBTM-2262	BTW what did these scripts do in the old days? Did they add the jacorb jar explicitly?	Ask function of existing code	https://issues.jboss.org/browse/JBTM-2262
74	JBTM-2259	One of the reasons we haven't done this to date is that it's not standards compliant. Why isn't the ordering that JTA allowed in the interposed option sufficient?	Ask reason of req	https://issues.jboss.org/browse/JBTM-2259
75	JBTM-2245	Hi Tom, Mike, What is the planned fixed version of this task? Is there already a support of shutdown(timeout) in Narayana? I am asking as we have RFE https://issues.jboss.org/browse/EAP7-86 resolved and ready for QA. Thanks, Hayk.	Ask status & time	https://issues.jboss.org/browse/JBTM-2245
76	JBTM-2231	Hi Jesper, Please can you give me some ideas of what this for? Is it the IJ tracer? I have no problem with adding it, please just provide some context as there could be more I can give you. Thanks, Tom	Ask reason/use case of req	https://issues.jboss.org/browse/JBTM-2231
77	JBTM-2231	Hi Jesper, Is the attached PR what you were thinking? Thanks, Tom	Ask req clarification	https://issues.jboss.org/browse/JBTM-2231
78	JBTM-2230	As mentioned on the call, we need to be mindful of when we introduce this new interface method as it will require implementations to add the new method when they wish to upgrade. We could add it to a new interface. As there are various options here Mike we should create a design discussion and mention it to various teams. Can you create the discussion and flag it to Jesper (and Andy) to make sure known XARW impls are aware of the requirements (Andy obviously knows but he might want to contribute if Jesper has any observations.	Ask external/offline discussion	https://issues.jboss.org/browse/JBTM-2230
79	JBTM-2224	Can you point me at the forum posting where this was discussed? Then I can take my comments there rather than in JIRA.	Ask external resource (forum)	https://issues.jboss.org/browse/JBTM-2224
80	JBTM-2222	Forum discussion? Could be around, I just can't find it.	Ask external resource (forum)	https://issues.jboss.org/browse/JBTM-2222

#	Requirement	Question	Justification	URL
81	JBTM-2213	I think this should be 5.next version?	Suggestion version of req	https://issues.jboss.org/browse/JBTM-2213
82	JBTM-2213	That's what me and Mike where not sure about yesterday. So do I simply put all the stuff to 5.next which I'm working on? Or do I work on 5.later and then we move them to 5.next during the planning?	Ask fixed version of req	https://issues.jboss.org/browse/JBTM-2213
83	JBTM-2210	Granted I haven't updated my local copy for a few weeks, but these warnings seemed to come all of a sudden and a lot of them to boot. Was this expected?	Ask function of existing code/implementation	https://issues.jboss.org/browse/JBTM-2210
84	JBTM-2210	So all of these warnings have been there for 7 months?!	Ask req clarification	https://issues.jboss.org/browse/JBTM-2210
85	JBTM-2210	The reason its not been detected is because most of our builds are done in the cloud or (certainly for myself) an IDE. If you think its an issue, maybe we should get the build to fail on warning? http://jira.codehaus.org/browse/MCOMPILER-120 is available, so its just a case of enabling -Werror.	Suggestion	https://issues.jboss.org/browse/JBTM-2210
86	JBTM-2210	No, failing on warn is probably not a good idea. I'm surprised that I didn't see it earlier, since I built the source quite a bit only a few weeks ago. I'm also surprised no one else caught it if it's been around that long. Do we regularly check the build logs? Even building in the cloud or an IDE should generate the same warnings (e.g., in Eclipse).	Ask req clarification	https://issues.jboss.org/browse/JBTM-2210
87	JBTM-2209	Hi Jesper, Please can you add a few more details to the Jira on the reasons for the change? Tom	Ask req clarification	https://issues.jboss.org/browse/JBTM-2209
88	JBTM-2209	Thanks, but that's about as much use to me as a chocolate fireguard! ;) Or giving me a scalar number when I need a vector. What was the test used? How many iterations? We could do this ourselves of course, but it'd be so much simpler if you've already done that work :-). And yes, I see the test code you added, but I was looking for some test that worked in situ with the Stack usage within Narayana.	Ask code (testing)	https://issues.jboss.org/browse/JBTM-2209
89	JBTM-2206	I don't think we should change ArjunaCore/arjuna/classes/com/arjuna/ats/arjuna/coordinator/CheckedAction.java as it is in the public API so unless it has considerable performance win the cost of update apps to use the new API is probably not justified. As to the rest of the changes, can you give us a brief indication of the performance benefit of this particular change?	Ask reason of existing code/implementation	https://issues.jboss.org/browse/JBTM-2206
90	JBTM-2206	Mike, using our performance toolkit please can you provide some details on here about the performance improvement that we see with the commits applied? I am trying to make a decision on whether its worth having to inconvenience users by deprecating the checked action API and replacing it with the updated one.	Ask benefit of req	https://issues.jboss.org/browse/JBTM-2206
91	JBTM-2206	So no perceivable difference at all?	Ask code changes	https://issues.jboss.org/browse/JBTM-2206
92	JBTM-2206	Can you include a link to the forum posting here for future reference?	Ask req outside discussion	https://issues.jboss.org/browse/JBTM-2206
93	JBTM-2203	Hi Jesper, Thanks for creating the separate issues! Can you now share the relative performance benefit of the enhancement please, same for JBTM-2209, JBTM-2207, JBTM-2208 and JBTM-2206? Thanks! Tom	Ask benefit of req	https://issues.jboss.org/browse/JBTM-2203
94	JBTM-2203	I don't think we can remove the *Mutex() methods from StateManager as these are part of the public API so people that inherit from that class would no longer have access to them (as you saw in the tests you had to modify). Please can you provide a few more details on what the performance win was with removing the Reentrant lock to help evaluate the benefit vs the drawback of the API change?	Ask req clarification	https://issues.jboss.org/browse/JBTM-2203
95	JBTM-2203	Maybe delay creation of lock until first use? Might help the performance issue, whatever that is.	Suggestion solution	https://issues.jboss.org/browse/JBTM-2203
96	JBTM-2203	Hold on ... doesn't this changed the public API of StateManager and LockManager? Whilst the latter is a superset of the previous API, the former isn't and potentially breaks other classes from users that may inherit from StateManager. Did we deprecate as discussed last July? Did we determine what the cost/benefit is of doing this?	Ask req clarification	https://issues.jboss.org/browse/JBTM-2203
97	JBTM-2202	Presumably you mean what we've discussed in the past: having the ability to bypass the SOAP layers and go directly to the underlying Arjuna classes which they wrap?	Ask req clarification	https://issues.jboss.org/browse/JBTM-2202
98	JBTM-2202	What does this have to do with REST transactions? If it's local-only, then distribution over REST is not an option.	Ask req clarification	https://issues.jboss.org/browse/JBTM-2202
99	JBTM-2189	For clarity in future if others come to this issue, can you please describe what problem you've seen and hence why you've submitted a patch. "fails to bootstrap properly" isn't sufficient. Thanks.	Ask req clarification	https://issues.jboss.org/browse/JBTM-2189
100	JBTM-2189	Hi Chris, Many thanks for your interest in Narayana. Normally for something like this it is best to start with a discussion on our forums over here: https://community.jboss.org/en/jbosstm . That said, please can you add a little more detail into the description as to the nature of the bug. We do already function in Java SE and also complex classloader hierarchies so this may be a corner case which we would be happy to work with you to address. I am not familiar with OneJar, is it a maven plugin that will put all the dependency jar classes into the same Jar as the application jar? The main problem with your report as it stands is that we don't have a well defined reason for _why_ the change needs to be made, if you can express the motivation for the change in more general terms then we can get it merged. As you say, the fix looks like it should work in our currently supported options so it shouldn't be a problem to get it merged once we have more context to backup the need to merge it. Once we understand the problem better, and we proceed with your approach, please can you: 1. Sign the CLA over here: https://cla.jboss.org/ (project is JBoss Transactions) 2. Fork: https://github.com/jbosstm/narayana 3. Raise a pull request on the main repo with your fix in. It should be easier to review inline, plus it will get CI run on it. Many thanks for the contribution! Tom	Ask req clarification	https://issues.jboss.org/browse/JBTM-2189

#	Requirement	Question	Justification	URL
101	JBTM-2184	What's the procedure for ensuring our artifacts make it into maven central? How are you determining our missing dependencies?	Ask solution	https://issues.jboss.org/browse/JBTM-2184
102	JBTM-2184	Well, the missing dependency bit I qualify simply based on importing "org.jboss.narayana.jta:narayana-jta"; if this has classes that call out into jboss-transaction-spi and that's not a transitive dependency of narayana-jta, this is missing, correct? Unless you have some mechanism in place where some code paths have optional dependencies? So: *if* my issue above is valid and jboss-transaction-spi should be a transitive, runtime (or compile-time) dependency of narayana-jta, I'd suggest we add it as Step 1. And regarding Maven Central: again *if* this issue is valid, then any dependencies of a project in Central must also be in Central, as this remote repo is intended to be a complete repository. That's according to Sonatype rules and, if I'm remembering correctly, should be enforced by our Nexus installation when we go to promote/release builds.	Ask req clarification	https://issues.jboss.org/browse/JBTM-2184
103	JBTM-2184	So my guess is that we need to add the dependency to ArjunaJTA/narayana-jta/pom.xml - what is the easiest way to test that? I ran tattletale over narayana-jta and it reports transitive dependencies on other stuff as well (org.hornetq.core and org.jfree). Does that imply we need to add explicit dependencies on those too?	Ask completeness of solution	https://issues.jboss.org/browse/JBTM-2184
104	JBTM-2184	Just trying to work out, what were you doing to cause the issue? For example the ArjunaJTA/maven qs doesn't exhibit this behaviour after I fixed it to have the correct runtime dependencies. Some of the quickstarts do import the SPI correctly as they must need to use this behaviour: [tom@dhcppc1 quickstart](JBTM-2182) (15:40:29) \$ grep pom transaction-spi ./jca-and-tomcat/pom.xml ./jta-and-hibernate/pom.xml ./jca-and-hibernate/pom.xml It looks like "provided" scope dependencies of constituent modules have not been included in shaded poms. Some of this goes away if we start to recommend users to use the raw modules rather than the uber jars created by the shade plugin as then the poms are self documenting (like in WildFly). At the moment it is effectively assumed that you would know which jars to include if you use the uber jars as the Jars do not document the requirement. The disadvantage of the none-uber jar is only for JTS as you would be required to provide a configuration file to enable JTS, at the moment it is selected by the jar. Unfortunately I can't get a shaded pom to include provided scope at all, I could make it runtime but then it would still be dependent on the Jars being available which for some situations shouldn't be required. Will have a think...	Ask req clarification	https://issues.jboss.org/browse/JBTM-2184
105	JBTM-2183	Have we seen this behaviour elsewhere in the past? I don't recall it. A change in Oracle? Or just lucky for us so far not to have seen it? As to changing the default, we can't do that until the EAP 7 timeframe even if we wanted to. Plus, we'd better be able to show that this has no impact on performance, reliability etc. in environments where the RM doesn't exhibit this behaviour. Otherwise may be best to err on the side of caution and perhaps catch the error in recovery and swap in the alternate behaviour for the next sweep.	Ask RIM??	https://issues.jboss.org/browse/JBTM-2183
106	JBTM-2183	I think its that the EAP test suite now tests this and didn't before. It does sort of make sense that the user to commit has to have permissions on those tables. Maybe we could say the recovery user needs permissions on all the tables? but then we would need some additional config to indicate a recovery user. I will double check with Ondra that its not a regression.	Suggestion	https://issues.jboss.org/browse/JBTM-2183
107	JBTM-2180	As a start on this task I have enabled the Oracle® JRockit Flight Recorder for a subset of the performance tests in the perf CI job (http://172.17.131.2/job/narayana-performance/). Flight data is stored in the maven build directory (target/test-classes/*.jfr) and can be read using the jmc tool available in the standard JDK (edition 7u40 and above). What is the best way to evaluate other options given that there are many alternatives (and trying out each one involves engineering work): * JDK Collections (much of our code hasn't been updated to take advantage of newer releases (Jesper has been looking into this)); * Goldman Sachs Collections; * Google Collections (Guava); * javolution; * Highly Scalable Java; * ... Tests with flight data currently include: Core: com.hp.mwtests.ts.arjuna.performance.Performance1 com.hp.mwtests.ts.arjuna.performance.Performance2 com.hp.mwtests.ts.arjuna.performance.Performance3 com.hp.mwtests.ts.arjuna.performance.Performance4 JTA: com.arjuna.ats.jta.xa.performance.OnePhase2PCPerformanceDefaultUnitTest com.arjuna.ats.jta.xa.performance.OnePhase2PCPerformanceVolatileUnitTest com.arjuna.ats.jta.xa.performance.OnePhasePerformanceDefaultUnitTest com.arjuna.ats.jta.xa.performance.OnePhasePerformanceVolatileUnitTest JTS com.hp.mwtests.ts.jts.local.synchronizations.Performance com.hp.mwtests.ts.jts.orbspecific.local.performance.Performance1 com.hp.mwtests.ts.jts.orbspecific.local.performance.Performance2 com.hp.mwtests.ts.jts.orbspecific.local.performance.Performance3	Compare solutions	https://issues.jboss.org/browse/JBTM-2180
108	JBTM-2163	Hi [~zhfeng], Does your current pull request enable the recovery manager? If not we should make sure you have done that as a top priority before closing this issue. Thanks, Tom	Ask code details	https://issues.jboss.org/browse/JBTM-2163
109	JBTM-2163	Hi [~tomjenkinson], In the narayana-jta-osgi bundle, I just use the configuration from ArjunaJTA/narayana-jta/src/main/resources/jbossts-properties.xml I think it enables the recovery manager and does it need to add the additional entry ? Amos	Ask req clarification	https://issues.jboss.org/browse/JBTM-2163

