


AIRFLOW

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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.	Raphael Lopez Kaufman	0	0	1	I made some comments on the PR. The logic is sound, but some minor improvements wanted. I would raise it on the mailinglist as well, maybe someone solved it differently.	Bolke de Bruin	https://issues.apache.org/jira/browse/AIRFLOW-1414
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 ?	Raphael Lopez Kaufman	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-1414
AIRFLOW-1401	I'm particularly interested in thoughts from [~alexvanboxel], [~wwlian], and [~fenglu]	Peter Dolan	0	0	1	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.	Feng Lu	https://issues.apache.org/jira/browse/AIRFLOW-1401
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.	Feng Lu	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-1401
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.	Raphael Lopez Kaufman	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-1398
AIRFLOW-1398	[~bolke] I implemented your suggestion, would you mind reviewing?	Raphael Lopez Kaufman	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-1398
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.	Chris Riccomini	0	0	1	Yes airflow will only use ES if the user configures the logging_backend_url and S3/GCS won't be removed	Allison Wang	https://issues.apache.org/jira/browse/AIRFLOW-1325
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 - thats 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 dont 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?	Michal Dziemianko	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-1310
AIRFLOW-1019	[~aoen], changing this to blocker for 1.8.1. LMK if you don't think we should block the release on it.	Chris Riccomini	0	0	1	Probably not a blocker, just a performance degradation.	Dan Davydov	https://issues.apache.org/jira/browse/AIRFLOW-1019
AIRFLOW-983	Was this confirmed yet?	Bolke de Bruin	0	0	1	No, and I'll be out for a couple of weeks unfortunately.	Dan Davydov	https://issues.apache.org/jira/browse/AIRFLOW-983

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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]	Dan Davydov	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-983
AIRFLOW-979	Please review my Pull Request #2149: https://github.com/apache/incubator-airflow/pull/2149	Chris Sng	0	0	1	Issue resolved by pull request #2149 https://github.com/apache/incubator-airflow/pull/2149	Jeremiah Lowin	https://issues.apache.org/jira/browse/AIRFLOW-979
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. le. please explain why this is a blocker.	Bolke de Bruin	0	0	1	I'm still investigating and will update the ticket with more details as soon as I am able to. My first guess is that the regex field of backfills makes Airflow think that a task doesn't exist anymore if it's not included in the regex or something like that. This is a blocker because it overwrites existing task instances that were e.g. successful with the removed state.	Dan Davydov	https://issues.apache.org/jira/browse/AIRFLOW-932
AIRFLOW-932	Command: {quote}airflow backfill testg -t create_and_transfer_roster -s 2017-06-14 -e 2017-06-14{quote} DAG: {code} import airflow from airflow.executors import SequentialExecutor from airflow.providers.dummy_operator import DummyOperator from airflow.providers.python_operator import PythonOperator from airflow.providers.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) {code} 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) please	Dan Davydov	0	0	1	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. {code} MariaDB [airflow]> select * from task_instance; +-----+-----+-----+-----+-----+ +-----+-----+-----+-----+-----+ +-----+-----+-----+-----+-----+ +-----+-----+-----+-----+-----+ ----+ task_id dag_id execution_date start_date end_date duration state try_number hostname username job_id pool queue priority_weight operator queued_dttm pid +-----+-----+ +-----+-----+-----+-----+-----+ +-----+-----+-----+-----+-----+ +-----+-----+-----+-----+-----+ -----+-----+-----+ create_and_transfer_roster testg 2016-06-14 00:00:00.000000 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 +-----+-----+ +-----+-----+-----+-----+-----+	Bolke de Bruin	https://issues.apache.org/jira/browse/AIRFLOW-932

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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;</pre>  <pre>{code} create_and_transfer_roster testg 2016-06-14 00:00:00.000000 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 </pre>	Bolke de Bruin	0	0	1	This is off of master (88d9b0dc96e7528c87326c8070ee276e856554f). I just reproduced again (you are right about the start date though, should have been 2016).	Dan Davydov	https://issues.apache.org/jira/browse/AIRFLOW-932
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 airflow backfill testtg -t create_and_transfer_roster -s 2016-03-15 -e 2016-03-15	Dan Davydov	0	0	1	Ok. Confirmed now - really strange.	Bolke de Bruin	https://issues.apache.org/jira/browse/AIRFLOW-932
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> {code} 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). get_task_instances 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?	Bolke de Bruin	0	0	1	I think tasks should only be verified/removed in the scheduler dagrun loop. That would avoid this issue if I understand correctly.	Dan Davydov	https://issues.apache.org/jira/browse/AIRFLOW-932
AIRFLOW-932	PR is available with unit tests ["~aoen"]. Please review	Bolke de Bruin	0	0	1	Issue resolved by pull request #2122	Bolke de Bruin	https://issues.apache.org/jira/browse/AIRFLOW-932

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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)	Vijay Krishna Ramesh	0	1	0	https://github.com/apache/incubator-airflow/pull/2122		https://issues.apache.org/jira/browse/AIRFLOW-931
AIRFLOW-920	[~aoen] 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?	Bolke de Bruin	0	0	1	Bolke de Bruin I'm not seeing this behavior. Marking the SubDagOperator task instance successful doesn't mark the tasks in the subdag for that execution dat as successful. Let me know if I'm missing something.	Dan Davydov	https://issues.apache.org/jira/browse/AIRFLOW-920
AIRFLOW-908	Can i take this feature ?	Phawin Khongkhasawan	1	0	0	I'm not seeing this behavior. Marking the SubDagOperator task instance successful doesn't mark the tasks in the subdag for that execution dat as successful. Let me know if I'm missing something.		https://issues.apache.org/jira/browse/AIRFLOW-908
AIRFLOW-908	[~aoen] Would you give me an instruction ?	Phawin Khongkhasawan	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-908
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	Amit Pahwa	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-891
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?	James Santucci	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-890
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	Ultrabug	0	0	1	Hi Ultrabug, We are on RC5 now, and will release to PyPI once the current blockers are cleared, and a new vote on the release is taken. All of this combined might take nother couple of weeks. Best, Arthur	Arthur Wiedmer	https://issues.apache.org/jira/browse/AIRFLOW-846
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...	Ultrabug	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-846
AIRFLOW-679	I see this has a closing PR. Is this ticket still relevant or can it be closed?	Alex Van Boxel	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-679
AIRFLOW-539	Not sure if you saw the note on the related issue, but perhaps this could be implemented via a kwargs passed through to the API, so that one-off feature changes don't require end-to-end plumbing in the future?	Sam McVeety	0	0	1	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?	Chris Riccomini	https://issues.apache.org/jira/browse/AIRFLOW-539

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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?	Chris Riccomini	0	0	1	Opened https://issues.apache.org/jira/browse/AIRFLOW-559 .	Sam McVeety	https://issues.apache.org/jira/browse/AIRFLOW-539
AIRFLOW-539	[~criccomini]: although this ticket is closed (can you add a "gcp" label on this ticket), you're probably an admin. Thanks.	Alex Van Boxel	0	0	1	Chris Riccomini made changes - 11/Oct/16 15:24 Labels gcp	Chris Riccomini	https://issues.apache.org/jira/browse/AIRFLOW-539
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.	George Leslie-Waksman	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-513
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	Kevin Yuen	0	0	1	I think that sensors, including ExternalTaskSensors, should be counted in the core parallelism limit; they are after all using resources. I think the resource bottlenecks long-waiting ExternalTaskSensors can create should be managed with pools/queues or actual DAG dependencies. To the latter point I think in your case you may want to use the TriggerDagRunOperator instead so you are pushing to create a DagRun instance for DAG 2, instead of polling from DAG 2 for DAG 1. We do use some ExternalTaskSensors in the way you describe but we increase worker limits each time we add one, and in any case we have only 4 running at a time, so it hasn't reached your level of contention yet.	Laura Lorenz	https://issues.apache.org/jira/browse/AIRFLOW-513
AIRFLOW-511	[~criccomini] any thoughts on this?	Rob Froetscher	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-511
AIRFLOW-479	Can I take this issue?	Yongjun Park	0	0	1	Yes, please do. Thank you.	Jason DeCorte	https://issues.apache.org/jira/browse/AIRFLOW-479
AIRFLOW-461	This is for GCS to BQ operator, right?	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-461
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.	Siddharth Anand	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-454
AIRFLOW-445	[~jlowin][~criccomini] What do you think?	Siddharth Anand	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-445
AIRFLOW-409	[~mt dewulf] 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?	Robin Miller	0	0	1	Of course. Thank you.	Michael DeWulf	https://issues.apache.org/jira/browse/AIRFLOW-409

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
AIRFLOW-382	"=+1 [~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"	Ruslan Dautkhanov	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-382
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?	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-374
AIRFLOW-350	Could you add some more description to this ticket? I don't understand it.	Chris Riccomini	0	0	1	Close this since this is actually configurable. Sorry	Tongzhou Wang	https://issues.apache.org/jira/browse/AIRFLOW-350
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?	Kevin Deldycke	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-340
AIRFLOW-331	Ah, nice catch. Can you send a pull request, please?	Chris Riccomini	0	0	1	Commit 7dbc3cd40e5c4f072897bcae46929ad61a2b4590 in incubator-airflow's branch refs/heads/master from Peter Pang [https://git-wip-us.apache.org/repos/asf?p=incubator-airflow.git;h=7dbc3cd] AIRFLOW-331 modify the LDAP authentication config lines in 'Security' sample codes Closes #1674 from impangt/master	ASF subversion and git services	https://issues.apache.org/jira/browse/AIRFLOW-331
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?	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-329
AIRFLOW-325	Isn't this possible with an ExternalTaskSensor(same_dag_id, previous_task_id, allowed_states=[State.SUCCESS, State.FAILED])?	Matthias Huschle	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-325
AIRFLOW-312	Can you try running off of master? I'm wondering if my PR in AIRFLOW-200 fixes this.	Chris Riccomini	0	0	1	"OK, my bad. Thanks for your explanation."&"Right now I was just fighting to get BigQueryToCloudStorageOperator working and you just helped me to get out of this mess by pointing me out airflow[xxx]. Thanks again."	dud	https://issues.apache.org/jira/browse/AIRFLOW-312
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 the scheduler.	David Klosowski	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-288
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	Chad T Henderson	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-276
AIRFLOW-275	[~bolke], you committed the PR for this, right?	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-275

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
AIRFLOW-258	Makes sense to me. Want to send a PR?	Chris Riccomini	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-258
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!	Greg Neiheisel	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-258
AIRFLOW-249	This is really awesome analysis. Can you send a PR? (unless I missed it)	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-249
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?	Rob Froetscher	0	0	1	Rob Froetscher, totally fine. Go for it. Might want to ping Arthur Wiedner, as well, though.	Chris Riccomini	https://issues.apache.org/jira/browse/AIRFLOW-247
AIRFLOW-247	[~froetscher], are you still planning to work on this? If not, we can re-assign to [~gjreda].	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-247
AIRFLOW-247	[~gjreda], want to have a look at the PR?	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-247
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/1NG1P86HRIX43qTVucTOsFqIbCvYdOhq_np90VlbVRC/edit#slide=id.gd40eeee67_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/dags/operators/emr_spark_operator.py	Al Johri	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-247