#	Requirement	Question	Justification	URL
110	JBTM-2159	Huh? Assume I can't view the Bugzilla, which may go part way to explaining my confusion :-). Is this a race condition between recovery and application? If so it should be taken care of already within the transaction code.	Ask req duplication	https://issues.jboss.org/browse/JBTM-2159
111	JBTM-2151	Hi, The QS is merged is this "resolved"? Thanks, Tom	Ask status	https://issues.jboss.org/browse/JBTM-2151
112	JBTM-2151	Should I remove the fix version number as I don't believe this is work that is currently scheduled?	Ask fixed version	https://issues.jboss.org/browse/JBTM-2151
113	JBTM-2119	Hi Lami, Would you be interested in contributing a pull request for this enhancement? If so, please can you sign the CLA over here: https://cla.jboss.org/index.seam (the project is JBoss Transactions) The raise a pull request against either 4.17 or master over here: https://github.com/jbosstm/narayana/pulls Many thanks, Tom	Ask code (PR)	https://issues.jboss.org/browse/JBTM-2119
114	JBTM-2116	Is this something that needs to be meaningful in a distributed environment? What about a heterogeneous environment?	Ask req clarification	https://issues.jboss.org/browse/JBTM-2116
115	JBTM-2116	Hi Jesper, it would be great if you could provide a couple of details on the intended usage of the request, please?	Ask req clarification	https://issues.jboss.org/browse/JBTM-2116
116	JBTM-2116	I added a forum reference for this which I found. I am extrapolating a little, but is the issue related to calling hashCode in JTS afterCompletion in a subordinate node and are there other scenarios you have found it to be volatile?	Ask completeness of req	https://issues.jboss.org/browse/JBTM-2116
117	JBTM-2081	I think that if we are allowing defaults for all values (good idea) then it should be possible for someone to deploy a transactional application without a property file. Maybe give a warning or output in Info?	Suggestion	https://issues.jboss.org/browse/JBTM-2081
118	JBTM-2081	Hi Mark, The issue is that jts is not a strict superset of jta. To be able to provide a single binary containing both JTS and JTA, you therefore need a properties file. Otherwise you have to decide the java hardcoded defaults at compile time, which means using code templates in the jta module compile. We could look at only requiring the properties file once you get to JTA or greater level of the framework? What do you think? Tom	Ask req clarification	https://issues.jboss.org/browse/JBTM-2081
119	JBTM-2081	Or if there's no property file default to JTA (or JTS)?	Ask req clarification	https://issues.jboss.org/browse/JBTM-2081
120	JBTM-2081	Would the intention also be to no longer ship a default jbossts-properties.xml file in the jars, and to no longer be able to change the default name of the config file in the manifest.mf?	Ask intention	https://issues.jboss.org/browse/JBTM-2081
121	JBTM-2081	What values are in the default properties file these days?	Ask function of code(configuration file)	https://issues.jboss.org/browse/JBTM-2081
122	JBTM-2081	I think your idea of requiring it for JTA or JTS is a good start. Beyond that I wonder if we could infer these default values (e.g., JTA implementation) by what's in your class path? Maybe a separate JIRA?	Suggestion & new req	https://issues.jboss.org/browse/JBTM-2081
123	JBTM-2079	I am wondering about whether the reaper could process the XAResources asynchronously but process the afterCompletion from the application thread. It might help the JPA use case?	Ask solution	https://issues.jboss.org/browse/JBTM-2079
124	JBTM-2040	For some reason I don't spot the attachment? It does seem a little extreme but could the level be dropped?	Ask req clarification	https://issues.jboss.org/browse/JBTM-2040
125	JBTM-1972	JIRA is not the appropriate place to ask this question. Please can you raise it via the forums at http://www.jboss.org/ejb3	Ask discuss in external resource	https://issues.jboss.org/browse/JBTM-1972
126	JBTM-1955	We should make sure we are using the PRODUCT variable throughout. We should also add variables for the target app server (IMO). This will allow us to use Narayana and WildFly for the community docs and JBossTS and JBoss AS in the product docs. Does this last point make sense?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1955
127	JBTM-1955	I've read through the XTS docs. I've applied the minor changes (mostly typos) here: https://github.com/jbosstm/documentation/commit/238a06aa9e50ff03144c7f542c08f1b36a9e6b93 . Gytis, Can you apply this to the 4.17 branch please? I also have some ideas about how to layout the docs for the new feature-work done in Narayana 5. I will discuss this with Gytis when we meet to discuss the XTS docs.	Ask fixed version	https://issues.jboss.org/browse/JBTM-1955
128	JBTM-1942	Is this issues resolved now? /cc [~paul.robinson]	Ask status	https://issues.jboss.org/browse/JBTM-1942
129	JBTM-1930	Thanks Ondra, that looks like what I was missing. Mike is there any chance we can back port this into our community docs please?	Ask feasibility of operation	https://issues.jboss.org/browse/JBTM-1930
130	JBTM-1930	The definitive documentation for CLI subsystems are in the XSDs (docs/schema/jboss-as-txn_<version>.xsd). Rather than duplicate this documentation and risk it diverging my preference is to simply point the community docs at this file. Does anyone have any objections with this approach?	Ask review solution	https://issues.jboss.org/browse/JBTM-1930
131	JBTM-1925	[~paul.robinson] I think this is resolved? At least, I managed to debug the XTS issue we were having so I presume it is?	Suggest to change status	https://issues.jboss.org/browse/JBTM-1925
132	JBTM-1925	Gytis, I've had a look at this issue and updated the description with what should now work. Tom and I have tried it out. It's going to be quite time-consuming, and I still have a few M5 issues left to do. Can you take a look at this please. Re-assign it back if you don't think you will have time to do it.	Ask people to take look at req	https://issues.jboss.org/browse/JBTM-1925
133	JBTM-1915	Is this still likely for tomorrow?	Ask priority of req	https://issues.jboss.org/browse/JBTM-1915
134	JBTM-1909	Seems like a few things broke or changed behaviour due to 1875. Can we have a better vetting process in future to either prevent these things creeping in or at least document them such that they don't come as a surprise?	Ask improvement of RIM management	https://issues.jboss.org/browse/JBTM-1909
135	JBTM-1908	I notice qa and rts have printSummary set to true. Maybe do that for everything anyway? Whatever we do, we should be consistent though.	Suggestion solution	https://issues.jboss.org/browse/JBTM-1908
136	JBTM-1908	Did we upgrade the surefire version recently?	Ask project dependency	https://issues.jboss.org/browse/JBTM-1908

#	Requirement	Question	Justification	URL
137	JBTM-1908	I don't especially like the additional text either, but my detective skills have shown this is the reason it does it like this now: http://jira.codehaus.org/browse/SUREFIRE-968 Seems to make sense in the context of what they are trying to achieve, maybe we can raise an enhancement over on their Jira to make this only take affect when parallel is not false?	Suggestion	https://issues.jboss.org/browse/JBTM-1908
138	JBTM-1855	Thanks Mike to take a look. The StompConnect code will be moved to wildfly-blacktie subsystem with JBTM-1568, I just wonder it is any problem to use the bridge or maybe another way to associated JTS transaction ?	Ask feasibility of solution	https://issues.jboss.org/browse/JBTM-1855
139	JBTM-1855	[~gytis], [~zhfeng] mentioned that he had asked you to review: https://github.com/zhfeng/narayana/commit/b51faee71db0977898f471e806b5e88162100945 Have you had chance to do this yet?	Ask contributor's schedul	https://issues.jboss.org/browse/JBTM-1855
140	JBTM-1855	Hi Amos, I think this is resolved? At least, the pull was merged? Tom	Suggestion change status	https://issues.jboss.org/browse/JBTM-1855
141	JBTM-1852	I see. Now I understand the issue but not the context nor the solution, thanks. How does the participant recover and who rolls back the participant if there is no transaction.	Ask req clarification	https://issues.jboss.org/browse/JBTM-1852
142	JBTM-1816	How are you running the test. The following syntax used to work but not anymore: {code} ./build.sh -f ArjunaJTS/pom.xml test -Dtest=com.hp.mwtests.ts.jts.orbspecific.recovery.JavalDLRecoveryUnitTest -Didlj-enabled {code}	Ask code (testing)	https://issues.jboss.org/browse/JBTM-1816
143	JBTM-1816	The reason in the difference in the tests is because if JacORB runs first it has put its RC entry in the object store. You can replicate the issue by using the ObjectStore in the attached jacorb-rc-in-object-store.tgz. If you put it as ArjunaJTS/jts/target/test-classes/ObjectStore then even ". /build.sh -f ArjunaJTS/jts/pom.xml test -Dtest=JavalDLRecoveryUnitTest -Didlj-enabled" would fail. iirc the RC manager uses a known UID hence why running JacORB then JDKORB in the same ObjectStore will be a problem. If you notice, the specific test that fails isn't actually the one that is generated in the log you pasted, so success or fail you get those log messages. i.e. those log messages come from tests where failure is expected. That being said, I don't know if there are any tests where success is expected - we should make sure we have some of those as that log message implies that the JDK orb hasn't been configured correctly. i.e. do we have a test that correctly initializes a JDK ORB RC manager? The only tests I see in JavalDLRecoveryUnitTest are testing failure scenarios.	Ask code (testing)	https://issues.jboss.org/browse/JBTM-1816
144	JBTM-1809	Did this get resolved? I saw some commits on the 4.17 branch	Ask status	https://issues.jboss.org/browse/JBTM-1809
145	JBTM-1808	Tom, I was thinking that this would be good for Amos to investigate. Can you add any ideas you may have to the above text and then assign accordingly?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1808
146	JBTM-1808	Maybe this isn't worth the bother? How about we split the job up into 5 1hr pieces and defer this task until we are in a position where we are using too much CI time?	Suggestion	https://issues.jboss.org/browse/JBTM-1808
147	JBTM-1789	Hi Paul, Can you list the files that are in the uber jars but not in the individual jars? The AS currently just consumes the uber jars so I want to make sure when we move to using the individual jars in the AS we can anticipate any further modifications required. Thanks, Tom	Ask code	https://issues.jboss.org/browse/JBTM-1789
148	JBTM-1788	Hi Tom, [tom@dhcppc1 quickstart]((4.17.4.Final)) \$ git name-rev HEAD HEAD tags/4.17.4.Final [tom@dhcppc1 quickstart]((4.17.4.Final)) \$ find . -name \pom.xml -exec grep -l 4.17.5 {} \; [tom@dhcppc1 quickstart]((4.17.4.Final)) \$ I can't find any reference to 4.17.5 on that tag (SNAPSHOT or otherwise). Can you try to find which file has that reference in? Maybe somehow your repo is corrupted? Tom	Ask code	https://issues.jboss.org/browse/JBTM-1788
149	JBTM-1772	Is this not something that the build can automate?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1772
150	JBTM-1770	Hi Mike, This issue was marked as "Resolved, Done" but there is no "Fix Version" specified. Please can you update the Jira with the version of the software where you fixed this? Also, can you then close the issue (assuming its not an in development version). Thanks, Tom	Ask fixed version	https://issues.jboss.org/browse/JBTM-1770
151	JBTM-1754	Hi Paul, The dir name and artifact id are effectively fixed at the moment. However, as you say we have more flexibility on the "project name" field. Can you update your table with an example of your preferred project name for those modules you have referenced originally, please? Thanks, Tom	Ask code (example)	https://issues.jboss.org/browse/JBTM-1754
152	JBTM-1754	Tom, I have a question: Shall a rename the prefix from "Narayana: " to "JBossTS" for branch 4.17?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1754
153	JBTM-1754	Hi Weinan, I think that would be best, if you didn't mind? Thanks, Tom	Ask confirm	https://issues.jboss.org/browse/JBTM-1754
154	JBTM-1742	In what way did it break the build? I thought I had built everything when I changed this yesterday.	Ask req clarification	https://issues.jboss.org/browse/JBTM-1742
155	JBTM-1721	Is there a pull request for this on WildFly or does it require M5?	Ask code (PR)	https://issues.jboss.org/browse/JBTM-1721
156	JBTM-1707	Hi, Could you please confirm, can we use JBossTS 4.3.GA release for above mentioned problem. Thanks, Sube	Ask useability of specific release version	https://issues.jboss.org/browse/JBTM-1707
157	JBTM-1699	Should we link this to JBTM-1970 as well?	Ask RIM	https://issues.jboss.org/browse/JBTM-1699
158	JBTM-1695	Hi Weinan, We left as much of the old builds directory structure in place as maven could cope with. Please can you clarify if there is an EAP issue with the code in these folders as we have previously decided not to make cosmetic changes to the structure? Tom	Ask req clarification	https://issues.jboss.org/browse/JBTM-1695

#	Requirement	Question	Justification	URL
159	JBTM-1695	Hi Tom, I'm working on this with Paul one by one. After I change the dir structure of one component, I'll rerun mvn install on whole project to make sure I don't break anything. For what I've seen most of these legacy components could be restructured to maven standard one. If there are some more complex situations we could discuss on it then, wdyt?	Ask opinion about solution	https://issues.iboss.org/browse/JBTM-1695
160	JBTM-1694	{quote} we are using a custom fork but only for community testing {quote} That's what our thirdparty-releases repository [1] is for, so ideally we should upload the patched release there. Were the patches also submitted to upstream emma so they can get into a future release? Moving the emma tests to a profile would be good. [1] https://community.iboss.org/wiki/MavenDeployingaThirdpartyRelease	Suggest solution	https://issues.iboss.org/browse/JBTM-1694
161	JBTM-1689	Which profiles do your team regularly use? I guess we should keep just the most commonly used ones?	Ask code	https://issues.iboss.org/browse/JBTM-1689
162	JBTM-1680	Please can you add an estimate for this?	Ask estimated working hour of req	https://issues.iboss.org/browse/JBTM-1680
163	JBTM-1679	[~tomjenkinson], maybe this would be worth pushing to 5.next?	Suggest fixed version	https://issues.iboss.org/browse/JBTM-1679
164	JBTM-1679	How is this different to JBTM-1107? Should we close this as a duplicate?	Ask req duplication	https://issues.iboss.org/browse/JBTM-1679
165	JBTM-1679	I think we should just close it and add it into the details for JBTM-1107 as it seems somewhat intrinsic to the behaviour. If you agree, please can you make the change?	Suggest close req	https://issues.iboss.org/browse/JBTM-1679
166	JBTM-1672	Please can you clarify the reason to disable the maven-surefire-plugin? If it is still needed, please can you raise a separate Jira for this one?	Ask req clarification	https://issues.iboss.org/browse/JBTM-1672
167	JBTM-1672	Also, build-helper is required to create the uber jars, if this is not possible, please can you raise a separate issue and we can keep this one as the "findbugs" issue which I can resolve assuming the tests pass?	Suggest extent to new issue	https://issues.iboss.org/browse/JBTM-1672
168	JBTM-1664	Hi Weinan, I was under the impression that [~pgier] did not want specific EAP profiles creating? I am happy to create a profile that triggers these modules but as I say, I just want to check it is of use first. Thanks, Tom	Ask req clarification	https://issues.iboss.org/browse/JBTM-1664
169	JBTM-1658	Assign this to me, as I think I know what the problem is?	Ask assigne issue to him, so he can contribute	https://issues.iboss.org/browse/JBTM-1658
170	JBTM-1645	How about we do this: Split the PR job into a multi config with the same split as Narayana: MAIN XTS QA_JTA QA_JTS_JACORB QA_JTS_JDKORB These will all be scheduled in a random order, which is annoying if you are willing to merge based on a sub-set of tests. It wastes resources, and causes you to wait longer as the tests you are uninterested in may be scheduled first. The solution is to specify in the description which of the above groups you wish to skip. This is similar to what you mention above, but a bit more course-grained and fits in more neatly with the hudson multi-config approach. You can get the PR description with the following command. There's probably a better approach, but this works: curl -s https://api.github.com/repos/jbosstm/narayana/pulls/279 grep "\"body\"": I don't think it makes sense to allow MAIN to be skipped, and if you want to skip one lot of QA tests, you probably want to skip them all? So how about we just have these two options: skipXTS skipQA We could add more fine-grained options if it proves useful. I don't see the point in specifying who can do this. If a contributor skips some tests that need to pass, we simply tell them to update the description and re-test.	Suggest solution	https://issues.iboss.org/browse/JBTM-1645
171	JBTM-1645	How about if the pull just changes a JacORB qa test? Or just an XTS test, that would be rationale to skip MAIN wouldn't it? Agreed on the white listing, we just need to be vigilant :)	Suggest solution	https://issues.iboss.org/browse/JBTM-1645
172	JBTM-1645	That should be fine, I presume they wouldn't have to "comment_on_pull" as the variable can be checked before so the pull request comments won't include erroneous "Tests passed". Indeed, perhaps they could comment "XTS tests skipped as per request" if we wanted it to?	Ask req clarification	https://issues.iboss.org/browse/JBTM-1645
173	JBTM-1643	Why focus solely on JTS? Why not support something like this for each component?	Ask req clarification	https://issues.iboss.org/browse/JBTM-1643
174	JBTM-1635	Hi Jesper, Do you want me to pretty-print the flag or is the int enough? I don't think you can OR the flags on delist so the int should be enough? Tom	Ask req clarification	https://issues.iboss.org/browse/JBTM-1635
175	JBTM-1598	Hi Weinan, I can do this, but is it not ok for these all to be in the root pom in a dependency management section? I thought that was the preferred route for jboss projects? e.g. grep https://github.com/jbosstm/narayana/blob/master/pom.xml for "common<" Tom	Ask feasibility of solution	https://issues.iboss.org/browse/JBTM-1598
176	JBTM-1598	Hi Tom, because of EAP6 current policy, all the dependency management sections will be removed by VMan automatically (a EAP prod tool), so we have to refer \${project.version} explicitly in the poms. I have done that in EAP6 version, unfortunately I've tried but the patch could not be easily reused on upstream code. I believe we could first have a profile for EAP6, and then just add these version tags for the necessary components, wdyt?	Ask review solution	https://issues.iboss.org/browse/JBTM-1598
177	JBTM-1598	Hi [~weinanli] would you be able to review the pull today do you think? I am hoping to put out a 4.17.4? Thanks, Tom	Ask review code	https://issues.iboss.org/browse/JBTM-1598
178	JBTM-1591	ESB = Enterprise Service Bus?	Ask req clarification	https://issues.iboss.org/browse/JBTM-1591
179	JBTM-1584	Should this be done as part of the BLACKTIE-381 / BLACKTIE-202 package?	Ask RIM	https://issues.iboss.org/browse/JBTM-1584
180	JBTM-1583	I think it would make sense for the implementer of this to also work on BLACKTIE-381?	Suggest RIM	https://issues.iboss.org/browse/JBTM-1583
181	JBTM-1583	How about using rest rather than SOAP?	Suggest solution	https://issues.iboss.org/browse/JBTM-1583

#	Requirement	Question	Justification	URL
182	JBTM-1577	Hi Mike, you means quickstart/blacktie/integration1 ?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1577
183	JBTM-1576	the basic idea I think is to pack server.so(dll) in a .sar and deploy it in JBOSS like a MBean. so we can start/top it with jboss admin console ?	Ask review solution	https://issues.jboss.org/browse/JBTM-1576
184	JBTM-1576	Thanks Tom, 1. we need re-use MDBBlacktieTieService like following: {code} public class CPPXAtmiService extends MDBBlacktieService { native boolean tpervice(byte[] data); public Response tpervice(TPSVCINFO svcinfo) throws ConnectionException, ConfigurationException { /* convert svcinfo to byte[] and pass it to JNI, I'm not sure it is OK */ readEnv(btconfig.xml); // prepare to calculate the path of .so String navelibname = System.mapLibraryName("name of so"); System.load("/path/to/so" + navelibname); tpervice(data); /* but how to get the response from jni ? */ } } {code} 2. the C code to implement tpervice {code} JNIEXPORT jboolean JNICALL Java_CPPXatmiService_tpervice(...) { // we can not use socket call such as tpcall or tpsend // tpreturn is OK ? // and how to return or we don't need to return any response ? } {code} Amos	Ask solution	https://issues.jboss.org/browse/JBTM-1576
185	JBTM-1576	if we have this artifact for the C++ XATMI services, the tpcall in JBoss container don't need to send the message through the HornetQ and just only to load class and call tpervice. This is your means *the benefits of performance (not requiring socket comms between java and C)* ?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1576
186	JBTM-1576	So is this replacing sockets with JNI calls - I would have expected sockets (domain sockets on unix) to be faster than JNI. Have you considered java NIO memory mapped files which will be faster than both JNI and sockets.	Suggest solution	https://issues.jboss.org/browse/JBTM-1576
187	JBTM-1576	Great, though it may make the JNI side more complex? I was thinking we were going to pass a byte array and be able to re-use most of atmibroker-* to parse this. I wonder how much more performant ByteBuffers are? Perhaps Amos as we move forward you can evaluate both and update us via the Jira?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1576
188	JBTM-1576	Please can you add an estimate for this issue?	Ask estimated working hours for req	https://issues.jboss.org/browse/JBTM-1576
189	JBTM-1573	In HybridSessionImpl::send we call apr_socket_timeout_set(stompConnection->socket, 1); and abandon the send if we don't get the message receipt in time. Where in the code are we doing a blocking wait when the transaction timeout is set?	Ask trace to code	https://issues.jboss.org/browse/JBTM-1573
190	JBTM-1573	The attached file shows the problem in action taken from: http://albany/job/blacktie-windows2003/404/ Is it possible that the transaction is expired by the time the timeout for the receive is calculated and so it becomes blocking?? The particular lines of debug of interest are: 2010-08-19 12:56:03,212 [0x00002494] DEBUG (XATMlc :196) - receive session: 3 olen: 02B2EA28 flags: 0 2010-08-19 12:56:03,212 [0x00002494] TRACE (TxLogManagerc :125) - tx_ttl tx=02B832D8 2010-08-19 12:56:03,212 [0x00002494] TRACE (TxLogControl :231) - ENTER ttl=4 ctime=1282218958 now=1282218963 2010-08-19 12:56:03,212 [0x00002494] TRACE (TxLogControl :238) - > ttl=-1 2010-08-19 12:56:03,212 [0x00002494] TRACE (TxLogManagerc :132) - tx->ttl()=0 2010-08-19 12:56:03,212 [0x00002494] TRACE (XATMlc :183) - receive tx_ttl time=0	Ask feasibility of solution	https://issues.jboss.org/browse/JBTM-1573
191	JBTM-1568	Please can you add an estimate for this issue?	Ask estimated working hours for re	https://issues.jboss.org/browse/JBTM-1568
192	JBTM-1568	do you know the jbossas jira for this way to deploy a war as module ? does it need more configuration ? just for merge stempconnect and blacktie-admin-servers ear the estimate is 4h	Ask req clarification	https://issues.jboss.org/browse/JBTM-1568
193	JBTM-1568	I can not access ncl openvpn as my crt has been expired. can you mail me these txt ?	Ask resource	https://issues.jboss.org/browse/JBTM-1568
194	JBTM-1568	Hi Amos, Please find the log from one of the builds attached, it should point out the issue if you scroll to the bottom. Basically blacktie-wildfly is built _after_ a module called "build" that uses it. Maybe move the blacktie module up the list of <module>s or something? Tom	Suggest solution	https://issues.jboss.org/browse/JBTM-1568
195	JBTM-1568	Would it be possible for you to push the fix directly without the need for a PR (since it currently breaks all of our builds). Also why aren't we using the wildfly version that is built during the main narayana build.	Ask feasibility/possibility of solution	https://issues.jboss.org/browse/JBTM-1568
196	JBTM-1565	Hi Crispin, Please can you: 1. Sign our CLA: https://cla.jboss.org/index.seam (choose JBoss Transactions rather than BlackTie) 2. Raise a pull request for your change? This will then run CI over the proposed enhancement and double check it for us I will then merge the PR after it passes :) Thanks, Tom	Ask code (PR)	https://issues.jboss.org/browse/JBTM-1565
197	JBTM-1558	Can you provide more detail to this task, since I thought we already had an SPI?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1558
198	JBTM-1558	There are a few SPI related issues at the moment. Maybe cross-link them?	Ask RIM	https://issues.jboss.org/browse/JBTM-1558
199	JBTM-1556	What do you mean by a remote thread? We should be using interposition in all distributed cases, which means the Synchronisation will be registered with a subordinate coordinator.	Ask req clarification	https://issues.jboss.org/browse/JBTM-1556
200	JBTM-1556	When each coordinator handles the timeout, does it also cancel its subordinates (as if the subordinate is one of the resources)? Or is it more like the coordinator times out the subordinate (as a resource) which means blocking until the subordinate has completed the subordinate level (reaper) timing out of its (subordinate) resources? In other words, could the subordinate coordinator Synchronization.afterCompletion() be called from a transport thread ever or only from a reaper thread?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1556
201	JBTM-1526	Please can you itemize the dependencies that you are worried about - thanks	Ask dependencies	https://issues.jboss.org/browse/JBTM-1526