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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/f4c4146cd38afe21c517230bfff71d33ab041246/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?	Alexandr Nikitin	0	0	1	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 <code>airflow.contrib.executors</code>? 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?	Jeremiah Lowin	https://issues.apache.org/jira/browse/AIRFLOW-245
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 <code>airflow.contrib.executors</code>? 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?	Jeremiah Lowin	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-245
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.	Chris Riccomini	0	0	1	Chris Riccomini Thank you Chirs. This is exactly what I need. Problem closed.	Junwei Wang	https://issues.apache.org/jira/browse/AIRFLOW-199
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?	Sumit Maheshwari	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-177

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
AIRFLOW-169	What do you mean by expired schedules? You mean DagRuns that have finished? Which page in the UI are you referring to?	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-169
AIRFLOW-167	This sounds reasonable to me. Want to send a PR?	Chris Riccomini	0	0	1	Yup, already opened https://github.com/apache/incubator-airflow/pull/1541/files .	Sumit Maheshwari(Inactive)	https://issues.apache.org/jira/browse/AIRFLOW-167
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 <code>{{dag_state}}</code> command?	Chris Riccomini	0	0	1	Yup, already opened https://github.com/apache/incubator-airflow/pull/1541/files .	Sumit Maheshwari(Inactive)	https://issues.apache.org/jira/browse/AIRFLOW-167
AIRFLOW-165	I think this is already done, right? https://pythonhosted.org/airflow/concepts.html#task-documentation-notes You can do: <code>{code} t = BashOperator("foo", dag=dag) t.doc_md = ""\n #Title"</code> Here's a url (www.airbnb.com) <code>"""\n {code}</code>	Chris Riccomini	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-165
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.	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-161
AIRFLOW-161	I don't think that we want to embed Quobole 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?	Chris Riccomini	0	0	1	Fair enough, I can't challenge that decision, as Quobole is not as big as aws or gce .	Sumit Maheshwari(Inactive)	https://issues.apache.org/jira/browse/AIRFLOW-161
AIRFLOW-161	Sure, can you please cc top contributors to PR or here.	Sumit Maheshwari	1	0	0	However as that link will be visible only for quobole_operator type tasks, which implies that user is using quobole and having that link will help him. Also I think that airflow gonna need /redirect (or similar) route in near future.		https://issues.apache.org/jira/browse/AIRFLOW-161
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?	Chris Riccomini	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-135
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.	Chris Riccomini	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-129
AIRFLOW-128	[~bolke], OK to close this?	Chris Riccomini	0	0	1	Bolke de Bruin made changes - 02/Jun/16 07:41 Status Open [1] Closed [6] Resolution Fixed [1]	Bolke de Bruin	https://issues.apache.org/jira/browse/AIRFLOW-128
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.	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-123
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.	dud	0	0	1	Added it: https://issues.apache.org/jira/browse/AIRFLOW-625	Flex Gao	https://issues.apache.org/jira/browse/AIRFLOW-121
AIRFLOW-118	Was there a PR for this?	Chris Riccomini	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-118

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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)	Vitor Baptista	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-115
AIRFLOW-113	[~Shenghu], can you: # Test this against master? # Post your DAG (or some equivalent) Python code so we can try and replicate	Chris Riccomini	0	0	1	I will wait & test this on airflow 1.7.1, close this for now.	Shenghu Yang	https://issues.apache.org/jira/browse/AIRFLOW-113
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}	Chris Riccomini	0	0	1	I will wait & test this on airflow 1.7.1, close this for now.	Shenghu Yang	https://issues.apache.org/jira/browse/AIRFLOW-113
AIRFLOW-111	Thanks [~criccomini]. 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 ?	Shenghu Yang	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-111
AIRFLOW-101	[~bolke], which PR was this part of?	Chris Riccomini	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-101
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.	Jeremiah Lowin	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-93
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 !	Stanilovsky Evgeny	0	0	1	Your latest PR doesn't even include this line: configuration.get('webserver', 'ssl_keyfile') Which is in your stack trace (cli.py, line 689). Line 689 in your PR contains: ("r", "--report"), "Show DagBag loading report", "store_true"), I think your PR is missing some commits or something.	Chris Riccomini	https://issues.apache.org/jira/browse/AIRFLOW-91
AIRFLOW-91	i update related issue, pls take a look ones more ?	Stanilovsky Evgeny	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-91
AIRFLOW-91	hi for all, i already update PR, plz take a look on it? thanks !	Stanilovsky Evgeny	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-91
AIRFLOW-81	What do you think about writing a new sensor based off of the TimeSensor that blocks during the blackout period?	Chris Riccomini	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-81
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...	Jimbo Freedman	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-79

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
AIRFLOW-78	[~abridgett] Can you tell me more about your expectations of clear?	Siddharth Anand	0	0	1	Sure (sorry for the delay). My understanding was that it "clears" it - so that it's as if it's never run. The scheduler then comes along and says "oh, I'd better run that then" and runs the DAG (or task if you've just cleared a task within a DAG). Back when I first used airflow I expected "run" to force-run a task but that didn't seem to happen (even with "force" selected) and the code matched that (I see that it should work now if the task isn't successful or force is set now). The specific issue I had here was to do with that max_active_runs setting - I think it must have stopped any task from running at all (I don't think that the case that we had two runs simultaneously as it's unlikely I ran the clear when another run was running). The fact that I had to clear the dagrun seems to confirm this to me.	Adrian Bridgett	https://issues.apache.org/jira/browse/AIRFLOW-78
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?	Siddharth Anand	0	0	1	The idea behind 'clear' is more of a way to get task instances to re-run, by clearing their 'failed' state (problems has been fixed, we're ready to re-run) or by clearing a 'success' state ("hey, this was a false positive!", or "source data was wrong, let's rerun this"). When it is the intent, we want the scheduler to pick those tasks up and re-run them, this is why we re-activate the DagRuns. Maybe the docs need to be clarified. Maybe we need some sort of a '--dont-reactivate-dagruns' flag as well.	Maxime Beauchemin	https://issues.apache.org/jira/browse/AIRFLOW-78
AIRFLOW-78	[~abridgett]: Just to clarify, the issue here is still that the cleared dagruns get rerun without respecting the max_active_runs amount?	Norman Mu	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-78
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?	Norman Mu	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-78
AIRFLOW-71	[~asnir], can you describe what this is a bit more?	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-71
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?	Chris Riccomini	0	0	1	https://docs.docker.com/engine/reference/commandline/login/	Amikam Snir	https://issues.apache.org/jira/browse/AIRFLOW-71
AIRFLOW-71	Hi [~criccomini], 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="- ./bi_jobs/eti/example.py {{ execution_date }} {{ execution_date }}") user = 'test' pwd='pwd' if d.cli: d.cli.login(username=user, password=pwd, 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)	Felipe Tancredo	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-71
AIRFLOW-71	When will [~asnir]'s PR be approved? It's been 8 months.	Andrew Stewart	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-71
AIRFLOW-58	[~maxime.beauchemin@apache.org], any reason why this was closed?	Chris Riccomini	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-58

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
AIRFLOW-58	[~sanand], did you mean to put a real external URL in? You just put the JIRA in. ;)	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-58
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?	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-55
AIRFLOW-54	"=+1 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'"	Siddharth Anand	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-54
AIRFLOW-54	[~amustafa] Any PRs that we can look at?	Siddharth Anand	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-54
AIRFLOW-52	[~criccomini] Can we do an 'airbnb' release for now then? ie. 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)	Bolke de Bruin	0	0	1	AFAIK, entrance into the incubator does not require that you halt any release from the existing airbnb GitHub repo, using whatever process you've always used there. It's definitely good to show progress towards a first Apache release, but that's often easier to do decoupled from any time pressures about trying to ship a particular feature or bug fix. That way, you can focus on things like the licensing requirements, which are tricky for the first release in their own right.	Chris Nauroth	https://issues.apache.org/jira/browse/AIRFLOW-52
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.	Chris Riccomini	0	0	1	AFAIK, entrance into the incubator does not require that you halt any release from the existing airbnb GitHub repo, using whatever process you've always used there. It's definitely good to show progress towards a first Apache release, but that's often easier to do decoupled from any time pressures about trying to ship a particular feature or bug fix. That way, you can focus on things like the licensing requirements, which are tricky for the first release in their own right.	Chris Nauroth	https://issues.apache.org/jira/browse/AIRFLOW-52
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 pinging general@incubator.	Jakob Homan	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-52
AIRFLOW-52	[~aen], are you planning to burn in [~jlowin]'s fix before release?	Chris Riccomini	0	0	1	This one right https://github.com/apache/incubator-airflow/pull/1473 ? In this case yes.	Dan Davydov	https://issues.apache.org/jira/browse/AIRFLOW-52
AIRFLOW-52	This one right https://github.com/apache/incubator-airflow/pull/1473 ? In this case yes.	Dan Davydov	0	0	1	Dan Davydov I think Chris Riccomini refers to: https://github.com/apache/incubator-airflow/pull/1378 (as is also linked in AIRFLOW-56)	Bolke de Bruin	https://issues.apache.org/jira/browse/AIRFLOW-52
AIRFLOW-52	Should this be closed?	Jeremiah Lowin	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-52
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.	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-51

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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 *postgresql://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.	Siddharth Anand	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-48
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	Chris Riccomini	0	0	1	Chris Riccomini I agree the pendulum is currently at one extreme and I would like to not see it swing the other way but instead settle somewhere in the middle until a better solution can be reached. Currently, the entire Airflow app is located at /admin/. As such, all DAGs and DAG specific UI functionality, which can be considered as user-specific, is shared in the same UI space as admin-specific functionality (e.g. Variables, Connections, Configuration). In other words, there is no separation between Admin and User functionality. As an admin, I would like to see both connection and variable information. Currently, at Airbnb and Agari and a few other places, a reverse proxy sits in front of Airflow to only allow authenticated access to Airflow. Hence, only intended users can access it - e.g. engineers in the data team. Granted, they have access to everything - they are essentially admins. We do support encryption of Variable and Connection's Extra fields in the DB, so that people with access to the DB but who do not have access to the admin UI cannot access them. This is just a statement of the current state, not a statement of the ideal state. We should ideally break the app into separate admin and user functional areas. As you point out, the variables are still accessible within the DAGs. We have a requirement to make the variable value visible without clicking on the edit button. This goes against your requirement. I think we should have a configuration variable to hide/expose the encrypted values in the UI for both Variables	Siddharth Anand	https://issues.apache.org/jira/browse/AIRFLOW-45
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?	Chris Riccomini	0	0	1	"Note: this might lead to some confusion. Defaulting to 'True' means that when I add extras to a hook, or define any variable, by default, if I hit 'edit' afterwards, I'll see an empty box for the field." & "I think your proposal is sound. To address the confusion, there can be a message on the UI that says values are not show due to the 'True' config setting."	Chris Riccomini & Siddharth Anand	https://issues.apache.org/jira/browse/AIRFLOW-45
AIRFLOW-45	Sounds good. [~cheny258], would you like to try this?	Chris Riccomini	0	0	1	Chris Riccomini Yea will do. Creating a PR now.	Matthew Chen	https://issues.apache.org/jira/browse/AIRFLOW-45