#	Requirement	Question	Justification	URL
202	JBTM-1526	Hi Weinan, Please can you provide an update on which dependencies you are worried about? We might be cutting a 4.17.4 soon and would like to hear from you, thanks! Tom	Ask dependencies	https://issues.jboss.org/browse/JBTM-1526
203	JBTM-1526	Hi Weinan, No worries, hope you had good PTO. My question is more which of the dependencies that those modules you listed have are you concerned about (i.e. I guess there are some none-JBoss stuff? Thanks! Tom	Ask dependencies	https://issues.jboss.org/browse/JBTM-1526
204	JBTM-1526	Hi Weinan, Please can you raise a Jira for the second issue (orb.idl) as we don't see that on our machines, it might be environmental so let us know the OS etc and assign it [~mmusgrov] Please can you raise a pull request for the first one? Thanks! Tom	Ask code (PR) & seperate req	https://issues.jboss.org/browse/JBTM-1526
205	JBTM-1503	Just wondered what had necessitated this? Not that it is a problem but we should record in the Jira why this decision was made.	Ask reason/use case of req	https://issues.jboss.org/browse/JBTM-1503
206	JBTM-1498	Paul, is this now out of date? If not the description should be something like "add a TXFramework version of all relevant quickstarts"	Ask status	https://issues.jboss.org/browse/JBTM-1498
207	JBTM-1490	Elevated to critical, please do downgrade (optional?) if not required.	Suggestion	https://issues.jboss.org/browse/JBTM-1490
208	JBTM-1485	Does this change represent a possible backward incompatibility?	Ask change impact	https://issues.jboss.org/browse/JBTM-1485
209	JBTM-1479	It appears that the commit for JBICA-1011 made it into IronJacamar 1.0.16.Final. Is the quickstart working now? Where do I find it?	Ask code	https://issues.jboss.org/browse/JBTM-1479
210	JBTM-1479	I think that might be enough information to log a bug report against IronJacamar? Can you link this bug as dependent on the one you raise? Also can you provide the link to the stackoverflow article just in case. Thanks, Tom	Ask code	https://issues.jboss.org/browse/JBTM-1479
211	JBTM-1463	you are right in your concern and that is what I am trying to address but ineloquently I was thinking more about changes to narayana that need a change to jbosstm somehow being able to test both pulls at once. Maybe we could somehow put convention when you raise a pull request for jbosstm that you can put a comment to say something like: "requires: https://github.com/jbosstm/jboss-as/pull/217" We can then somehow use that info in the narayana.sh to pull in that branch rather than 5_BRANCH, what do you think? I know you did some work with using rest to talk to git hub to comment on pulls, could it possibly download the comment for the pull in order to grep it for that info?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1463
212	JBTM-1463	Please can you add an estimate for this?	Ask estimated working hours of req	https://issues.jboss.org/browse/JBTM-1463
213	JBTM-1454	Why? Also, will you be testing HTTP versus IOP initially to get a base?	Ask reason/use case of req	https://issues.jboss.org/browse/JBTM-1454
214	JBTM-1424	Please can you add a component for this?	Ask component of req	https://issues.jboss.org/browse/JBTM-1424
215	JBTM-1424	Shouldn't this also be fixed for 4.17?	Ask fixed version	https://issues.jboss.org/browse/JBTM-1424
216	JBTM-1423	Please can you add a component for this?	Ask component of req	https://issues.jboss.org/browse/JBTM-1423
217	JBTM-1407	It should be trivial to do that *today* for JDK orb with profiles as you can activate a profile by the presence of a file, e.g. JAVA_HOME/lib/orb.idl. However, JacORB would be in the classpath unless you specifically disabled it with a different profile so its presence would be guaranteed. Now, in the future when AS7-5582 is implemented, the "JDK" orb will be a maven dependency (same as JacORB), i.e. it will be available by default in the classpath. Lets say we made the change today as a convenience until AS7-5582 is available, the question is do we really want to have unpredictable results in the build. Lets say the user attempts this new default, detection based build but does still expect both a JDK orb and JacORB version but we as can't locate orb.idl (say they are on pi) and just build the JacORB version - what would you expect then, log a warning?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1407
218	JBTM-1398	Should org.jboss.as.weld.deployment.WeldAttachments be lifted into its own subsystem that doesn't depend on EJB3?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1398
219	JBTM-1398	I think this is resolved for now. Do you agree Tom?	Suggest to change status of req	https://issues.jboss.org/browse/JBTM-1398
220	JBTM-1382	Hi Mike, THANKS for raising this, I keep forgetting! Can we put the output in testoutput/ that way we can have a git ignore on just testoutput rather than testoutput TEST*.txt. Tom	Ask req clarification	https://issues.jboss.org/browse/JBTM-1382
221	JBTM-1382	The JIRA asks for output to be sent to qa/testoutput (not qa). For a qa run against both orbs if you look in the testoutput directory you will find two zips (testoutput-jacorb.zip and testoutput-idlj.zip) which contain the testoutput directory and the TEST*.txt files corresponding to each run. To see it working look at the latest narayana-dualstack job (I did notice that btny-pulls-narayana CI job does not archive testoutput so I have updated that job config). I think I may have misunderstood your requirement. Did you mean move all the TEST*.txt files into testoutput but leave the zip files in the qa directory.	Ask req clarification	https://issues.jboss.org/browse/JBTM-1382
222	JBTM-1369	Can we have an update on this?	Ask update of req	https://issues.jboss.org/browse/JBTM-1369
223	JBTM-1369	Can you paste in the commands you used to get those numbers, please?	Ask code	https://issues.jboss.org/browse/JBTM-1369
224	JBTM-1369	What version of JacORB was used please?	Ask running environment	https://issues.jboss.org/browse/JBTM-1369
225	JBTM-1369	What hardware architecture is your CI cluster?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1369
226	JBTM-1366	Just wondering if this should be in the Arquillian Transaction Extension Jira instance rather than here?	Ask project of req	https://issues.jboss.org/browse/JBTM-1366
227	JBTM-1359	Why has this been linked to the JDBC object store work? I'd like to understand the relative priority of this to other tasks.	Ask req clarification	https://issues.jboss.org/browse/JBTM-1359

#	Requirement	Question	Justification	URL
228	JBTM-1359	Thanks, but it still doesn't seem that there is clarity on what exactly is required. What is the problem that we're trying to fix? Is it only for JTA? Until we have answers to these and other questions there's no point in implementing anything.	Ask req clarification	https://issues.iboss.org/browse/JBTM-1359
229	JBTM-1357	It's worth pointing out that the script is inherited from AS7's build.sh, do you think you could try to raise this "upstream" so to speak?	Ask req clarification	https://issues.iboss.org/browse/JBTM-1357
230	JBTM-1340	I guess I need to look into this a bit more, but I think that should be a one-off cost? Not that we shouldn't optimize it as L 1. there could be use cases to deploy JBTM just to run a single transaction then to passivate somehow 2. we may want to support reloading config. 3. we may need to optimize boot time costs But to confirm the priority, I think this should be a one-off cost per boot of the TM? Also, please can you confirm the tool you are using to generate the profile (command line parameters too if possible) and maybe your test case if possible.	Ask req clarification	https://issues.iboss.org/browse/JBTM-1340
231	JBTM-1340	Hi Mark, Please can you upload the test you were using and the name of the tool that outputted the profiling information? We can look at tuning TxStats, I am thinking a listener pattern might help? Before that though, the test and profiler settings would be useful. Thanks, Tom	Ask code (testing)	https://issues.iboss.org/browse/JBTM-1340
232	JBTM-1339	Isn't this out of date now?	Ask status	https://issues.iboss.org/browse/JBTM-1339
233	JBTM-1339	OK, thanks. Does the title of the issue need changing now though? I'm assuming we're not back porting this to AS7 and only focussing on WildFly.	Ask req clarification	https://issues.iboss.org/browse/JBTM-1339
234	JBTM-1338	If there is commitAsync, there should also be rollbackAsync?	Ask req clarification	https://issues.iboss.org/browse/JBTM-1338
235	JBTM-1338	How do I get access to BaseTransaction in AS7? As JNDI lookup (java:jboss/UserTransaction) returns ServerVMClientUserTransaction.	Ask req clarification	https://issues.iboss.org/browse/JBTM-1338
236	JBTM-1328	Was this a one-off?	Ask req clarification	https://issues.iboss.org/browse/JBTM-1328
237	JBTM-1328	Hi Mike, This issue was marked as "Resolved, Done" but there is no "Fix Version" specified. Please can you update the Jira with the version of the software where you fixed this? Also, can you then close the issue (assuming its not an in development version). Thanks, Tom	Ask fixed version	https://issues.iboss.org/browse/JBTM-1328
238	JBTM-1313	Actually, maybe we need that job? it might be worth renaming it as it looks like an old "unit test" repeat job named as it is?	Ask req clarification	https://issues.iboss.org/browse/JBTM-1313
239	JBTM-1310	Does this affect all versions or 4.17.1+ in particular?	Ask affected veriosns	https://issues.iboss.org/browse/JBTM-1310
240	JBTM-1277	Hi Mike, In the (ideal) future where everything is in Maven we will need to support -DskipTests in qa/ anyway and this is just one step towards that. It is integrated into trunk now, we can always back it out if it becomes an issue? Maybe the priority can be like this: User sets: -DskipTests - skip all tests User sets: QA_TESTS=1 ./scripts/hudson/narayana.sh -DskipTests - override skipTests for the QA suite? I am trying to make it appear more maven like, it wasn't clear to me that skipTests would not skip all the tests sorry :(I am also thinking what is the most intuitive behaviour for the build. Tom	Ask req clarification	https://issues.iboss.org/browse/JBTM-1277
241	JBTM-1271	Hi Paul, this is a duplicate of JBTM-1072 (assigned to you) I think?	Suggestion req duplication	https://issues.iboss.org/browse/JBTM-1271
242	JBTM-1266	[~tomjenkinson] Not sure if this is relevant anymore. Let me know if you want to me to lookup these "updates" that I spoke of.	Ask needs of change status	https://issues.iboss.org/browse/JBTM-1266
243	JBTM-1254	Mike, Does that mean that you agree with me that we should resolve this as "won't fix"?	Suggestion fixed version	https://issues.iboss.org/browse/JBTM-1254
244	JBTM-1251	Thanks for raising this Paul. I would go with jbossts-EAP61-branch417 for the existing jobs. We can never tell what the next EAP is, perhaps Narayana will fold into EAP8 or EAP6.x so I am not sure of the naming for the future job though. Maybe narayana-500 is OK (i.e. don't pin it to an EAP in terms of name)?	Ask req clarification	https://issues.iboss.org/browse/JBTM-1251
245	JBTM-1251	You don't need to branch the quickstart, narayana, or docs branches to create these jobs. Currently 5.x and 4.17 share a common branch (master) so both sets of jobs build against master. The place in which they differ is that 4.17 builds against 4_BRANCH of our jboss-as fork and 5.x against 5_BRANCH. They also differ in what is tested as TXF is only tested in 5.x. I would suggest we create the jobs now, but the problem is that we are going to double the load on Hudson as any change is going to trigger both the EAP61 and the narayana-500 version of the job. The reason why I'm keen to get some narayana 5.x CI is that TXF is currently not getting any CI. Maybe I could just create some TXF jobs in the mean time until this is resolved?	Ask review of solution	https://issues.iboss.org/browse/JBTM-1251
246	JBTM-1247	What is the use case for this?	Ask reason/use case of req	https://issues.iboss.org/browse/JBTM-1247
247	JBTM-1227	Did this make it into PropertiesFactory?	Ask req clarification	https://issues.iboss.org/browse/JBTM-1227
248	JBTM-1217	Seems to be related to REM3-158. Ivo, I'm on PTO until August 28th. Can you try rebuilding AS7 with XNIO 3.0.6 and then re-testing on the IPv6 node?	Suggestion solution	https://issues.iboss.org/browse/JBTM-1217
249	JBTM-1212	Amos, Has this been done for the docs yet?	Ask status	https://issues.iboss.org/browse/JBTM-1212
250	JBTM-1212	Can we mark this as resolved now?	Suggestion change status	https://issues.iboss.org/browse/JBTM-1212
251	JBTM-1208	Paul, I wonder that the contents of recovery records have only the text or the binary characters ?	Ask req clarification	https://issues.iboss.org/browse/JBTM-1208
252	JBTM-1199	http://172.17.131.2/view/Narayana+BlackTie/job/narayana-java6/1237/ Seemed to start happening for no obvious reason - I guess it is an underlying change in AS/arq?	Ask status	https://issues.iboss.org/browse/JBTM-1199

#	Requirement	Question	Justification	URL
253	JBTM-1199	Amos, please can you investigate why https://github.com/jbossas/jboss-as/commit/435c862b9514bfcf0a036bb2b916ac3dc63b4f4a would stop the XTS crash recovery tests from working? I noticed somethings to do with "Reload" perhaps AS7 has some behaviour that reloads the server automatically now?	Ask some one investigate sth	https://issues.jboss.org/browse/JBTM-1199
254	JBTM-1188	Where did the requirement for a per-tx override come from? Is the existing global config option not sufficient?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1188
255	JBTM-1187	Amos, What made you think "the client handler chain can be added by txframework automatically"? This is a feature we have planned but we don't have it yet. Did you read it in some docs somewhere?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1187
256	JBTM-1187	Paul, I just wonder why we need ClientStub in the tests of wsat-simple and wsba-participant-completion-simple currently. I know this feature which have not in txframework yet. But once we have it, the client stub codes will not be needed ?	Ask reason of req	https://issues.jboss.org/browse/JBTM-1187
257	JBTM-1184	Status update please?	Ask status	https://issues.jboss.org/browse/JBTM-1184
258	JBTM-1182	With recent changes to transactional fileio repo, is this still relevant? Also what about the AS7 RPMs that have already been produced for Fedora: don't they duplicate this?	Ask status & duplication	https://issues.jboss.org/browse/JBTM-1182
259	JBTM-1173	How to run the tests: # ipv4: {noformat} mvn clean test {noformat} # ipv6: {noformat} mvn clean test -DipV6 -Parq-jbossas-managed-7 {noformat} Note that for IPv6 the XTS demo app needs to be compiled with HOSTNAME set to [::1] in jboss.properties file.	Ask code (testing)	https://issues.jboss.org/browse/JBTM-1173
260	JBTM-1173	Please can you update the status of this Jira to indicate when it will be ready as we intend to release M2 on Monday. Thanks, Tom	Ask status	https://issues.jboss.org/browse/JBTM-1173
261	JBTM-1173	I can do this on Sunday (hopefully). The diff looks quite simple. The problem will be testing it. I'm not 100% sure i have IP6 only-mode working properly on my machines and I don't know where Mike is with getting IP6 working on Hudson. Mike: will I be able to test this on Hudson before the Monday release?	Ask task status	https://issues.jboss.org/browse/JBTM-1173
262	JBTM-1173	I won't be able to test this before Weds. If it's urgent, could someone else test the pull? It sounds like we have many other IPv6 (JBTM-1156) issues that will be pushed to 5.final, so we could push this too and I'll look at it and JBTM-1156 on Weds when I hope to be back to work.	Ask some on help to test PR	https://issues.jboss.org/browse/JBTM-1173
263	JBTM-1166	Hi Mike, This issue was marked as "Resolved, Done" but there is no "Fix Version" specified. Please can you update the Jira with the version of the software where you fixed this? Also, can you then close the issue (assuming its not an in development version). Thanks, Tom	Ask fixed version	https://issues.jboss.org/browse/JBTM-1166
264	JBTM-1156	Paul, do you have any experience with setting up jobs on the QE cluster?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1156
265	JBTM-1156	Please can you update the status of this Jira to indicate when it will be ready as we intend to release M2 on Monday. Thanks, Tom	Suggestion	https://issues.jboss.org/browse/JBTM-1156
266	JBTM-1156	Great, are you working with Paul on resolving the XTS issues? +1 on the JDBC tests	Ask some one's current task on this req	https://issues.jboss.org/browse/JBTM-1156
267	JBTM-1147	Please can you add an estimate for this issue?	Ask estimated working hours of req	https://issues.jboss.org/browse/JBTM-1147
268	JBTM-1082	As workaround is possible to define java opt argument -XX:-UseSplitVerifier Put it to file sharedbuild.xml to the tag <junit> like: <jvmarg value="-XX:-UseSplitVerifier" /> I suppose that this problem is caused by the fact that the emma library was compiled with older version of compiler that we are trying to use for compiling TS sources. Some links e.g. here http://stackoverflow.com/questions/7970622/java-7-jvm-verifyerror-in-eclipse Do you think that the workaround is enough for solving this jira or should be tried to recompile emma or dive to some other investigation?	Ask completeness of solution	https://issues.jboss.org/browse/JBTM-1082
269	JBTM-1073	Worth a link to where this was done?	Ask code	https://issues.jboss.org/browse/JBTM-1073
270	JBTM-1072	I commented on the pull re the password in narayana.sh, not ideal, can it be externalized into the job config? Also, I commented to say what happens if you are running this locally on a checked out pull, can we make the default that the commenting is disabled and job config enables it?	Ask feasibility of solution	https://issues.jboss.org/browse/JBTM-1072
271	JBTM-1072	Tom, btny-pulls-blacktie and btny-pulls-quickstarts are tied to Windows2008. btny-pulls-documentation will run on any machine (including windows). It's not going to be trivial to get this working on windows: {code}function comment_on_pull { if ["\$COMMENT_ON_PULL" = ""]; then return; fi PULL_NUMBER=\$(echo \$GIT_BRANCH awk -F 'pull' '{ print \$2 }' awk -F '/' '{ print \$2 }') if ["\$PULL_NUMBER" != ""] then JSON="{ \"body\": \"'\$1'\" }" curl -d "\$JSON" -ujbosstm-bot:\$BOT_PASSWORD https://api.github.com/repos/\$GIT_ACCOUNT/\$GIT_REPO/issues/\$PULL_NUMBER/comments else echo "Not a pull request, so not commenting" fi } {code} I guess the only way to do this is by having the job run cygwin with a bash script that does the above. I don't fancy finding Windows commands for all of the above! WDYT?	Ask review solution	https://issues.jboss.org/browse/JBTM-1072
272	JBTM-1072	in terms of the quickstarts one, I wonder how complex that would be in bat files? It is mainly curl (which I believe can be installed as an exe). We really just need to translate: PULL_NUMBER=\$(echo \$GIT_BRANCH awk -F 'pull' '{ print \$2 }' awk -F '/' '{ print \$2 }') Again, that looks like it is mainly awk. I bet that can be installed too.	Ask complexity of solution	https://issues.jboss.org/browse/JBTM-1072
273	JBTM-1063	Hi Amos, Please can you backport this to the branch? Thanks! Tom	Ask code	https://issues.jboss.org/browse/JBTM-1063

#	Requirement	Question	Justification	URL
274	JBTM-1063	Amos, {quote} and now jboss-as version is 7.1.2.SNAPSHOT and jbossts version is 4.16.2.Final {quote} What do you mean by this?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1063
275	JBTM-1063	Hi, I have updated the config on the jobs to use this new script thanks Amos! Paul, please can you give Amos a pointer on the xts commands? Tom	Ask some one provide help	https://issues.jboss.org/browse/JBTM-1063
276	JBTM-1063	Paul, I checkout 5_BRANCH on jboss-as and build fail with following: [ERROR] Failed to execute goal on project jboss-as-transactions: Could not resolve dependencies for project org.jboss.as:jboss-as-transactions:jar:7.1.1.Final: The following artifacts could not be resolved: org.jboss.jbossts:jbossts:jar:5.0.0.M2-SNAPSHOT, org.jboss.jbossts:jbossts-integration:jar:5.0.0.M2-SNAPSHOT: Failure to find org.jboss.jbossts:jbossts:jar:5.0.0.M2-SNAPSHOT in http://repository.jboss.org/nexus/content/groups/public/ was cached in the local repository, resolution will not be reattempted until the update interval of jboss-public-repository-group has elapsed or updates are forced It looks like we should change groupId from "org.jboss.jbossts" to "org.jboss.narayana" can you confirm this change? Amos	Ask confirmation of change	https://issues.jboss.org/browse/JBTM-1063
277	JBTM-1063	Amos, Can you give us an update on when this will be ported to the 4.16 branch?	Ask time of fix issue	https://issues.jboss.org/browse/JBTM-1063
278	JBTM-1063	Paul, I ran these test commands on my machine and it failed with <testcase classname="com.arjuna.qa.junit.LocalJUnitRunner" name="testCallServlet" time="0.01"> <error message="org/apache/commons/logging/LogFactory" type="java.lang.NoClassDefFoundError">java.lang.NoClassDefFoundError: org/apache/commons/logging/LogFactory at org.apache.commons.httpclient.HttpClient.<clinit>(Unknown Source) at com.arjuna.qa.junit.HttpUtils.accessURL(HttpUtils.java:97) at com.arjuna.qa.junit.LocalJUnitRunner.testCallServlet(LocalJUnitRunner.java:40) Caused by: java.lang.ClassNotFoundException: org.apache.commons.logging.LogFactory at java.net.URLClassLoader\$1.run(URLClassLoader.java:202) at java.security.AccessController.doPrivileged(Native Method) at java.net.URLClassLoader.findClass(URLClassLoader.java:190) at java.lang.ClassLoader.loadClass(ClassLoader.java:306) at sun.misc.Launcher\$AppClassLoader.loadClass(Launcher.java:301) at java.lang.ClassLoader.loadClass(ClassLoader.java:247) </error> </testcase> did I miss some jars ?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1063
279	JBTM-1063	I think that was fixed by https://issues.jboss.org/browse/JBTM-1086 Have you done an 'svn up' on the 4.16 branch?	Suggestion solution	https://issues.jboss.org/browse/JBTM-1063
280	JBTM-1063	I think that the branch version needs to hook in the txbridge tests?	Suggestion	https://issues.jboss.org/browse/JBTM-1063
281	JBTM-1063	how about commands on the txbridge tests?	Suggestion	https://issues.jboss.org/browse/JBTM-1063
282	JBTM-1061	I just put this on HORNETQ-274: Hi I just trawled the the Jira and forum postings I ls'd my .m2/repository/org/hornetq find . -name \pom.xml in the AS7 version of hornetq but can't find the dependency to import for this. Do you know that this work was done? If so, do you happen to know what I should replace: <dependency> <groupId>org.hornetq</groupId> <artifactId>hornetq-core</artifactId> </dependency> with? I also can't find it by grepping the internet, e.g. http://mvnrepository.com/artifact/org.hornetq	Ask status	https://issues.jboss.org/browse/JBTM-1061
283	JBTM-1057	1. (Slightly unrelated) can we move the XTS demo to quickstarts/? 2. (on topic) Is the idea that the user builds upon the XTS demo or is the restaurant service used for hack/simplicity? If the idea is that we build upon the XTS demo then I think that the dependency (assuming XTS demo is moved to quickstarts) is OK, if the restaurant service is however used out of convenience and the XTS demo is not referenced in any way in associated meta-data (README.md etc) then I think it is a good idea to duplicate the code in the TXBridge demo and break the dependency.	Ask req clarification	https://issues.jboss.org/browse/JBTM-1057
284	JBTM-1052	Is this still required? We could do with a better description to explain the actual motivation or leave this closed as "incomplete description".	Ask status	https://issues.jboss.org/browse/JBTM-1052
285	JBTM-1051	A better solution would be to have this problem detected in the AS7 testsuite. This would prevent the error from even being accepted into the AS7 codebase. Ivo, can this type of check be made in the allTests junit suite?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1051
286	JBTM-1051	PS, shouldn't this issue be moved to an AS7 Jira?	Ask project of req	https://issues.jboss.org/browse/JBTM-1051
287	JBTM-1051	I don't even know where standalone-jts.xml is referenced from :) In my quickstarts I show people how to change the config manually. I wonder who added the jts one? Maybe that is an option for XTS? Explain the manual configuration?	Ask req clarification	https://issues.jboss.org/browse/JBTM-1051
288	JBTM-1051	Ivo, Of course, I forgot these tests would do this also. Do your tests fail if the server boots with errors, even though XTS is still functioning correctly? We still have the problem where a minor change is forgotten; one which doesn't break XTS and doesn't fail the server boot. However, these are relatively minor and could probably be checked manually every now and then.	Ask req clarification	https://issues.jboss.org/browse/JBTM-1051
289	JBTM-1051	Can this be moved to an AS7 Jira?	Ask project of req	https://issues.jboss.org/browse/JBTM-1051
290	JBTM-1051	Sure, but can we re-open issues like this in AS7 project do you think?	Ask project of req	https://issues.jboss.org/browse/JBTM-1051
291	JBTM-1031	I've just added 4.17.4 to the fix version. Can you cherry-pick your commit from master and raise a pull request on 4.17?	Ask some one to do sth	https://issues.jboss.org/browse/JBTM-1031
292	JBTM-1016	I suggest we add a dependency on BYTEMAN-190 for this task. Amos, if you agree, can you add the dependency please? More Actions -> Link -> Depends on	Suggestion	https://issues.jboss.org/browse/JBTM-1016
293	JBTM-1016	Please can you add an estimate for this issue?	Ask estimated working hours of req	https://issues.jboss.org/browse/JBTM-1016