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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.	Siddharth Anand	0	0	1	Siddharth Anand will do that for the list view of variables.	Matthew Chen	https://issues.apache.org/jira/browse/AIRFLOW-45
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?	Matthew Chen	0	1	0	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?		https://issues.apache.org/jira/browse/AIRFLOW-45
AIRFLOW-45	Could we not handle this with a regexp of fields that should be hidden rather than a boolean?	Paul Rhodes	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-45
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?	Chris Riccomini	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-45
AIRFLOW-39	Can you provide sample code?	Siddharth Anand	0	0	1	<pre> from airflow import DAG from airflow.operators import BashOperator from datetime import datetime, timedelta default_args = { 'owner': 'airflow', 'depends_on_past': False, 'start_date': datetime(2016, 5, 4, 16, 40), 'end_date': datetime(2016, 5, 4, 16, 50), 'email': ['airflow@airflow.com'], 'email_on_failure': False, 'email_on_retry': False, 'retries': 1, 'retry_delay': timedelta(minutes=5), # 'queue': 'bash_queue', # 'pool': 'backfill', # 'priority_weight': 10, # 'end_date': datetime(2016, 1, 1), } dag = DAG('tutorial4', default_args=default_args, schedule_interval=timedelta(minutes=5)) # t1, t2 and t3 are examples of tasks created by instantiating operators t1 = BashOperator(task_id='print_date'</pre>	dud	https://issues.apache.org/jira/browse/AIRFLOW-39
AIRFLOW-39	Oh, so not version 1.7.0? What does 'git log head' look like for you?	Siddharth Anand	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-39
AIRFLOW-35	[~sanand], any particular reason you closed this?	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-35

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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 <code>{{import}}</code> syntax etc). Maybe some of what I did for plugins would be reusable?	Jeremiah Lowin	0	0	1	"This is causing problems. See AIRFLOW-200 for more examples. Based on some reading, it seems like we might be able to clean it up, at least, using <code>sys.metapath</code> . I'm not familiar with this part of Python, so I could be misguided. What I'm thinking is just implementing a series of <code>sys.meta_path</code> classes that sit in the <code>_init_.py</code> files, and look when imports are called (lazily). When they're called it tries to import, using similar logic to what's in <code>import_module_attrs</code> , and pukes (with a real error this time) if it can't load."&"Please have a look at AIRFLOW-200, and this PR. While it does not use standard imports, it does improve the experience for the current import style, so that people can see why hooks/operators can't import, and makes loads lazy. "	Chris Riccomini	https://issues.apache.org/jira/browse/AIRFLOW-31
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.	Chris Riccomini	0	0	1	Chris Riccomini 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 With my suggestion the transaction would not span multiple operators — <code>GenericTransfer</code> has a <code>preoperator</code> argument (see the docs) which just holds a query that's executed before the insertions. I was suggesting that we allow binding these together, with all of them being part of the same task instance. In your use cases, are the source and destination DBs part of the same DB cluster? Nope, different clusters.	Bence Nagy	https://issues.apache.org/jira/browse/AIRFLOW-30
AIRFLOW-30	In your use cases, are the source and destination DBs part of the same DB cluster?	Chris Riccomini	0	0	1	Chris Riccomini 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 With my suggestion the transaction would not span multiple operators — <code>GenericTransfer</code> has a <code>preoperator</code> argument (see the docs) which just holds a query that's executed before the insertions. I was suggesting that we allow binding these together, with all of them being part of the same task instance. In your use cases, are the source and destination DBs part of the same DB cluster? Nope, different clusters.	Bence Nagy	https://issues.apache.org/jira/browse/AIRFLOW-30
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).	Chris Riccomini	0	0	1	Chris Riccomini: 1. I would consider that unlikely — Python as far as I know has continuous garbage collection and no batch GC runs that would cause a huge slowdown for only some of the imports. 2. Even if there would be some rare periodic overloading, it would affect only one scheduler run which is not a huge deal. 3. If someone is frequently causing the Python process to be really, really slow, they can just change this config variable. But at that point they might as well hit the 30s limit too — I would assume the median import time (of correctly implemented DAGs) is in the 0.001s - 0.01s order of magnitude. Considering this there's not that huge of a difference between limits of 1s and 30s.	Bence Nagy	https://issues.apache.org/jira/browse/AIRFLOW-29

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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?	Chris Riccomini	0	0	1	"We do a fair amount of "dynamic DAG generation" and have many large DAGs at Airbnb, some generated by parsing yaml which can be surprisingly slow. 1 second is a relatively short time, 5 seconds may be a more sensible default. When we do alter defaults, we need to be upfront about it in the release notes, we should introduce a new section at the top that says "Changes the may affect you""&"We're using dynamically generated DAGs as well; the most complex one is defining 5 DAGs with 4 tasks each, and I already feel like this makes me enough of a power user that it would be reasonable to require me to raise this setting from the default in the config manually. But the import time doesn't come close at all actually, my benchmarks show it at ~15ms (if the airflow package is cached already, but it surely is when the scheduler is importing DAGs). I'd argue that a decrease to 5s wouldn't really accomplish what my goal was here — to prevent people from wasting resources by obliviously writing operators like the S3FileTransformOperator in 1.7.0. I stand by my stance that 1s is a setting that would help a lot more people than it would (very slightly) inconvenience."	Maxime Beauchemin & Bence Nagy	https://issues.apache.org/jira/browse/AIRFLOW-29
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?	Siddharth Anand	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-29
AIRFLOW-28	Okay, but what's preventing someone else from implementing this? How come the issue was closed as won't fix?	Bence Nagy	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-28
AIRFLOW-23	Hey [~alexvanboxel], checking in to see how this and AIRFLOW-24 are coming.	Chris Riccomini	0	0	1	Would like to tests all of the components, but I don't use all of them: https://github.com/apache/incubator-airflow/pull/1532	Alex Van Boxel	https://issues.apache.org/jira/browse/AIRFLOW-23
AIRFLOW-22	Can public buckets/BQ datasets (thinking about GDELT) be used without a project?	Jeremiah Lowin	0	0	1	What's GDELT? 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.	Chris Riccomini	https://issues.apache.org/jira/browse/AIRFLOW-22
AIRFLOW-22	What's GDELT? 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.	Chris Riccomini	0	0	1	A very cool global dataset that is freely available through BQ: http://gdeltproject.org/data.html	Jeremiah Lowin	https://issues.apache.org/jira/browse/AIRFLOW-22
AIRFLOW-21	Do you need some help? I could do it.	Alex Van Boxel	0	0	1	[~alexvanboxel], that would be awesome. Can you wait until AIRFLOW-16 goes in, though, to avoid merge conflicts? Should be today.	Chris Riccomini	https://issues.apache.org/jira/browse/AIRFLOW-21

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
AIRFLOW-21	[~alexvanboxel], that would be awesome. Can you wait until AIRFLOW-16 goes in, though, to avoid merge conflicts? Should be today.	Chris Riccomini	0	0	1	I tried to support the old P12 as well but I think it's broken on GCP itself. If I request cloud scope it says it's not a valid scope. Removing it says "Request had insufficient authentication scopes". I'm going to give an explicit error when using a p12 key.	Alex Van Boxel	https://issues.apache.org/jira/browse/AIRFLOW-21
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	Bolke de Bruin	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-20
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.	Bolke de Bruin	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-20

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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.	Bolke de Bruin	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-20
AIRFLOW-17	Got it. Is there any way to fix it?	Chris Riccomini	0	0	1	"Unfortunately, I have very little knowledge of how the license check actually works. The code here seems very simplistic : https://github.com/airbnb/airflow/blob/master/scripts/ci/check-license.sh#L98 Maybe we can disable this check in the case of a revert." & "It is very simplistic . The most simple way to fix this to do a PR that adds a couple of license headers to some of the files that still require it. I'm not sure (ie. didn't check) if travis knows about reverts. As we need to add more licenses anyway the above is maybe best. I was actually thinking of proposing to do an extra -1 of every PR to make sure we get compliant sooner without one giant PR that requires everyone to sync, although that might be unavoidable anyway ."	Arthur Wiedmer & Bolke de Bruin	https://issues.apache.org/jira/browse/AIRFLOW-17
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.	Chris Riccomini	0	0	1	"I don't have an easy way to test the Datastore or BigQuery hooks, but I can confirm that I was able to very easily create a 'google_cloud_platform' connection in the UI and use it to write/read remote logs. The DS/BQ code didn't seem to raise any errors, for what it's worth."&"Got it working with the new hooks. I have a few suggestions though, have all GCP operations and hooks be prefixed with gcp. I'm going to add at least 2 services (dataflow and dataproc) maybe a few others. So example: gcp_dataflow gcp_dataproc gcp_storage gcp_datastore gcp_bigquery It a good moment to do, because we are breaking things anyway. Could be a ticket, I'll pick it up (after this is merged)."	Jeremiah Lowin & Alex Van Boxel	https://issues.apache.org/jira/browse/AIRFLOW-16
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.	Chris Riccomini	0	0	1	No bug, had to adapt my POC. Anyway +1	Alex Van Boxel	https://issues.apache.org/jira/browse/AIRFLOW-16
AIRFLOW-15	Pull request: https://github.com/airbnb/airflow/pull/1448 ["~jlowin], could use a review.	Chris Riccomini	1	0	0			https://issues.apache.org/jira/browse/AIRFLOW-15
AIRFLOW-15	["~criccomini] is this issue resolved?	Bolke de Bruin	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-15