#	Requirement	Question	Justification	URL
294	JBTM-1016	Paul, can I delete XTS/localjunit/crash-recovery-tests/check.sh as now we can use maven plugin to check rule ?	Ask feasibility of solution/ change impact analysis	https://issues.iboss.org/browse/JBTM-1016
295	JBTM-989	Playing Devil's Advocate: and what happens to certain modules that are only every touched by one person who may not like the specific code style?	Ask req clarification	https://issues.iboss.org/browse/JBTM-989
296	JBTM-989	I think this should apply to all modules hanging off the root narayana pom. The following statements should be taken assuming we do decide to reformat: I don't think we have any modules that fulfil the property of being isolated to the single developer degree, some could be at the moment but in the long term are very unlikely to be. Incubator style modules - e.g. STM - where there is a single developer working on it at the moment, probably don't need the formatting *right now*, but before they become .Final and as such become community supported I believe they would need to follow the consistent style as the developer has contributed that work to the public domain and adhering to a formatting will make it easier for other developers to then submit patches/pull reqs. Also, a note to the above statement, if we don't like the AS7 style potentially we could develop our own (negating one positive which would be consistency with AS7). I am of the opinion of where I don't really mind what the code looks like (within reason) so long as it all looks the same?	Ask req clarification	https://issues.iboss.org/browse/JBTM-989
297	JBTM-989	Also, I think it's worth stressing that the reason why this is being suggested is not because I/we don't like the current code style. It's because using a common code-style brings efficiencies. Hence the reason why I'm willing to adopt a style that I don't like.	Ask req clarification	https://issues.iboss.org/browse/JBTM-989
298	JBTM-968	Quick notes after Andrew explained the XTS structure to me: {quote} XTS project layout WS-AS, Activity service (needed? - yes but need to know why) WS-CF (general purpose context management), useful WS-C Client and coordinator side impl of the oasis standard (current) Activation and registration service WS-T oasis TX protocol specific services for client, coord, ws negotiation (WS-AT and WS-BA) WS-TX JBoss specific APIs which drive WS-C and WS-T services {quote}	Ask reason of req	https://issues.iboss.org/browse/JBTM-968
299	JBTM-956	Hi Amos, is this closed?	Ask status	https://issues.iboss.org/browse/JBTM-956
300	JBTM-948	Ah, ok. So what you are saying is that this will work on Hudson as you have fixed this problem before for other jobs?	Ask req clarification	https://issues.iboss.org/browse/JBTM-948
301	JBTM-948	Paul, what CI environment did you mean in your comment [1]? Yes, I was talking about Hudson as it is the primary environment for QE testing. If I run the tests locally on my laptop I usually don't need headless mode as I want to see how the test is going and thus see its actions performed in Firefox in front of my eyes. This is the reason why the tests look so at the moment. But I can change them to better fit to your requirements. On Hudson if you want to run a selenium test you need to enable 'Run Xvnc during build' and select some of selenium slaves at 'Label Expression'. Then the selenium test should work. [1] https://issues.iboss.org/browse/JBTM-948?focusedCommentId=12661667&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-12661667	Ask req clarification	https://issues.iboss.org/browse/JBTM-948
302	JBTM-940	Please can you update the status of this Jira to indicate when it will be ready as we intend to release M2 on Monday. Thanks, Tom	Ask status	https://issues.iboss.org/browse/JBTM-940
303	JBTM-934	Do we need pseudo support for idlj? [tom@tomsfc15 JBOSSSTS_4_15_0_Final]\$ find . -name *idl -exec grep -l pseudo {} \; ./install/jacorb/idl/omg/Corba.idl ./install/jacorb/idl/omg/orb.idl ./install/jacorb/idl/omg/CosTSPortability.idl ./install/jacorb/idl/omg/CosNaming.old.idl ./install/jacorb/idl/omg/PortableServer.idl ./ArjunaJTS/build/install/jacorb/idl/omg/Corba.idl ./ArjunaJTS/build/install/jacorb/idl/omg/orb.idl ./ArjunaJTS/build/install/jacorb/idl/omg/CosTSPortability.idl ./ArjunaJTS/build/install/jacorb/idl/omg/CosNaming.old.idl ./ArjunaJTS/build/install/jacorb/idl/omg/PortableServer.idl ./ArjunaJTS/jacorb/idl/omg/Corba.idl ./ArjunaJTS/jacorb/idl/omg/orb.idl ./ArjunaJTS/jacorb/idl/omg/CosTSPortability.idl ./ArjunaJTS/jacorb/idl/omg/CosNaming.old.idl ./ArjunaJTS/jacorb/idl/omg/PortableServer.idl	Ask req clarification	https://issues.iboss.org/browse/JBTM-934
304	JBTM-934	Please can you update the status of this Jira to indicate when it will be ready as we intend to release M2 on Monday. Thanks, Tom	Suggestin	https://issues.iboss.org/browse/JBTM-934
305	JBTM-934	Will it work on OpenJDK?	Ask req clarification	https://issues.iboss.org/browse/JBTM-934
306	JBTM-933	The main questions I have are: 1. where would this file be written to, is it ok to write to current working directory? 2. the XID would have this integer in its bqual, therefore, are there possibly none-TS tools that would want to read the bqual and need to read meaningful information regarding the EIS without the benefit of our utility classes?	Ask req clarification	https://issues.iboss.org/browse/JBTM-933
307	JBTM-932	Yeah, I considered the tx reaper, but what about if the timeout is 0?	Ask req clarification	https://issues.iboss.org/browse/JBTM-932
308	JBTM-928	[~tomjenkinson], these test still don't work. Shall I investigate further or keep them ignored?	Ask solution	https://issues.iboss.org/browse/JBTM-928
309	JBTM-914	Hi Paul, I don't think EAP5 support is something we should drive any longer, if you are happy this is resolved on EAP6 can we update the fixversions, subject etc and resolve and close this one? Thanks, Tom	Suggest update status, fixed version etc.	https://issues.iboss.org/browse/JBTM-914
310	JBTM-909	Yeah I think we can either defer this issue or simply close it (and re-open later if we ever want to revisit the issue). Agreed?	Ask req clarification	https://issues.iboss.org/browse/JBTM-909
311	JBTM-883	Does this work? Are you ready to create a pull request?	Ask status	https://issues.iboss.org/browse/JBTM-883

#	Requirement	Question	Justification	URL
312	JBTM-877	Thanks Andrew, I think that we need to make this clearer to the user somehow. Maybe we can update the readme.txt to clarify where the emma.jar in ext is coming from?	Suggestion solution	https://issues.jboss.org/browse/JBTM-877
313	JBTM-852	Firstly some background. It is possible to configure different stores for different purposes. Each store is a singleton so each store type is shared. We currently use at least four stores (XTS uses some others): - an action store for transaction and participant logs; - a recovery store (which defaults to the action store); - a state store for use with TxOJ and is not used with JTA (or JTS); - a communications store for storing connectivity information such as ORB endpoints or transaction manager status items. Calling <code>arjPropertyManager.getObjectStoreEnvironmentBean()</code> looks up the default store and is equivalent to: <code>BeanPopulator.getNamedInstance(ObjectStoreEnvironmentBean.class, "default")</code> and is used to configure the action store. The other stores can be configured by looking up a specific named <code>ObjectStoreEnvironmentBean</code>. Use "default" or "stateStore" to configure the action and state stores respectively. Use "communicationStore" to configure the communications store. So the actual bug/feature is that looking up the default store environment bean via the properties manager (<code>arjPropertyManager.getObjectStoreEnvironmentBean()</code>) should return a bean that will configure all stores. However this could be risky (for example it could interfere with XTS behaviour) so my preference would be to force the user to be specific about which stores he wants an environment bean for by making the <code>getObjectStoreEnvironmentBean</code> call take the store name as a parameter. But that leaves the difficult issue of how to specify which store to use in the properties file <entry key="ObjectStoreEnvironmentBean.objectStoreType">XYZStore</entry> since this entry is only applied to the default store. We could simply document the behaviour in <code>jboss-props.xml</code> and tell users to explicitly set the communications store programmatically if required. So the communications store type would have to be set as follows: <code>BeanPopulator.getNamedInstance(ObjectStoreEnvironmentBean.class, "communicationStore").setObjectStoreType(com.arjuna.ats.internal.arjuna.objectstore.VolatileStore.class.getName());</code> Is this latter approach acceptable to all concerned?	Ask completeness of solution	https://issues.jboss.org/browse/JBTM-852
314	JBTM-852	Thanks Mike, I presume you have confirmed nothing is created on disk in that scenario? Perhaps your resolution could be to update the documentation in the ArjunaCore programmers guide (or wherever it is we document ObjectStores) in the github (head/trunk/master whatever you want to call it) version. Note, do not add it to the documentation under docs/ as that is platform docs and this ObjectStore feature is not supported in EAP	Ask req clarification	https://issues.jboss.org/browse/JBTM-852
315	JBTM-846	Further to previous comment ... for instance what do we (or customers) expect from the statistics gathering routines in this new environment? Should it be measuring all transactions created and terminated in the JVM, or on a per transaction manager basis?	Ask req clarification	https://issues.jboss.org/browse/JBTM-846
316	JBTM-846	To address Davids requirement we need to support 1xArjunaCore (for the MSC) and 1xArjunaJTA ArjunaJTS (for the EJB container). The problem is the recovery manager, are we expecting to launch two of those in the same VM? I presume we have to in order to cleanly separate the MSC from the EJB container (e.g. different object stores)	Ask req clarification	https://issues.jboss.org/browse/JBTM-846
317	JBTM-846	Why would we need to recovery managers? As long as the transactions are unique and things like the nodename are shared by recovery, it should be possible to have one instance.	Ask reason/use case of req	https://issues.jboss.org/browse/JBTM-846
318	JBTM-846	Allowing multiple TMs in the same JVM is a generic statement and although WF needs this, there may be other users who want it for different reasons. I think we should either define this as a generic problem which we can try to solve for other use cases (do we know them all?) or we should refocus this solely on the WF use case and then come back to a more general requirement if/when it comes back up.	Ask completeness of req	https://issues.jboss.org/browse/JBTM-846
319	JBTM-845	Isn't this duplicated by JBTM-1451?	Ask req duplication	https://issues.jboss.org/browse/JBTM-845
320	JBTM-829	Hi, Please can you confirm that I am right in thinking you mean that we should remove the version numbers from the remaining dependencies in <code>maven/pom-template.xml</code> where they have a corresponding version number in the AS7 pom tree? Ta	Ask req clarification	https://issues.jboss.org/browse/JBTM-829
321	JBTM-817	I'm in the process of validating the trace output from these tests and I think doing this automatically is not going to be a trivial process. I have some ideas around how the byteman scripts can be re-factored to simplify the validation process and to also improve scalability so that we can introduce more complex, multi-party tests. Therefore I am removing the requirement to validate the trace output from this issue. Validation will come in a later release. Hopefully this will turn this issue into a much more straight forward task. So to confirm, the requirements for this issue are to: # Run the tests as part of the AS7 test suite (as per Ivo's pattern: https://github.com/istudens/jboss-as/compare/xts_tests#diff-0) # Automate the running of each crash recovery test without any Byteman rules. This provides some test coverage and ensures we have a valid system before the Byteman rules are applied. # Automate each valid combination of Byteman script to TestCase. The Byteman scripts specify which tests are available. # Gather the output from all test runs in separate files for "eyeball validation". Until this validation is automated, I (with the help of Andrew initially) will do this validation. The output files should be named "<byteman script name>:<test class name>". For example: "ATCrashDuringCommit.MultiParticipantPrepareAndCommitTest". Amos: does this make sense? Is this in line with how you are currently implementing the feature?	Ask req clarification	https://issues.jboss.org/browse/JBTM-817

#	Requirement	Question	Justification	URL
322	JBTM-817	Also, I don't think we need Arquillian to crash the app server as Byteman is doing that for us. This is good as Arquillian is unlikely to have this feature as part of a JBoss AS release soon enough. Amos: Does this make sense, or am I missing something?	Ask req clarification	https://issues.jboss.org/browse/JBTM-817
323	JBTM-813	Email from Jonathan re ArjunaTA: On 10/20/2010 11:39 PM, Misty Stanley-Jones wrote: > However, what would be IMMENSELY helpful is for someone to go through the Installation and Administration guides and highlight differences for me. well the JTA ones are simple... Install Guide: identical to core except for using ArjunaTA instead of ArjunaCore as the name throughout. Assuming that token is already parameterized in the core's docbook.xml the job's pretty much done already. P.S. IIRC the ArjunaTA name arose because of some legal hassle with using the term 'JTA' in that context. Anyone happen to know if that's still an issue and if so, have any good ideas for a better name? Admin Guide: as above, except: - extra section: 'prerequisites' basically saying you need to read the core docs too. Eventually we'll actually merge the docs and that paragraph will vanish but for now it probably needs keeping. - extra list entry in the xml code block for the RecoveryEnvironmentBean.recoveryExtensions in the periodic recovery section. - extra para at the end: 'selecting the JTA implementation' I'll take a look at the JTS ones tomorrow. Jonathan.	Ask req clarification	https://issues.jboss.org/browse/JBTM-813
324	JBTM-813	Email from Jonathan re ArjunaJTS: JTS install guide is same as core except name and: - add first and last bullet point paras for 'preparing your system' section - add 'specifying the ORB to use' and 'the interface repository' sections' - the 'loading the distribution' section is obsolete, just junk it. - add 'installing into non-default location' section - operating system services needs additional note on PUT_JACORB_HOME_HERE - additional 'web services component' section. JTS admin guide is same as core except name and: - intro paragraph adds mention of web services - add 'OTS/J2EE transaction service management' (except it's now 'Java EE' not 'J2EE') - add 'web service transaction management' section - in periodic recovery section, RecoveryEnvironmentBean.recoveryExtensions module list is wrong for JTS - it should be per the xml in ArjunaJTS/jbossts-properties-arjunajts.xml - add 'orb specific configurations' section - add 'initialising JBossTS applications' section side note: the docs dialect is actually 'en_GB' not 'en_US', hence 'initialising' not 'initializing' and such. What is the company policy on this - do we need to standardise it? The docbook layout seems to assume the canonical version is en_US. Likewise A4 vs. US letter for pdf page size. Jonathan.	Ask req clarification	https://issues.jboss.org/browse/JBTM-813
325	JBTM-811	Release Note: Cause: The finalizer was not needed in the StateManager. Consequence: ?? Fix: The finalizer was removed from the StateManager. Result: The finalizer is no longer in the StateManager.	Ask req clarification	https://issues.jboss.org/browse/JBTM-811
326	JBTM-810	The controlling property is com.arjuna.ats.arjuna.recovery.recoveryListener (in bean RecoveryEnvironmentBean). With a socket based recovery manager the TM/RM can easily detect whether another RM is managing an object store. Changing the default puts the onus on the user to make sure that only one RM is managing it. So my question is: is it the intention to make the user responsible for getting it right?	Ask intention of existing code/function	https://issues.jboss.org/browse/JBTM-810
327	JBTM-810	Would the same problem exist if you opened up two RMs on different ports but against the same ObjectStore? If so, +1 on the user getting it right.	Ask req clarification	https://issues.jboss.org/browse/JBTM-810
328	JBTM-809	Is there any documentation that explains what is going on here? Will JBossTS be still usable for JDBC connections? In what would consist the migration? We're using JBossTS in a JTA infrastructure where database queries are done through JDBC. The client code just uses JDBC connections (obtained through the transactional driver), while the global JTA transaction is controlled using Spring JTA transaction manager. The JBossTS backend, through the transactional driver and some of our glue code, is obtaining JTA connections through XA data sources created using the DynamicClass mechanism. We're working outside JEE, just plain Tomcat + JBossTS + Spring. How will the scenario change with JCA?	Ask req clarification & documentation	https://issues.jboss.org/browse/JBTM-809
329	JBTM-785	Hi Paul, Is there any update on this one? Is it covered by any QS? Tom	Ask updates	https://issues.jboss.org/browse/JBTM-785
330	JBTM-785	Hi [~paul.robinson], Is there any update on this? We can't understand what is needed here. What is a standalone client in this context? Tom	Ask updates	https://issues.jboss.org/browse/JBTM-785
331	JBTM-741	[~tomjenkinson], what would you think about postponing this issue until JBTM-2571 is resolved? I think it'd be too dangerous to remove public classes or methods without having code coverage from all available tests. What do you think?	Ask status of related req	https://issues.jboss.org/browse/JBTM-741
332	JBTM-728	Are you sure the remote tests are actually running? Looks to me like they are turned off, hence no coverage. <target name="run.tests"> ... <exclude name="**/remote/**"/> ... </target>	Ask req clarification	https://issues.jboss.org/browse/JBTM-728
333	JBTM-713	Marks description from the duplicate issue: {quote} We're working on WS-BA based compensation annotations. That's great, but can we refactor so that we can support compensations outside of WS-BA? ArjunaCore can support compensations without Web Services. Maybe the WS-BA annotations vanish entirely, or collapse down to a subset of what currently exists to support what is truly unique about WS-BA? {quote}	Ask req clarification	https://issues.jboss.org/browse/JBTM-713

#	Requirement	Question	Justification	URL
334	JBTM-700	I looked at this from the perspective of EAP, i.e., what docs are needed there versus only needed in the community. TS docs outline for EAP distribution: In: 0..1 InstallGuide 1..1 ReleaseNodes (JIRA generated) FailureRecovery Guide Administration Guide (ArjunaCore as basis, with diffs from JTA and JTS) Programmers Guide (ArjunaCore with all TXOJ aspects removed; start with JTA) Separate Programmers guide for JTS and XTS (maybe separate Admin guides then?) Separate all JTS and XTS information into separate documents? JTS Quickstart Guide (should be expanded to cover JTA) Configuration guide (Administration Guide update) Out: Debugging documentation All TXOJ aspects from documentation Maybe: None. Duplicated: Tooling discussions Debugging discussions ObjectStore setup Errors	Ask req clarification	https://issues.jboss.org/browse/JBTM-700
335	JBTM-659	Hi Mike, Trying to work out whether this is something we want to support on 4.16.x branch right now. Do you know if it is fixed in 4.16.x yet? e.g. hoping to put 4.16.1 out on Friday and hopefully that would be the last version for EAP6. Tom	ask version details	https://issues.jboss.org/browse/JBTM-659
336	JBTM-659	When I come to release 4.16.1 it is going to ask me what to do with this bug, is it in a state that it can be closed or does it need to move to 4.16.2 (or ideally 5.0.0.M2)?	Ask fixed version	https://issues.jboss.org/browse/JBTM-659
337	JBTM-659	Is this bug resolved in 4.6.1.CP13?	Ask fixed version	https://issues.jboss.org/browse/JBTM-659
338	JBTM-659	I fixed it in a private branch of JBOSSTS_4_6_1_GA_CP12. I am not aware of any CP13 branch. The details are in https://issues.jboss.org/browse/JBTM-522 including a link to the hudson job that tests the functionality. Has there been a decision on whether the backport will go into a 4.6.1 release? When/if it does then we should ask QA to include it in there test plan.	Ask fixed version	https://issues.jboss.org/browse/JBTM-659
339	JBTM-659	Can we mark this as complete?	Ask status	https://issues.jboss.org/browse/JBTM-659
340	JBTM-636	Having BasicAction put a timestamp in to its serialized data should work. Tricky questions like: should it be the time the tx was started? prepared? last written out, i.e. may change if the tx is partially recovered and then rewritten to the store? But the basic read/write of the timestamp should be pretty straightforward.	Ask req clarification	https://issues.jboss.org/browse/JBTM-636
341	JBTM-595	LogFactoryInterface yields a AbstractLogInterface - how does one get from LogI18n to AbstractLogInterface? I confess this logging infrastructure is pretty baffling to me	Ask req clarification	https://issues.jboss.org/browse/JBTM-595
342	JBTM-595	I stuck with LogI18n initially because it was unclear how LogFactory can receive resource bundle information. I guess the key to solving resource bundles will indeed be some kind of nifty subclassing as we are limited to a single resource bundle instance per category. Would you consider simply moving to java.util.logging wholesale, and removing all the facade stuff completely? And by "you" I guess I mean "me", since it's my dumb idea. The only practical difficulty I see would be that I'd have to add a Level class so that you have DEBUG/TRACE instead of JUL's FINE/FINER/FINEST, but that's not too big a deal. The upside would be removing the dep on commons-logging and deleting a bunch of classes. At first glance I think the resource generation stuff would work OK with JUL.	Suggest solution	https://issues.jboss.org/browse/JBTM-595
343	JBTM-575	Hi Richard This does not make much sense to me. JBAS-4238 and JBTM-66 are similar but not equivalent. The JBAS issue is addressed in both releases, but JBossAS 4.2.3 uses TS 4.2.3.SP7, which does not have the JBTM-66 changes. JBossTS AS 5.1.0.GA uses TS 4.6.1.GA, which does. So, you seem to be having problems due to a behaviour change that you requested in the first place?! Ah well, maybe I misunderstood what you wanted. Do you have any preference between the implementation alternatives I outline in http://community.jboss.org/message/515473 assuming either will address your needs?	Ask req clarification	https://issues.jboss.org/browse/JBTM-575
344	JBTM-575	at com.arjuna.ats.jbosstx.BaseTransactionManagerDelegate.setRollbackOnly(BaseTransactionManagerDelegate.java:123) at org.hibernate.ejb.AbstractEntityManagerImpl.markAsRollback(AbstractEntityManagerImpl.java:421) at org.hibernate.ejb.AbstractEntityManagerImpl.throwPersistenceException(AbstractEntityManagerImpl.java:576) at org.hibernate.ejb.AbstractEntityManagerImpl.throwPersistenceException(AbstractEntityManagerImpl.java:614) at org.hibernate.ejb.AbstractEntityManagerImpl\$1.beforeCompletion(AbstractEntityManagerImpl.java:513) What's key here is that, from the transaction manager's perspective, the setRollbackOnly call precedes the exception - hibernate sets the tx to rollback before throwing the exception. So, the setRollbackOnly call IS the root cause of the transaction rolling back. JBTM-66 is not only present but functioning exactly as designed. As mentioned in the forum thread, this is arguably misuse of the api by hibernate. Of the alternatives proposed in the forum thread, chaining the deferredThrowable to the rollbackOnly trace won't work, as there is no deferredThrowable at the point in time that the setRollbackOnly call is made. So the only viable alternative seems to be using the deferredThrowable as the initCause of the rollback in preference to the rollbackOnlyCallerStackTrace in cases where both are available. Will that meet your needs?	Suggest solution	https://issues.jboss.org/browse/JBTM-575
345	JBTM-575	Jonathan, Actually I have no problem asking Hibernate to stop messing with causality, if that is what is happening. Perhaps this broke in JBoss 5.1.0 not because of a JBTM change but because of a Hibernate change? Could you please tell me which classes (even line numbers?) are suspect within the Hibernate code, and I will raise a JIRA with that team. Regards, Richard.	Ask code	https://issues.jboss.org/browse/JBTM-575

#	Requirement	Question	Justification	URL
346	JBTM-575	Johathan, can I ask you another question on a related topic(one that does not involve the client at all). If a business method in EJB session bean throws a RuntimeException in the JBoss log I see only this line: 11:32:12,550 ERROR [TxPolicy] javax.ejb.EJBTransactionRolledbackException No cause, no message, no nothing - the RuntimeException is however successfully relayed to the client, but I remember that in older JBoss version I used to get I full stack trace with the javax.ejb.EJBTransactionRolledbackException. What might be causing this problem?	Ask req clarification	https://issues.jboss.org/browse/JBTM-575
347	JBTM-575	Jonathan, As you asked, I asked someone from the Hibernate team (Steve Ebersole) about causality. Steve has since responded, as have a couple other people. What is our next step to getting this resolved? Regards, Richard.	Ask solution	https://issues.jboss.org/browse/JBTM-575
348	JBTM-575	Was this fixed or not? I can't tell from the JIRA or the forum post. Assuming that it was, I will write this release note: Cause: In ArjunaCore which was included in JBoss Application Server 4.2 and previous, when a transaction rolled back, the <systemitem>RollbackException</systemitem> included the root cause of the exception. In JBoss Application Server 5.1, the exception no longer included this information, even though the information could still be found in a different area of the log file. Consequence: Developers who relied on the information included in the <systemitem>RollbackException</systemitem> in the older versions of JBoss Application Server found that their applications no longer worked as expected. Fix: The debugging information has been put back into the exception, so that it works as it did previously. Result: Developers can once again write code to take advantage of the debugging information found in the <systemitem>RollbackException</systemitem>.	Ask status	https://issues.jboss.org/browse/JBTM-575
349	JBTM-573	This says "Done" but there is no information. Is it done or WONTFIX? If it is done, what is the command-line tool and what is the syntax for using it?	Ask status	https://issues.jboss.org/browse/JBTM-573
350	JBTM-535	Another thought about the Listener.finalize() method: it is useless also because if the Listener is running, it can't be finalized, so finalize() will never be called. Viceversa, if finalize() is called it means that the listener has already been stopped, so its finalize() implementation is useless... am I wrong?	Ask req clarification	https://issues.jboss.org/browse/JBTM-535
351	JBTM-535	Could this be targetted for 4.7.0? Resolution should be quite straightforward... Thanks in advance.	Ask fixed version	https://issues.jboss.org/browse/JBTM-535
352	JBTM-535	Given that right now this problem is causing a memory leak whenever JBossJTA is restarted, any thought on a candidate target release? Thanks in advance.	Ask fixed version	https://issues.jboss.org/browse/JBTM-535
353	JBTM-527	Hi Romain, If you still would like this fix, please can you consider raising a pull request on the Narayana project and we can take a look at it. The new build system uses maven so you could add pmd into a profile and don't enable that profile by default? Thanks, Tom PS If you no longer require this feature please consider closing this issue and marking as deferred	Ask code (PR)	https://issues.jboss.org/browse/JBTM-527
354	JBTM-503	Hi Andrew, any news on this? Could it be considered for 4.7.0?	Ask status & fixed version	https://issues.jboss.org/browse/JBTM-503
355	JBTM-465	Please can you review this in context of JBTM-1230, I think the changes we made there might have addressed your requirements	Ask code review & RIM	https://issues.jboss.org/browse/JBTM-465
356	JBTM-465	Hi Tom, please let me understand something. From JBTM-1230 I understand that JDBC connections for the Object Store are now retrieved using JNDI, is this correct? If so, as I already stated already in other circumstances, this is a problem for us, since we're using JBoss TS embedded in an application which is meant to work outside a JEE environment. We would need to setup a fake JNDI repository with Spring to let the object store find JDBC connections. Do you confirm this? Or is there an alternative way to still provide connections without JNDI? If not, I personally consider this, together with the push towards the mandatory use of a JCA connector (in place of the plain JDBC implementation) in JBoss TS 5.x, a step backward. Please note that I've not followed this evolution so tightly, but this is what I understood based on the feedback I received on the bugs I opened myself (and for some of which I also contributed with patches that were never integrated). Thanks in advance.	Ask req clarification	https://issues.jboss.org/browse/JBTM-465
357	JBTM-465	Hi Mauro, I found the link, it doesn't seem like you have signed the CLA, please can you consider doing to assist in integrating your future patches? https://cla.jboss.org (choose "JBoss Transactions") Thanks, Tom	Ask some one do sth	https://issues.jboss.org/browse/JBTM-465
358	JBTM-465	Hi Tom, thanks a lot for your feedback. Two questions: - I'm evaluating the CLA with my employer to sign a corporate CLA, thank you; just a question: in the form at http://cla.jboss.org there's also a text field where I can specify a JIRA ticket; should I need to fill in the form for each ticket I want to contribute or is it just a hint to the first/most important ticket I want to contribute? - is (or will be) there any documentation available where there is an explanation of the migration path from JBoss TS 4.x to 5.x? I mean, from a system integrator point of view, rather than from a consumer point of view (as I said, we're embedding JBoss TS in our own application and we had to integrate it in our Spring+JDBC+pooling environment - not a 2-minute trivial task) Thanks in advance.	Ask req clarification	https://issues.jboss.org/browse/JBTM-465
359	JBTM-437	Jonathan, can you take this?	Ask some one take look at req	https://issues.jboss.org/browse/JBTM-437