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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.	Chris Riccomini	0	0	1	Excellent question. I've tried to address this with an addition to the existing kill_zombies() function. Basically, multiple DRJ's can be trying to run multiple DagRuns – that's not an issue. They just take a lock when they're actually executing so the other DRJ's know to skip that one. (this is a slight repeat to something I put in the email, but better to have it here than there). The lock could have been a simple True/False flag, but that could lead to exactly the situation you're describing – the DRJ locks the DR, dies, and then no one ever touches that DR again because of the lock. So the new kill_zombies() method looks at the lock_id and looks for an active job with that id. If the job exists, it takes no action. If the job is gone/ended, it unlocks the DagRun which makes it available for other DRJs (like a new Scheduler). See: https://github.com/jlowin/airflow/blob/dagrun-refactor/airflow/models.py#L337	Jeremiah Lowin	https://issues.apache.org/jira/browse/AIRFLOW-14
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"]	Chris Riccomini	0	0	1	That is exactly correct. Multiple DRJ are not just supported but encouraged.	Jeremiah Lowin	https://issues.apache.org/jira/browse/AIRFLOW-14
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?	Chris Riccomini	0	0	1	DagRuns are primary keyed by (dag_id, execution_date), so there is only one canonical version. DRJ needs to make sure to refresh from the db to check for a lock immediately before running it, however. The mechanism is very similar to TaskInstance – you can create as many TI objects as you want, but they all point at the one canonical version and can be refreshed at any time to reflect the "true" state. DRJs pick up DagRuns in two ways: 1. explicitly via DagRunJob.submit_dags(). This is used for example by BackfillJob; it generates a bunch of DagRuns and calls submit_dags() to submit them to itself. Then it enters its loop. Scheduler also uses this after scheduling a DagRun, though it's actually redundant because of the second way (below) 2. automatically via DagRunJob.collect_dags(). This is called inside each DRJ loop and is used by SchedulerJob to look for any active DagRuns and add them to its set of DagRuns to execute.	Jeremiah Lowin	https://issues.apache.org/jira/browse/AIRFLOW-14
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?	Chris Riccomini	0	0	1	Yes, that situation could arise... I'm not sure of the best way to handle this in a database-agnostic way (i.e. SQLAlchemy). I'd appreciate any suggestions!	Jeremiah Lowin	https://issues.apache.org/jira/browse/AIRFLOW-14

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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?	Jeremiah Lowin	0	0	1	there are two DRJs running simultaneously, this race condition can occur, as far as I know: DRJ0 and DRJ1 both see DAG0 needs to be rnu DRJ0 checks if the lock is set. It's not. DRJ1 checks if the lock is set. It's not. DRJ0 sets lock to be owned by DRJ0. DRJ0 checks to see who the owner is. It's DRJ0. DRJ0 starts running the DAG. DRJ1 sets the lock to be owned by DRJ1. DRJ1 checks to see who the owner is. It's DRJ1. DRJ1 starts running the DAG. The SQL that I wrote above prevents this from happening because the second update (6) doesn't happen because at that point, the lock_id isn't null, it's DRJ0.	Chris Riccomini	https://issues.apache.org/jira/browse/AIRFLOW-14
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?	Paul Yang	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-14
AIRFLOW-14	While I understand that for being truly distributed it would be great to have a distributed scheduler, are not looking for a distributed executor instead ["riccomini"]? Now the scheduler and the executor are coupled, but that does not need to be and does not even make sense imho.	Bolke de Bruin	0	0	1	{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.	Chris Riccomini	https://issues.apache.org/jira/browse/AIRFLOW-14
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.	Chris Riccomini	0	0	1	Now the scheduler is, in case of using the LocalExecutor, tied to the executor: ie. both are same process. This is nor 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 anuwhere. ie. producer/consumer like. Does this make sense?Bolke de Bruin	Bolke de Bruin	https://issues.apache.org/jira/browse/AIRFLOW-14
AIRFLOW-14	Now the scheduler is, in case of using the LocalExecutor, tied to the executor: ie. both are same process. This is nor 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 anuwhere. ie. producer/consumer like. Does this make sense?Bolke de Bruin	Bolke de Bruin	0	0	1	"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?" & "Here's how I'd characterize the current system: Executors take in task instructions and then make sure the tasks are run – either in the executor's own process (SequentialExecutor, in child processes (LocalExecutor) or in distributed processes (CeleryExecutor)." Scheduler figures out what tasks need to be run and passes those instructions to the executor "master", which is running in its own process. So I think the piece you guys are discussing centers on the hand off from the Scheduler to the executor "master". Today that's in-process, for convenience but I see no reason that couldn't be done via DB (with care). However this is a larger issue and just to be clear, not part of this proposal.	Chris Riccomini & Jeremiah Lowin	https://issues.apache.org/jira/browse/AIRFLOW-14

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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?	Chris Riccomini	0	1	0			https://issues.apache.org/jira/browse/AIRFLOW-14

ANY23

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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/lewismc/any23 ANY23-297 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 ----	ASF GitHub Bot	0	0	1	GitHub user ansell commented on the issue: https://github.com/apache/any23/pull/30 I will run a build to verify that it works and get back to you.	ASF GitHub Bot	https://issues.apache.org/jira/browse/ANY23-297
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	ASF GitHub Bot	0	0	1	GitHub user ansell commented on the issue: https://github.com/apache/any23/pull/30 I will run a build to verify that it works and get back to you.	ASF GitHub Bot	https://issues.apache.org/jira/browse/ANY23-297
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?	Lewis John McGibbney	0	0	1	I attach a modified test file, adapted from the original, where I moved one statement in /html/body/div/script. The corresponding test EmbeddedJSONLDEXtractorTest.java then can still test that "assertStatementsSize(null, null, 3);" to test if the piece inserted in the body was correctly parsed. I am not familiar at all with how to fix things or create a patch...	Thomas Francart	https://issues.apache.org/jira/browse/ANY23-291

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
ANY23-280	Github user lewisinc 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.	ASF GitHub Bot	0	0	1	Github user ansell commented on a diff in the pull request: https://github.com/apache/any23/pull/24#discussion_r67792344 — Diff: plugins/office-scraper/src/test/java/org/apache/any23/plugin/officescraper/ExcelExtractorTest.java — @@ -40,6 +42,8 @@ import org.slf4j.Logger; import org.slf4j.LoggerFactory; +import com.fasterxml.jackson.databind.introspect.WithMember; — End diff — This import seems to be unused" & "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." & "Github user ansell commented on the issue: https://github.com/apache/any23/pull/24 If we are going to be modifying the public API we probably should be aiming for a 2.0 release, otherwise the version numbers are arbitrary" & "Github user ansell commented on the issue: https://github.com/apache/any23/pull/24 Given how broad this pull request is, it needs to be completed before I can work on some of the issues I have assigned to	ASF GitHub Bot	https://issues.apache.org/jira/browse/ANY23-280
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?	ASF GitHub Bot	0	0	1	Github user lewisinc commented on the issue: https://github.com/apache/any23/pull/27 Hi @ansell ... [INFO] ----- [ERROR] Failed to execute goal org.apache.maven.plugins:maven-javadoc-plugin:2.8:jar (attach-javadocs) on project apache-any23-core: MavenReportException: Error while creating archive: [ERROR] Exit code: 1 - /usr/local/any23/core/src/main/java/org/apache/any23/extra ctor/rdfa/RDFA11Parser.java:324: error: @param name not found [ERROR] * @param curieOrIRIList list of CURIE/URI. [ERROR] ^ [ERROR] /usr/local/any23/core/src/main/java/org/apache/any23/extra ctor/rdfa/RDFA11Parser.java:328: warning: no @param for curieOrIRIList [ERROR] protected IRI[] resolveCIRIleOrIRIList(Node n, String curieOrIRIList, boolean termAllowed) [ERROR] ^ [ERROR] /usr/local/any23/core/src/main/java/org/apache/any23/extra ctor/rdfa/RDFA11Parser.java:328: warning: no @param for termAllowed [ERROR] protected IRI[] resolveCIRIleOrIRIList(Node n, String curieOrIRIList, boolean termAllowed)	ASF GitHub Bot	https://issues.apache.org/jira/browse/ANY23-276
ANY23-276	Github user lewisinc 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.	ASF GitHub Bot	0	0	1	Implemented in version2	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-276

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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.	ASF GitHub Bot	0	0	1	Github user ansell commented on the issue: https://github.com/apache/any23/pull/27 I will work on the appassembler still being present for core and check the cli modules assembly-plugin and then merge to master. Can you do a silent retarget (no notifications sent to mailing list) of all Jira issues to 2.0 to reflect the version bump?	ASF GitHub Bot	https://issues.apache.org/jira/browse/ANY23-276
ANY23-259	[~p_ansell] is this the issue which is leading to the unstable builds?	Lewis John McGibbney	0	0	1	Sorry, this issue was filed as a bug when it should have been an improvement. Have fixed that now. The main issue still seems to be the meta tags not being closed which is ANY23-267 There also appear to be a changed (lower, unless the junit assertEquals was applied backwards) number of extractors than expected.	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-259
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 :(	Lewis John McGibbney	0	0	1	The last two out of the 5 test failures (build 1464) on Jenkins are likely not related to meta elements. https://builds.apache.org/job/Any23-trunk/1464/testReport/	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-259
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/owls/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?	Peter Ansell	0	0	1	As OWL isn't an input format, per se, which particular format were you wanting to target? Correct, right now however I was wanting to implement support for those particular ontologies. We use them extensively within the earth science community. so it is a little surprising that it isn't being picked up right now. ACK. I am going to try debugging some of this on test file and see where the hang up is.	Lewis John McGibbney	https://issues.apache.org/jira/browse/ANY23-257
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.	ASF GitHub Bot	0	0	1	Github user asfgit closed the pull request at: https://github.com/apache/any23/pull/13	ASF GitHub Bot	https://issues.apache.org/jira/browse/ANY23-250

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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	ASF GitHub Bot	0	0	1	GitHub user lewisinc commented on a diff in the pull request: https://github.com/apache/any23/pull/17#discussion_r27442017 --- 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.	ASF GitHub Bot	https://issues.apache.org/jira/browse/ANY23-247
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?	Lewis John McGibbney	0	0	1	I think the only place they are defined right now is in DefaultValidator.loadDefaultRules, and the only place they are applied is in DefaultValidator.validate. You may need to create an instance of Rule to match documents that have 'itemscope' and then use the Fix implementation that you have written already to patch them with 'itemscope="itemscope"'. You pair the Rule with the Fix in DefaultValidator.loadDefaultRules Ideally we would have a FixFactory interface that is implemented for each combination of a Rule with an optional Fix. The FixFactory can then be registered as a service using META-INF/services, to avoid having them hardcoded into DefaultValidator.loadDefaultRules.	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-247

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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 [RDFa11Extractor](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).	ASF GitHub Bot	0	0	1	Github user ansell commented on the pull request: https://github.com/apache/any23/pull/17#issuecomment-201545776 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.	ASF GitHub Bot	https://issues.apache.org/jira/browse/ANY23-247
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	ASF GitHub Bot	0	0	1	Github user ansell commented on the pull request: https://github.com/apache/any23/pull/17#issuecomment-201545776 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.	ASF GitHub Bot	https://issues.apache.org/jira/browse/ANY23-247
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.	ASF GitHub Bot	0	0	1	Github user asfgit closed the pull request at: https://github.com/apache/any23/pull/17	ASF GitHub Bot	https://issues.apache.org/jira/browse/ANY23-247