#	Requirement	Question	Justification	URL
360	JBTM-435	Since I'm going thru this now, I'll summarize what I was looking for in the docs that this JIRA can hopefully address. I have a J2EE app running in JBossAS 4.2 - it can support either PostgreSQL or Oracle (I deploy a custom datasource in my own -ds.xml file) Because I am not limited to just Oracle, I want to use JDBCXARecovery (the title of this JIRA mentioned "OracleXARecovery" - so I mention this so we can have docs for the recoverer that uses JNDI to obtain the XADatasource). I need to know what settings I have to put in jbossjta-properties.xml to enable XA-recovery for my datasource. This documentation is more targeted to the JBossAS integration with JBossTM, since involves talking about what settings to put in the conf/jbossjts-properties.xml.	Ask code (settings)	https://issues.jboss.org/browse/JBTM-435
361	JBTM-400	Done: JBTM-1472. Can you add some more detail to the description? As it stands, my understanding of this issue is that it's a duplicate of JBTM-1472.	Ask req clarification	https://issues.jboss.org/browse/JBTM-400
362	JBTM-389	Which svn tag are your src changes based on?	Ask code	https://issues.jboss.org/browse/JBTM-389
363	JBTM-369	Can you give a use case for this? I'm not seeing any real advantage over just doing e.g. -Dcom.arjuna.ats.arjuna.xa.nodeIdentifier=foo ?	Ask reason/use case of req	https://issues.jboss.org/browse/JBTM-369
364	JBTM-369	Jonathan, How to end-up with a "smooth" integration within the AS? I understand that the TM has its own configuration file when working inside some other runtime, but when running inside the AS, can't he use the JBoss config files (like jboss-services.xml file) or similar? If not, can we at least implement the same feature set as JBoss AS has, including the property replacer that Roland was suggesting to add? Thanks for your feedback. Cheers, sach	Suggest solution	https://issues.jboss.org/browse/JBTM-369
365	JBTM-369	Hi Sacha Please clearly define what you are trying to achieve. So far nobody has given me a convincing, specific real world use case for this feature. As far as I can tell the existing ability to override values from the config file by specifying system properties is equivalent to value substitution, except for cases where you may want to use the same value for multiple parameters. The only such case I can think of is the IP address, for which we have a separate solution anyhow. With the exceptions of the port bindings (which will be done via SBM), recovery (which will be done via the new tx integration API) and multiple 1PC (which is an horrifically bad idea anyhow), users should not need to change values in the tx properties file. I believe this feature will benefit only a small proportion of users. I agree it's better to have consistent behaviour for the properties files, it makes for a better user experience. But that's secondary to having a product that works. Because of these factors, I consider this issue a low priority. It will get done as and when we have time, but not until e.g. crash recovery and JTS integration are finished. That's why it's still not scheduled against any particular release. Perhaps I'm missing something though. If you feel it's worth taking time out from managing the division to comment on the issue, you must feel it's important. Can you say why?	Ask reason/use case of req	https://issues.jboss.org/browse/JBTM-369
366	JBTM-355	Can you please explain why the value of this timeout needs to be configured with a larger value and how this value relates to your transaction's 'size'? Also, by 'size' are you referring to the number of participants or the amount of work done by each participant? The transport timeout is only ever employed when dispatch of a WS-C/T protocol message must wait on receipt of a corresponding protocol message. For example, the WS-AT Coordinator engine responds to a prepare request by sending a PREPARE message to the participant and then waiting for a PREPARING message to be returned in response. Processing of the incoming message and dispatching the corresponding response will normally happen very quickly as it only involves a small amount of work on the part of the transaction system. The other variables in the equation are the communication transport delays and, in some cases, a certain amount of processing time added by the participant implementations (in the example above the participant's prepare method will be invoked). The participant may conceivably be delaying the response for more than 30 seconds. If so then this is likely to be due to an error in the participant.	Ask req clarification	https://issues.jboss.org/browse/JBTM-355
367	JBTM-355	I'm still not convinced that configuring this timeout is the correct thing to do. Can you explain i) which database you are using and ii) what your code is doing in the prepare call? I am assuming that this problem arises because entities created in your subordinate TX only get flushed at prepare time due to the delegation involved in bridging. The persistence manager install a Synchronization to flush cached SQL updates to the database. This normally runs before the 2PC is entered but with bridging the flush cannot be performed (i.e. the Synchronization cannot be executed) until the delegated prepare arrives which stalls the parent TX. One possibility might be for you to manually flush changes inside your EJB code as you make them. This should leave little or no work to do at prepare time because the DB updates will already be out on disk. Can you try this? Another option might be to do break the create operations into lots of little transactions. Do you have to do this update all in one go? Anyway leaving these questions aside for the moment, if you want to reset the timeout you can easily do so by invoking TransportTimer.setTransportTimeout(long) from your application code.	Ask req clarification	https://issues.jboss.org/browse/JBTM-355
368	JBTM-345	What happens if a module goes away between first and second passes, or during a recovery pass? Do we need blocking unregister calls for modules so that does not happen?	Ask req clarification	https://issues.jboss.org/browse/JBTM-345
369	JBTM-334	Just a clarification: this is not about deprecating what NORMAL does, just changing the name from NORMAL to PERIODIC. Correct?	Ask req clarification	https://issues.jboss.org/browse/JBTM-334

#	Requirement	Question	Justification	URL
370	JBTM-331	See http://jira.jboss.com/jira/browse/JBTM-20?page=all for the backstory on this. We already downgrade this stuff from INFO to DEBUG when running integrated with the AS. The arjuna level won't be changed, so the options are to change the global rewriting of arjuna messages in AS from INFO->DEBUG to INFO->TRACE which will impact other messages too, or install a filter in log4j. Why is AS even logging at DEBUG by default? It should be INFO, at least for production releases.	Ask req clarification	https://issues.jboss.org/browse/JBTM-331
371	JBTM-319	Yes, we definitely don't want to duplicate code from elsewhere. How about getting the ManagedConnectionFactory instance from JNDI, since that's what creates the WrapperDataSource in the first place?	Suggest solution	https://issues.jboss.org/browse/JBTM-319
372	JBTM-319	Or the ConnectionFactory?	Suggest solution	https://issues.jboss.org/browse/JBTM-319
373	JBTM-319	Can you take any queries and discussion to the forums, where it's easier to track? Thanks.	ask to talk more in forum	https://issues.jboss.org/browse/JBTM-319
374	JBTM-312	Would this work as a general design pattern if you wanted to implement a HA recovery coordinator? 1. Use the JDBC object store. 2. Configure a recovery coordinaor (RC) on each node to recover all (*) nodes. You can only *start* one of these RC at any one time 3. Develop some "watchdog" process runs on every node and monitors the current RC process and if it fails, uses a group protocol to elects one of the other RC's to start 4. Rather than starting the RC on one of the nodes, start the watchdog processes and allow them to enable one of the RC's Would this work for JTS?	Ask req clarification	https://issues.jboss.org/browse/JBTM-312
375	JBTM-300	Thank you for backporting. I put 4.2.3.SP5_CP07-brew jars into EAP 4.3.0 CP06 (SOA-P 4.3.0 CP02) and tested AtomicActionExpiryScanner. Then I found that expired tx-logs were copied to "Expired" directory but the original files were not removed. > find data/tx-object-store/HashedActionStore/defaultStore/StateManager/BasicAction/TwoPhaseCoordinator/AtomicAction/#197#/7f000001_9020_4b025a1e_2f9 data/tx-object-store/HashedActionStore/defaultStore/StateManager/BasicAction/TwoPhaseCoordinator/AtomicAction/Expired/#197#/7f000001_9020_4b025a1e_2f9 ... It's hitting JBTM-418. Could you take care of JBTM-418 for 4.2.3 branches too?	Ask some one take care of req	https://issues.jboss.org/browse/JBTM-300
376	JBTM-280	Do you mean simply saving them as .odt files? I tried to check one in but got the following error : svn: CHECKOUT of '/!svn/ver/15549/labs/jbosstm/trunk/ArjunaCore/docs/install_guide': 403 Forbidden (https://svn.labs.jboss.com) Am I using the wrong path or don't I have the correct permissions to commit perhaps. Thanks, Daniel	Ask req clarification	https://issues.jboss.org/browse/JBTM-280
377	JBTM-264	Is this resolved as part of the CLI work possibly? e.g. can the AS7 CLI be deployed externally?	Ask req clarification	https://issues.jboss.org/browse/JBTM-264
378	JBTM-264	Hi Mike, This issue was marked as "Resolved, Done" but there is no "Fix Version" specified. Please can you update the Jira with the version of the software where you fixed this? Also, can you then close the issue (assuming its not an in development version). Thanks, Tom	Ask fixed version	https://issues.jboss.org/browse/JBTM-264
379	JBTM-252	Daniel, would you please clarify the scope of the quantifiers 'some' and 'all' in your issue description? i.e. in which documents have you noticed the format problem? which other documents do you suspect are affected?	Ask req clarification	https://issues.jboss.org/browse/JBTM-252
380	JBTM-252	Hi Andrew, As I'm planning on going through the docs anyway, I'd be happy to take this JIRA and update the doc if you like? /Daniel	Suggestion	https://issues.jboss.org/browse/JBTM-252
381	JBTM-250	What about JTXA BaseTransaction.java?	Ask req clarification	https://issues.jboss.org/browse/JBTM-250
382	JBTM-204	http://www.nabble.com/Spring---JBosTS-integration-t3901668.html You might want to mark this as Started ;-)?	Suggestion	https://issues.jboss.org/browse/JBTM-204
383	JBTM-144	Mark The DTF QA tests have their own (JBoss internal) repo, due to different licencing for the DTF. Can we close this one?	Ask status	https://issues.jboss.org/browse/JBTM-144
384	JBTM-116	Which lower level? JBossTS understands -1. It means "transactions never time out".	Ask req clarification	https://issues.jboss.org/browse/JBTM-116
385	JBTM-103	Mark I'm not clear on this one. The effective default timeout appears to be 0 regardless of which TransactionManagerImple is used. Am I missing something?	Ask req clarification	https://issues.jboss.org/browse/JBTM-103
386	JBTM-98	what is the setup of your test? If I have a single JBoss server talking to oracle then it works fine. When I try the following scenario is when I get the problem (both server 4.0.3 SP1 and JBossTS 4.2): Server1 : starts the global distributed transaction, performs local oracle database operations and calls a server2Bean on server2 // this method has the tx "required" attribute in ejb-jar.xml performGlobalTx(){ doSomething(); // do oracle operations server2Bean.doSomething(); // call remote jboss Server2 } Server2: the bean on this server updates the same oracle database: doSomething(){ performDBOperations...; // here is where I get the exceptions reported in the forum <<<<-----here }	Ask code (testing)	https://issues.jboss.org/browse/JBTM-98
387	JBTM-95	We have such a thing in place now as far as I can see. Let's close this?	Suggest change status of req	https://issues.jboss.org/browse/JBTM-95
388	JBTM-86	Could you please increase the priority of this bug, or at least give it some sort of schedule? Defect analysis becomes a nightmare without any sort of exception chaining.	Ask change priority of req	https://issues.jboss.org/browse/JBTM-86
389	JBTM-86	Hi Andrew, hi Jonathan, what is the fix version that includes the whole fix? 4.14? Thanks in advance!	Ask fixed version	https://issues.jboss.org/browse/JBTM-86
390	JBTM-49	Does this still make sense? We should probably ensure our tooling is consistent with AS7 as much as anything.	Ask feasibility of req	https://issues.jboss.org/browse/JBTM-49
391	JBTM-49	Hi Mike, This issue was marked as "Resolved, Done" but there is no "Fix Version" specified. Please can you update the Jira with the version of the software where you fixed this? Also, can you then close the issue (assuming its not an in development version). Thanks, Tom	Ask fixed version	https://issues.jboss.org/browse/JBTM-49
392	JBTM-46	FYI. I am going to mark TxManager as deprecated. See the linked task. Any reason why you didn't use links?	Ask reason	https://issues.jboss.org/browse/JBTM-46

#	Requirement	Question	Justification	URL
393	JBTM-20	Can you add something to the documentation (the section around application server integration)? If you want, send me text and I'll put it in place.	Ask documentation	https://issues.jboss.org/browse/JBTM-20
394	JBTM-5	From Adrian Brock: "Adrian comment: "Will this also be integrated with the JBoss installer and the JBoss Network for "ease" of installation/choice? i.e. we include JBoss Transactions in the installer distrubtion and let them choose between the two implementations."	ask about future plans	https://issues.jboss.org/browse/JBTM-5