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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 for 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 output 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	ASF GitHub Bot	0	0	1	Hi Lewis. The BNode ids will be of the same format. The problem was that the MD5 of the string "0" was being used in some cases where the source contained 'itemid=""'. For example, in one of our extractions this lead to over 140,000 type relations to the BNode _:nodecfc208495d565ef66e7dff9f98764da as 'cfc208495d565ef66e7dff9f98764da' is the MD5 of "0". I'm not objecting to the use of an MD5 hash as the BNode id as long as it has an extremely low probability of collisions. In Any23 the MD5 is often generated directly on the Java hashcode, which when extracting billions of tuples can lead to collisions. Especially if there is a problem with the hashcode implementation.	Timothy Potter	https://issues.apache.org/jira/browse/ANY23-238
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.	Lewis John McGibbney	0	0	1	JSONLD-Java doesn't currently parse HTML, whereas Any23 does have HTML-based formats that it supports. Hence, it may be more appropriate to pick out the <script type="application/ld+json">...</script> elements in Any23 and pass their content to JSONLD-Java for parsing. If it was going to be added to jsonld-java, then it would probably end up being in an integration module to pull in the otherwise unused HTML parsing libraries without adding those dependencies to the core module. Mostly the reasoning for that is that it is a non-normative section in the spec so it is an optional and not well defined feature. One immediate question I have is what would happen if there are two or more JSON-LD elements in a single HTML document. I would think that they would need to be processed to RDF to be combined, which would lose the context information, but should preserve the other information.	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-226
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*	ASF GitHub Bot	0	0	1	Github user ansell commented on the pull request: https://github.com/apache/any23/pull/16#issuecomment-84257478 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 "<div itemscope>"	ASF GitHub Bot	https://issues.apache.org/jira/browse/ANY23-226

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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.	Michele Mostarda	0	0	1	Hey Michele Mostarda, I am 100% for the following 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? I totally understand the basis of this issue and will actually log in just now to the analytics account and see some results. It would be nice if we could report these out to dev@ I'll see if I can address that. Are you able to address the suggestions you've put forward? Thank you Michele.	Lewis John McGibbney	https://issues.apache.org/jira/browse/ANY23-202
ANY23-185	hi [~nisala12], 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."?	Lewis John McGibbney	0	0	1	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-equiv 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	Nisala Mendis	https://issues.apache.org/jira/browse/ANY23-185
ANY23-177	Hey [~lanthaler], do you have any interest in working on this issue? If so I can assign it to you.	Lewis John McGibbney	0	0	1	Lewis John McGibbney, 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. Peter 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!	Markus Lanthaler	https://issues.apache.org/jira/browse/ANY23-177
ANY23-155	Hi Debmalya, can you submit a patch we can apply against trunk? Thanks for opening this issue.	Lewis John McGibbney	0	0	1	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 >	Debmalya Jash	https://issues.apache.org/jira/browse/ANY23-155

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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	Debmalya Jash	0	0	1	Hi Debmalya, you can grab the Any23 source here https://github.com/apache/any23 , once you've made changes to the codebase you can create a patch as follows git diff --no-prefix trunk > ANY23-\$JiraNumber.patch The time line is whenever you can Any Any23 contributions are very welcome. Thank you	Lewis John McGibbney	https://issues.apache.org/jira/browse/ANY23-155
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	Peter Ansell	0	0	1	java version "1.7.0_10" Java(TM) SE Runtime Environment (build 1.7.0_10-b18) Java HotSpot(TM) 64-Bit Server VM (build 23.6-b04, mixed mode) Apache Maven 3.0.4 (r1232337; 2012-01-17 16:44:56+0800) Maven home: C:\apache-maven-3.0.4-bin\apache-maven-3.0.4\bin\.. Java version: 1.7.0_10, vendor: Oracle Corporation Java home: C:\Program Files\Java\jdk1.7.0_10\jre Default locale: en_SG, platform encoding: Cp1252 OS name: "windows 7", version: "6.1", arch: "amd64", family: "windows"	Debmalya Jash	https://issues.apache.org/jira/browse/ANY23-155
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	Peter Ansell	0	0	1	Please find the attachments.	Debmalya Jash	https://issues.apache.org/jira/browse/ANY23-155
ANY23-155	So RDFSchemaUtilsTest is working now for you after you reverted your changes to it?	Peter Ansell	0	0	1	I reverted the changes. It is not working now. Failed tests: testGetClassesInPackage(org.apache.any23.util.DiscoveryUtilsTest)) testGetClassesInJAR(org.apache.any23.util.DiscoveryUtilsTest) testSerializeVocabulariesNTriples(org.apache.any23.vocab.RDFSchemaUtilsTest): expected:<1622> but was:<0> testSerializeVocabulariesRDFXML(org.apache.any23.vocab.RDFSchemaUtilsTest): expected:<4241> but was:<0> Tests in error: testBlankNodesViaURL(org.apache.any23.Any23Test): www.usarab.org Tests run: 391, Failures: 4, Errors: 1, Skipped: 4 Regarding testBlankNodesViaURL, www.usarab.org is not available now.	Debmalya Jash	https://issues.apache.org/jira/browse/ANY23-155
ANY23-147	http://purl.org/dc/terms/ doesn't look to be deprecated. Are you thinking of http://purl.org/dc/elements/1.1/ ?	Peter Ansell	0	0	1	Well its a redirect right? Which URI so we use as the NS?	Lewis John McGibbney	https://issues.apache.org/jira/browse/ANY23-147

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
ANY23-147	Well its a redirect right? Which URI so we use as the NS?	Lewis John McGibbney	0	0	1	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?	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-147
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?	Peter Ansell	0	0	1	{bq}Is that part of the confusion? {bq} Yes, entirely. Can you also suggest a suitable Class namign convention for the vocabulary as I am going to submit and commit the naming for ANY23-144	Lewis John McGibbney	https://issues.apache.org/jira/browse/ANY23-147
ANY23-144	What do you suggest then Peter? I am happy to go with a positive consensus on this one.	Lewis John McGibbney	0	0	1	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. That is mostly based on positive experiences with the Sesame vocabulary classes, where you can use RDF.TYPE etc. The any23 vocabulary design requires access to entire vocabularies, which isn't a design constraint currently for the Sesame vocabulary classes, so there may end up being some difference still. The current Any23 pattern has access to vocabulary properties using <code>DOAP.getInstance().PROPERTY</code> , which tends to result in <code>DOAP.getInstance()</code> being stored in a private static final variable inside of each of the relevant extractors. The static final variable in the current extractors tends not to be named in upper-case to follow the java convention for static final variables. Instead they are named in hungarian case, with a "v" prefix to indicate they are vocabularies. To clear up the naming I would prefer if we avoided the <code>getInstance()</code> pattern. We might be able to convert the vocabularies into an extensible enumeration pattern with direct access to fields along with an interface to access the entire vocabulary using a base implementation. Then we could legitimately get access to the vocabulary properties programmatically as we currently can, and we could still have <code>RDF.TYPE</code> like access to the properties without storing them in a local private field to get around the <code>getInstance()</code> call.	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-144

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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?	Lewis John McGibbney	0	0	1	DOAP and LKIFCoreAction match my suggestion. I was rambling about the entire vocabulary design pattern, as I thought it would be useful to have a discussion about that also, however, you are right that the getInstance()+private static final fields, etc, are a separate issue.	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-144
ANY23-144	patch for trunk which addresses comments and concerns. If someone can review we can sort this one out. Thank you	Lewis John McGibbney	0	0	1	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 the apply the patches?	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-144
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?	Peter Ansell	0	0	1	I* think we should release 0.8.0 now.	Lewis John McGibbney	https://issues.apache.org/jira/browse/ANY23-141
ANY23-141	Sesame-2.7.0 is now available, and the Git repository is now writable. What should we do about this issue?	Peter Ansell	0	0	1	Hi Peter. I am really looking forward to getting any23 back on track. I think you should progress with the upgrade as you see fit. There is no reason to delay now as the goal posts shifted somewhat over the last month or so. I'll work to get the remaining tests ignored or passing tomorrow and we can push an RC ASAP. Thanks for your persistence Peter it is really important for us all.	Lewis John McGibbney	https://issues.apache.org/jira/browse/ANY23-141
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.	Lewis John McGibbney	0	0	1	Hi Lewis. It should fix all related issues and failed test cases at [1]. Semargl v0.4 will be released in 1-2 weeks and will be available at maven central, so I can update proposal for further review. [1] http://rdfa.info/test-suite	Lev Khomich	https://issues.apache.org/jira/browse/ANY23-137
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?	Lewis John McGibbney	0	0	1	I am waiting on Lev to either deploy Semargl-0.5 to Maven Central, or release and deploy a 0.6 release. The 0.4 release does not integrate easily into Sesame-2.7.5 as mentioned above. This was fixed in 0.5 so either solution is fine with me as far as I remember. https://github.com/levkhomich/semargl/issues/29	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-137
ANY23-137	What's the status of this issue to merge semargl into Any23? Is there anything I can test to make progress?	stephane corlosquet	0	0	1	Hi stephane corlosquet, AFAIK this issue is not far from being suitable for integration in to the codebase. There is a ANY23-137 branch which you can checkout and work on. As Peter Ansell noted above, there are some test failures which need to be addressed before we can merge in to trunk. If you are able to work on some of these it would be excellent.	Lewis John McGibbney	https://issues.apache.org/jira/browse/ANY23-137

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
ANY23-121	Has this got anything to do with Any23?	Lewis John McGibbney	0	0	1	Hi Lewis, Apologies for this, I am new to this site and dont able to find exact group to post this issue. Appreciate your help if you could move this issue to relevant group. Thanks	Koushik Biswas	https://issues.apache.org/jira/browse/ANY23-121
ANY23-110	I honestly didn't realise that this was committed. Have you verified it works Peter? I will try tomorrow	Lewis John McGibbney	0	0	1	I just committed it when I made the comment. There is very little to verify about a Vocabulary, other than fixing the numbers in RDFSSchemaUtilsTest. It compiled and did not break any tests other than the RDFSSchemaUtilsTest, which I fixed. If there are predicates/classes etc., that are still missing feel free to add them. I committed it to the api module, which was a slight change to the patch that was to the core module, as the vocabularies are all exported from the api module now to avoid users relying on the core module where possible.	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-110
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?	Peter Ansell	0	0	1	No Peter. I was not sure whether TDB should be supported just in the webservice module or it should be supported within a more core part of the code. At this stage any suggestion is better than nothing.	Lewis John McGibbney	https://issues.apache.org/jira/browse/ANY23-105
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?	Lewis John McGibbney	0	0	1	TripleHandler's could be pluggable with a little effort to add TripleHandlerFactory implementations for each of the existing TripleHandler implementations. Given that the WebService can have its TripleHandler setup manually though, it would not require them to be pluggable using SPI to use them in web service. TDB may be supported in a similar way to the current RepositoryWriter that exports statements to an OpenRDF RepositoryConnection. It would just require a simple translation. Or, on another level, it may even be possible to directly use RepositoryWriter with the jena-sesame code that Andy and I have been collaborating on: https://github.com/ansell/JenaSesame	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-105
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.	Lewis John McGibbney	0	0	1	Yes, once all of the test resources were moved to a separate jar file, this bug was fixed.	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-103

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
ANY23-83	Having had a browse of both o.a.a.vocab.RDFSchemUtils, 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	Lewis John McGibbney	0	0	1	I was hoping that Any23 would be extensible to more than the current list of supported parsers. I understand that when Any23 was originally developed that it would have needed to restrict itself to a small set of supported parsers. I am not familiar enough with the rest of the codebase to know why the shortlist of supported parsers is in place, but it would be good if it could be made generic in the future. In terms of a generic factory type implementation, one could use the SPI based <code>Rio.getParser(RDFFormat)</code> or any of the other static methods in the <code>Rio</code> class. <code>RDFFormat</code> is not a fixed enum, so it can be extended, and the parsers are instantiated using <code>RDFParserFactory</code> implementations instead of directly created [1]. The mime-types are described inside of <code>RDFFormat</code> as a list, with preferred and alternate mime types, and the <code>RDFParserFactory</code> implementations that are found using SPI are each tested to see if the <code>RDFFormat</code> that they support fits with a given mime-type before the parser is instantiated. In terms of <code>RDFSchemUtils</code>, it is not clear why the list is restricted the way it is, as it would be very useful to be able to patch any <code>RDFHandler</code> in, and would not take much coding to make the change as far as I can tell. [1] http://www.openrdf.org/doc/sesame2/api/org/openrdf/rio/RDFParserFactory.html	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-83
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?	Lewis John McGibbney	0	0	1	My personal preference (as you may have noticed at ANY23-19) is to stay with the OpenRDF API, even if you don't use the Sesame implementation of it. All of the elements of the OpenRDF API are specified as interfaces and there are no hidden static singletons in the background that you can't reimplement another way when you need to. In addition, you can pull in the API using Maven without having to use the implementations which I view as a large point in its favour. I made quite a few changes yesterday to basically suit my use cases. These including fixing the cause for this issue by switching to <code>RDFFormat</code> as the central point of reference for the utils classes. I pushed the changes to GitHub at https://github.com/ansell/any23 . I am not sure how many of the changes you will want to pick up but you may find something interesting. The most visible change is that I split the core module up into a large number of smaller modules. You probably won't want to do that just yet, but I found that I could work with it and visualise it more easily if everything wasn't in a single core module. It may seem a little convoluted, particularly having the test resources in their own module, separate from the core unit tests, but it makes it simpler to use the test resources in maven to unit test the other modules, without having to either split up the test resource bundle or replicate some resources in other modules. The other very visible change is that I deleted the Any23 NQuads library in favour of importing the MIT licensed Sesame toolkit NQuads library.	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-83

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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	Lewis John McGibbney	0	0	1	"I read the patch, it sounds really clear, it is an optimization to improve performances while processing detecting nested node paths. I will integrate it in a couple of days." & "Thanks Michele. I'll explain the change anyway: The current implementation is slow on pages with lots of items because it does nested iterations over all itemscope and itemprops under the given scopeNode. In the inner loop it builds XPath strings for each node to test if the nodes are related. This test condition is inherent in the tree structure of the DOM. The patch changes the code to do a traversal of the DOM tree limited to nodes only in the given scopeNode. Reading the code itself is probably the best way to understanding the change."	Michele Mostarda & Timothy Potter	https://issues.apache.org/jira/browse/ANY23-75
ANY23-69	Does this duplicate ANY23-7 ?	Lewis John McGibbney	0	0	1	ANY23-7 is about general extractor performances while this issue is about specification conformance (how much the implementation is compliant with the specification).	Michele Mostarda	https://issues.apache.org/jira/browse/ANY23-69
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?	Michele Mostarda	0	0	1	Hola! META-INF/services metadata are always generated at compile time, indeed plugins are annotated with @MetalnfService 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 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, WDYT? Anyway, please double-ack if you agree I can commit the modification! TIA!	Simone Tripodi	https://issues.apache.org/jira/browse/ANY23-50
ANY23-50	Hola! META-INF/services metadata are always generated at compile time, indeed plugins are annotated with [MetalnfService 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 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, WDYT? Anyway, please double-ack if you agree I can commit the modification! TIA!	Simone Tripodi	0	0	1	Hi Simo, sorry for delay. I had a quick check to org.kohsuke.MetalnfServices . It seems to fit all our needs: use the quickest and more standard way to declare plugins in a transparent fashion. I don't have any objection with this patch, for me it is ok. Thanks Mic	Michele Mostarda	https://issues.apache.org/jira/browse/ANY23-50
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)	Lewis John McGibbney	0	0	1	sorry for rejoining late the discussion, been busy with yet another podling setup (Apache Syncope). @Lewis: of course, actually we are using a built-in Java engine to load classes - plugin in our cases - so it can be taken as template to manage in a quick and easy way components dynamic loading. Going to commit it, thanks -Simo	Michele Mostarda	https://issues.apache.org/jira/browse/ANY23-50

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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?	Lewis John McGibbney	0	0	1	It's worth mentioning both but they do also have slightly different risks associates, from infra's POV so we may have to split them be so that one does not hold up the other,	Andy Seaborne	https://issues.apache.org/jira/browse/ANY23-46
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.	Michele Mostarda	0	0	1	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	Lewis John McGibbney	https://issues.apache.org/jira/browse/ANY23-38
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	Lewis John McGibbney	0	0	1	he extractor factory patterns work is currently on the following branch: https://github.com/ansell/any23/compare/apache:trunk...ansell:spi-extractors You can port the patch by downloading the following URL: https://github.com/ansell/any23/compare/apache:trunk...ansell:spi-extractors.diff And running something like the following (it is a while since I ran this command myself, as any23 is the only project that I need to use it for): patch -p1 < spi-extractors.diff It will apply to the current SVN trunk HEAD as I have been regularly rebasing it onto the current Git trunk branch. The tika-1.2 patch is independent of the spi-extractors patch, so I do not know whether they will both apply. To satisfy Andy, I am downloading the patch and attaching it here, as apparently there may be some legal complaint apparently if you do it yourself!	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-26
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.	Paolo Castagna	0	0	1	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 crufiest 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	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-19

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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 crufiest 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 recognizes two types of triples, those with	Peter Ansell	0	0	1	Hi Peter, thanks for your comments (which I truly appreciate). > ... representing everything as strings As I said, java-rdfa do not provide any data checking/validation. Does Any23 currently guarantee that all the RDF triples generated are "valid" (checking syntax of URIs/IRIs, literals, etc.)? If it does, using strings would not certainly be an improvement, I agree. If it does not, then using an RDF library vs. plain strings does not provide a big advantage. > How often is Any23 currently being used as a library? I do not know the answer to this question, to me Any23 is valuable as a library (in particular if it can be used with projects such as Apache Nutch and Hadoop/MapReduce). You should also ask "how often Any23 could be used as a library?", I don't know the answer to this either. > 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. Moving the creation of Sesame specific objects to a single place in Any23 is a good idea. > Switching to another library may cause the bloat that you say you do not want. Exactly, I did not propose that, certainly not at this point in time. > Jena and its immediate dependencies is quite large compared to the modular sesame jar files I know (and I always said it in the past). Things might change in future and indeed I think that for a project such as Jena 3	Paolo Castagna	https://issues.apache.org/jira/browse/ANY23-19
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.	Lewis John McGibbney	0	0	1	Shouldn't we have a section on the web site listing all the products/initiatives that are using Any23? Starting from sindice.com and http://webdatacommons.org/ for example.	Michele Mostarda	https://issues.apache.org/jira/browse/ANY23-11
ANY23-4	I think this certainly duplicates ANY23-69 ?	Lewis John McGibbney	0	0	1	Sure, linking to duplicate and closing (the latest is more detailed).	Michele Mostarda	https://issues.apache.org/jira/browse/ANY23-4
ANY23-3	Is this not addressed?	Lewis John McGibbney	0	0	1	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?	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-3
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?	Peter Ansell	0	0	1	> Is Any23 interested in the non-RDF components of N3? Good question, Any23 should not be interested in non-RDF components of N3 and/or in N3, since it does not seem to produce anything that is not RDF. Microdata and RDFa give you RDF. Lewis, perhaps as Peter noted, you meant Turtle instead of N3 here. (The difference is important, but from a first look not so clear and people get confused or assume, wrongly, that N3 and Turtle are the same).	Paolo Castagna	https://issues.apache.org/jira/browse/ANY23-3
ANY23-2	@Peter, Are you interested in committing this in line with the new SPI Factory implementation you wrote?	Lewis John McGibbney	0	0	1	"Okay, working on this now" & "Git branch for this issue, and to fix license headers for previous extractor factories from ANY23-26 is at: https://github.com/ansell/any23/compare/ANY23-2 Attaching patch here to satisfy legal minded people."	Peter Ansell	https://issues.apache.org/jira/browse/ANY23-2

IMMUTANT

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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 {{^C}}, 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?	Toby Crawley	0	0	1	Toby 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	Ronald Suwandi	https://issues.iboss.org/browse/IMMUTANT-628
IMMUTANT-628	[~tcrawley] 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	Ronald Suwandi	0	0	1	Correct, there's no way to kill the job threads themselves. If they are long running and loop, you could have them check an atom on each iteration, and change the value of that atom when you shutdown the app.	Toby Crawley	https://issues.iboss.org/browse/IMMUTANT-628
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	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-580
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. [~tcrawley] 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	Fabio Bernasconi	0	0	1	What I was trying to say is I don't want to change the value of the `:singleton` option based on the value of `:allow-concurrent-exec?`. But there won't be a need to, since `:singleton` is true by default.	Toby Crawley	https://issues.iboss.org/browse/IMMUTANT-568
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?	Toby Crawley	0	0	1	I've finally found the time to test this and can confirm it works! Thanks!	Fabio Bernasconi	https://issues.iboss.org/browse/IMMUTANT-568

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
IMMUTANT-562	Ah, that's a good idea. Would you mind sharing an example of how you are doing that?	Toby Crawley	0	0	1	Sure! Here's a reduced example for SSE based on core.async. The basic idea is to have a fixed buffer channel of size n to which one has to put before a send!. The :on-success and :on-error callbacks of each send! then have to take from that channel. If the channel is full (i.e. n send operations are queued), further sends will thus block until at least one of the queued sends has completed. I have also created a drop-in wrapper for immutable.web.sse implementing this strategy. Perhaps it is of use to you!	Moritz Heidkamp	https://issues.iboss.org/browse/IMMUTANT-562
IMMUTANT-456	Immutable 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 Immutable's configurator put the HttpSession and cookies in the user properties?	Jim Crossley	0	0	1	Yeah, making Immutable's configurator put HttpSession and cookies in the user properties is good for my needs, but I am not sure if this will be enough for everyone else's. Might as well put all headers in to cover all the bases.	Feng Zhou	https://issues.iboss.org/browse/IMMUTANT-456
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.	Jim Crossley	0	0	1	That's a good idea. I looked at other methods of Configurator, they mostly are get methods, so it should be enough to have just the :handshake function (maybe better to call it :modify-handshake to be consistent?). I still would like Immutable's configurator to put all the http headers in userProperties for convenience though, since that's all I want and I'd guess that is the majority of use cases for modifyHandshake(). I'd suggest just making Immutable's configurator put in the headers for now, since it won't change the immutable.websocket API and will be a good enough default for a lot of people. If the need arises from other developers, then consider the :handshake function. And because :handshake can be optional, it shouldn't be a problem with compatibility if you do introduce it in the future.	Feng Zhou	https://issues.iboss.org/browse/IMMUTANT-456
IMMUTANT-453	I've fixed Immutable to honor META-INF in resources (see IMMUTANT-455), so you no longer need to put the 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, #12...	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-453
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 Immutable namespace...	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-436
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.	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-436

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
IMMUTANT-415	I just pushed (via http://github.com/immutable/immutable/commit/047042ac0) a function that exposes this setting: {{{immutable.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.	Toby Crawley	0	0	1	Seems to work as expected. Thank you!	Christian Karlsen	https://issues.iboss.org/browse/IMMUTANT-415
IMMUTANT-410	Kanwei: I tried to recreate this issue with https://github.com/tobias/immutable-410/blob/master/src/cookie_store_test/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.	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-410
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?	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-405
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?	Vlad Bokov	0	1	0			https://issues.iboss.org/browse/IMMUTANT-390
IMMUTANT-359	Have we marked them as deprecated somehow? A grep of the codebase says "nope".	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-359
IMMUTANT-334	What would the workaround be for this? I added {{{-lein-profiles [:default]}}} to the {{{:immutable}}} 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...)	Chas Emerick	0	0	1	You shouldn't have seen target/user+dev+base/classes being generated I don't believe - were you? The issue is that since Immutable 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 Immutable 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 Immutable uses internally.	Toby Crawley	https://issues.iboss.org/browse/IMMUTANT-334

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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.	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-334
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?	Toby Crawley	0	0	1	Looks good here, thanks!	Chas Emerick	https://issues.iboss.org/browse/IMMUTANT-334
IMMUTANT-334	Mariano: can you confirm as well? (I forgot that you were the original reporter).	Toby Crawley	0	0	1	switched to lein 2.3.2 and it's working ok here.	Mariano Guerra	https://issues.iboss.org/browse/IMMUTANT-334
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/56f59b98de1441727d9e797a7db57b65183724c...bd2a5fc891417c355690fe0ee436b6a73b542a46 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.	Toby Crawley	0	0	1	Looks good here w/ build 987. Thanks!	Chas Emerick	https://issues.iboss.org/browse/IMMUTANT-329
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.	Jim Crossley	0	0	1	Hi Jim, Yeah, in this short lived scenario, it behaves more like a topic. On other MQs, I've resolved this by just spinning up a new topic (usually called my.queue.name.dev-topic), and made sure that a consumer published every message it saw. Another way I've worked around this in the past is hooking into the broker's journal. I'm more than happy to do the duplicate topic here - I was just trying to make the best use of the network resources. The reason you can't use a topic for the origin destination, is that you actually want queuing properties in the production system (as an easy example, so you can do request/respond). I'm totally ok with JMS. I haven't hit a hard limitation I wasn't able to easily work around. Thanks for looking into this!	Paul deGrandis	https://issues.iboss.org/browse/IMMUTANT-326

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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.	Jim Crossley	0	0	1	I've been toying with the internals a little more and I found something that works fairly well for me. Auto-ack is fine (since, like I said early, maps nicely to the transactions you'd want to ack anyway). In scenarios where you want to squeeze as much performance out of the MQ as possible, it's actually best to use the lower level JMS/Core APIs directly and hand control your ACKing and transactions. Still, on my laptop, I was easily processing 2300 msgs/sec with the immutable.messaging stuff. I'm ok with you removing this ticket	Paul deGrandis	https://issues.iboss.org/browse/IMMUTANT-326
IMMUTANT-317	I looked into what it would take to do this at deploy-time inside Immutable, and it's currently not straightforward, because: 1) we explicitly strip any Immutable 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 Immutable 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) {{leininingen.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.	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-317
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.	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-287
IMMUTANT-287	cache.org: the default locking mode is nil, but the description states: "One of :optimistic or :pessimistic". so what does nil mean?	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-287
IMMUTANT-287	iboss.org: do we need to update the logging section?	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-287

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
IMMUTANT-219	Should this be resolved as 'Done'? We didn't actually add scheduled delivery to publish.	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-219
IMMUTANT-195	Pragmatically, I'm guessing that filing 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?	Norman Richards	0	0	1	I agree that this is a bug in catalina. It looks like this exact same issue came up on TorqueBox recently (TORQUE-955), and the fix was to pass a different, unescaped path-info to the handler. I'd like to implement the same thing in Immutant so the path-info will at least be correct. I'd like to think about & discuss the implications of enabling ALLOW_BACKSLASH by default.	Toby Crawley	https://issues.iboss.org/browse/IMMUTANT-195
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.	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-195
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 AWSPIG, 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?	Ian Eslick	1	0	0			https://issues.iboss.org/browse/IMMUTANT-182
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.	Jim Crossley	0	1	0			https://issues.iboss.org/browse/IMMUTANT-107
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.iboss.org/reteasy/hornetq-rest/1.0-beta-1/userguide/html_single/index.html#d0e162 The URI for a topic/queue resource 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?	Matthew Giannini	0	0	1	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 #".^.?queue". This would make queues starting with "/queue", ".queue", "queue", "xqueue" all legal. How does that sound?	Toby Crawley	https://issues.iboss.org/browse/IMMUTANT-96


Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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 #"?queue". This would make queues starting with "/queue", ".queue", "queue", "xqueue" all legal. How does that sound?	Toby Crawley	0	0	1	It seems fine to make it very flexible with #"?queue". If you don't want it to be as flexible i think #"[\\. \\/]queue" would be equally good. That is, only allow queue or topic names to start with ".queue" or "/queue" instead of <anything>queue. But i don't feel strongly about that.	Matthew Giannini	https://issues.iboss.org/browse/IMMUTANT-96
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?	Toby Crawley	0	0	1	v0.8.2 seems to have fixed all issues that I observed...thanks!	Tom Hicks	https://issues.iboss.org/browse/IMMUTANT-84
IMMUTANT-78	Is this one still relevant?	Jim Crossley	0	1	0			https://issues.iboss.org/browse/IMMUTANT-78
IMMUTANT-55	What does this even mean?	Toby Crawley	1	0	0			https://issues.iboss.org/browse/IMMUTANT-55
IMMUTANT-27	Jim: Isn't this kind of done? Or are is this jira waiting on more formal Immutant support on OS?	Toby Crawley	0	0	1	I kinda want to leave this open until we have an Immutant cartridge built.	Jim Crossley	https://issues.iboss.org/browse/IMMUTANT-27
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.	Bruno Oliveira	1	0	0			https://issues.iboss.org/browse/IMMUTANT-24

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
IMMUTANT-21	Hey Jimbotron - is there anything else we need to do here?	Toby Crawley	0	0	1	Reloading works as long as you pass absolute paths to ring.middleware.reload/wrap-reload. Note the function name and its signature has changed in recent versions of Ring, but as of Ring 1.1.1, this should work: <pre>(ns my-project-name.init (:require [immutant.web :as web] [immutant.utilities :as util] [ring.middleware.reload :as ring] [noir.server :as server])) (server/load-views (util/app-relative "src/my_project_name/views")) (web/start "/" (ring/wrap-reload (server/gen-handler {:mode :dev :ns 'my-project-name})) {:dirs [(util/app-relative "src")]}))</pre> The relevant bit is the wrapping of the noir handler function with an additional check for the fully-qualified path to "src". And the :dirs key to the options map is the new signature mentioned above.	Jim Crossley	https://issues.iboss.org/browse/IMMUTANT-21
IMMUTANT-21	Hey Jimbotron - is there anything else we need to do here?	Toby Crawley	0	0	1	Reloading works as long as you pass absolute paths to ring.middleware.reload/wrap-reload. Note the function name and its signature has changed in recent versions of Ring, but as of Ring 1.1.1, this should work: <pre>(ns my-project-name.init (:require [immutant.web :as web] [immutant.utilities :as util] [ring.middleware.reload :as ring] [noir.server :as server])) (server/load-views (util/app-relative "src/my_project_name/views")) (web/start "/" (ring/wrap-reload (server/gen-handler {:mode :dev :ns 'my-project-name})) {:dirs [(util/app-relative "src")]}))</pre> The relevant bit is the wrapping of the noir handler function with an additional check for the fully-qualified path to "src". And the :dirs key to the options map is the new signature mentioned above.	Jim Crossley	https://issues.iboss.org/browse/IMMUTANT-21

DASHBUILDER

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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?	David Gutierrez	0	0	1	David Gutierrez 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?	Liz Clayton	https://issues.iboss.org/browse/DASHBUILDE-237
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?	Liz Clayton	0	0	1	Hello David Gutierrez Liz Clayton Just getting started on this one. But, generally, I agree that local filters and global filters can coexist. Local filters should be placed on the chart and global filters should be placed at the top of the page. https://kiegroup.github.io/kie-uxd-sandbox/prototype/uxd-playground/metrics.html I still have some gaps. 1) What are the three icons above the Commit History table? 2) I would rather the dropdown filters say something like No Filter ("Nothing selected" is the default and I am not sure how to change it). 3) I was not able to get the first two charts to be the same height. Nor was I able to reproduce David's Commits Per Author chart. In any case, this is a first pass but it will help us evolve the design over time. Thank you, Matthew	Matthew Stevens	https://issues.iboss.org/browse/DASHBUILDE-237
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?	David Gutierrez	0	0	1	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?	Matthew Stevens	https://issues.iboss.org/browse/DASHBUILDE-237
DASHBUILDE-237	I seem to be losing the thread here a bit. A quick conversation	Matthew Stevens	0	1	0			https://issues.iboss.org/browse/DASHBUILDE-237
DASHBUILDE-237	Hello [~davidredhat] do you have time to chat on Thursday or	Matthew Stevens	0	1	0			https://issues.iboss.org/browse/DASHBUILDE-237
DASHBUILDE-237	The following screenshot shows the toolbar implemented in	David Gutierrez	1	0	0			https://issues.iboss.org/browse/DASHBUILDE-237
DASHBUILDE-232	Hi Team, It would be a great help if a workaround is provided	Munishan Narang	0	1	0			https://issues.iboss.org/browse/DASHBUILDE-232

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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? :-)	David Gutierrez	0	0	1	David Gutierrez 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. 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!	Liz Clayton	https://issues.iboss.org/browse/DASHBUILDE-230
DASHBUILDE-230	[~davidredhat] I sent my comments to a design related jira,	Liz Clayton	1	0	0			https://issues.iboss.org/browse/DASHBUILDE-230
DASHBUILDE-229	Initial implementation done. Preview video =>	David Gutierrez	1	0	0			https://issues.iboss.org/browse/DASHBUILDE-229
DASHBUILDE-228	[~davidredhat] or [~porcelli], can you guys please put this on Sprint5?	Eder Ignatowicz	0	0	1	Alexandre Porcelli made changes - 31/May/17 3:20 PM Field Original Value New Value Sprint Sprint5 [6126]	Alexanfre Porcelli	https://issues.iboss.org/browse/DASHBUILDE-228

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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 ). 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?	David Gutierrez	0	0	1	David Gutierrez overall your proposal sound fine to me. Except for one detail: I'm not sure we can assume fixed width for characters (e.g. with some crazy chinese font this assumption that 1 char == 10px might break.) But if we're going to set button width based on width of labels displayed, you could calculate (# chars in label) / (# chars in all labels) and set the width to percentage perhaps?	Jan Hrcek	https://issues.iboss.org/browse/DASHBUILDE-226
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?	David Gutierrez	0	0	1	David Gutierrez this is great improvement! I personally like option #2 better, as the available actions are less "scattered" and more "all things in one place".	Jan Hrcek	https://issues.iboss.org/browse/DASHBUILDE-218
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: IScreen Shot 2017-05-24 at 10.40.20 AM.png thumbnail!	Liz Clayton	0	0	1	Option #2 wins. Liz Clayton I will also implement the latest changes suggested.	David Gutierrez	https://issues.iboss.org/browse/DASHBUILDE-218
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?	Jan Hrcek	1	0	0			https://issues.iboss.org/browse/DASHBUILDE-93
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?	Roger Martínez	0	0	1	Done (master) https://github.com/dashbuilder/dashbuilder/commit/8287fd70362c0c9e7de8c4a10e8ea2a1ce13d307	David Gutierrez	https://issues.iboss.org/browse/DASHBUILDE-70

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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?	Jan Hrcek	0	0	1	Hi Jan Hrcek Honestly I don't have the details about it. We are building our applications on top of UberFire, and it's UF who provides the new integration with patternfly and bootstrap3, so as we are just using uberfire based components, we do not aware of the underlying changes at all... But as I far as I know, for the Patternfly integration side, I think there are some new stylesheets that have been added/modified, and for the bootstrap3 side, there are new widgets that we can use. But i'm only using new widgets provided by bootstrap 3, from the patternfly side I'm not aware of it as it's provided by the framework layer (UF), from the application side we only have to depend on the new "UF View for Patternfly", and that should be all... You can find more information about the Patternfly integration in UF at https://github.com/uberfire/uberfire/tree/master/uberfire-workbench/uberfire-workbench-client-views-patternfly. You can ask too Alexandre Porcelli or any guy from UF team, they have the details about it. Hope clarifying a little bit... Thx and regards	Roger Martinez	https://issues.iboss.org/browse/DASHBUILDE-28
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?	Diogo Pontual	0	1	0			https://issues.iboss.org/browse/DASHBUILDE-10

DROOLS

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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...	Duncan Doyle	0	0	1	"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." & "Maciej makes good points, I think this should really only be used for stateless DMN services for dynamic environments. I don't think we would be able to preserve it between restarts though, so it would really be a "demo" feature. " & "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...."	Maciej Swiderski & Edson Tirelli & Maciej Swiderski	https://issues.iboss.org/browse/DROOLS-1652
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...	Maciej Swiderski	0	1	0			https://issues.iboss.org/browse/DROOLS-1652
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 serv...	Maciej Swiderski	0	1	0			https://issues.iboss.org/browse/DROOLS-1652
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 uni...	Edson Tirelli	1	0	0			https://issues.iboss.org/browse/DROOLS-1577
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.	Edson Tirelli	0	0	1	Edson Tirelli Thanks for your suggestion. The documentation is wrong. Its mentions in seconds. But its actually in millisecs. I have configured as 60000. It scans every minute now. Thats really a fantastic option. Sorry for the bug. Please close it. https://docs.iboss.org/drools/release/6.4.0.CR2/drools-docs/html/ch22.html Automatic Update: If you want a deployed Container to always have the latest version of your deployment without manually editing it, you will need to set the Version property to the value of LATEST and start a Scanner. This will ensure that the deployed provision always contains the latest version. The Scanner can be started just once on demand by clicking the Scan Now button or you can start it in the background with scans happening at a *specified interval (in seconds).*You can also set this value to LATEST when you are first creating this deployment. The Resolved Release Id in this case will show you the actual, latest version number.	Dhamodharan Krishnan	https://issues.iboss.org/browse/DROOLS-1561
DROOLS-1561	[~mfusco] can you please quickly fix this in the docs? Use this ticket for it.	Edson Tirelli	0	0	1	Documentation fixed by https://github.com/kiegroup/kie-docs/commit/a26761d9239bfdadec192bbc874e26f60541e5d8	Mario Fusco	https://issues.iboss.org/browse/DROOLS-1561

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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?	Mario Fusco	0	0	1	I attached an artificial example using java.io.StringWriter (which is AutoCloseable) (AutoCloseableSample.java). In one method (desiredState) I take advantage of autocloseable feature and the syntax is pretty compact. In another class, the autocloseable is not being used and I end up not even being able to compile the code. Of course, the example is again artificial, just to make the point. My actual scenario is that: I use drools in a Java 8 multi-user (it is a webapp) scenario I create one KieSession every time an user invokes that specific scenario and I want to release the resources by calling the "dispose" method. Inserting the objects in the KieSession using lambdas is pretty elegant (at least in my case). Generally, the drools documentation emphasizes the need for the developer to call the "dispose" method once he is done with the KieSession (assuming he created it using a new* method on the KieBase). I see this no different that invoking close on input streams or database result sets... Now, the AutoCloseable has been introduced in Java 7. If the intention is to keep the binaries compatible with 1.6 or below, then please forget about all this...	Tech Meshter	https://issues.iboss.org/browse/DROOLS-1444
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 th...	Mario Fusco	1	0	0			https://issues.iboss.org/browse/DROOLS-1444
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 re...	Mario Fusco	0	0	1	Mario Fusco added the PR. The hashcode on ProjectClassLoader is very simple - it delegates to hashcode of its wrapped classloaders making the hashcode for a projectclassloader equal for all instances wrapping the exact same instances for underlying loaders - which in here is just whats needed	Hans Lund	https://issues.iboss.org/browse/DROOLS-1443
DROOLS-1419	[~geoffrey] it seems this JIRA might need an update to resolved?	Kris Verlaenen	1	0	0			https://issues.iboss.org/browse/DROOLS-1419
DROOLS-1398	[~Rikkola] sorry, I am confused. Does this change need to be done in the workbench or in core?	Edson Tirelli	0	1	0			https://issues.iboss.org/browse/DROOLS-1398
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, b...	Tibor Zimányi	0	0	1	Tibor Zimányi, thank you for integrating the test case. I'm new to JMH so I'm not 100% sure how kie-benchmarks is supposed to work but as far as Drools1387AccumulateBenchmark/Drools1387ExistsBenchmark are concerned, it looks correct and your result matches the original reproducer as well.	Toshiya Kobayashi	https://issues.iboss.org/browse/DROOLS-1387

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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 tell...	Toshiya Kobayashi	0	0	1	Toshiya Kobayashi, there are 2 other aspects that are not immediately clear of your benchmark suite. 1. I don't think you should run the benchmark in a multithreaded environments, unless you don't want to check how drools performs in this specific situation. However benchmarking in a concurrent use case is a totally different matter and I believe that in this case thread contention can negatively affect the reliability of your benchmark. 2. KieSession creation and dispose shouldn't be part of the elapsed benchmark time. I saw you're using a pool of KieSession that you put back into a concurrent data structures after having removed all facts from there, but also these operations are part of the elapsed time of your benchmark and I believe that this is wrong. More in general writing a reliable benchmark is a tricky task and that's why you should always use a well established framework like JMH. As said by Tibor Zimányi feel free to join us on IRC in case you want to discuss this. Regarding the rules you have in your benchmark, I can confirm that RETEEO is a bit faster than PHREAK, but the difference is below the 10% and however not even close to the 30% you claimed. More important I believe that the difference is quite unrelated by the type of nodes but caused by the fact that in your use case all segments and path are linked (there are facts satisfying all rules) and then the lazy nature of phreak is not an advantage in this case. I'm now profiling the latest benchmarks developed by Tibor Zimányi but at this moment I'm skeptical that I will be able to find any low hanging fruit to improve PHREAK performances in this specific case.	Mario Fusco	https://issues.iboss.org/browse/DROOLS-1387
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...	Tibor Zimányi	0	1	0			https://issues.iboss.org/browse/DROOLS-1387
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...	Mario Fusco	0	0	1	Finally, I understand that all rules are linked in this test case and it results in the performance difference. Wrote a blog post from the investigation (https://tkobayas.wordpress.com/2017/01/17/phreak-unlinking/). Closing this JIRA. Thanks, Mario and Tibor!	Toshiya Kobayashi	https://issues.iboss.org/browse/DROOLS-1378
DROOLS-1355	[~porcelli] can you please take a look at this and check how that relates to the work you are doing?	Edson Tirelli	1	0	0			https://issues.iboss.org/browse/DROOLS-1355
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...	Petr Široký	1	0	0			https://issues.iboss.org/browse/DROOLS-1345
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 incorpor...	Marek Czernek	1	0	0			https://issues.iboss.org/browse/DROOLS-1344
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	Kavi sri	0	1	0			https://issues.iboss.org/browse/DROOLS-1319
DROOLS-1270	Can you please clarify your request and provide an example?	Edson Tirelli	0	0	1	Edson Tirelli I need to delete or update facts. Normally I would use GetFactHandlesCommand to get a list of handles then I would iterate to find the relevant fact to delete or modify. However the current list of supported commands does not include fact handle based commands: http://docs.iboss.org/drools/release/6.4.0.Final/drools-docs/html/ch22.html#d0e24762 The list of commands available in Java are here: https://github.com/droolsjbpm/drools/tree/master/drools-core/src/main/java/org/drools/core/command/runtime/rule	Dan Climpoesu	https://issues.iboss.org/browse/DROOLS-1270

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
DROOLS-1250	I fixed the issue by adding <code>_*dialect "mvel"*_</code> at the two rules in <code>Sample.drl</code> . Is it a bug in the drools wizard that dialect is missing in the examples? or it supposes to work without dialect? Thanks.	Tianwei Chen	1	0	0			https://issues.iboss.org/browse/DROOLS-1250
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.	Sarit Adhikari	1	0	0			https://issues.iboss.org/browse/DROOLS-1250
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?	Sean Wu	0	0	1	The fix for this will not be available until 6.5.Final community (6.4 product) is released, shortly. At that time I will add a comment to this JIRA indicating where the new bits can be downloaded.	Robert(Bob) Brodt	https://issues.iboss.org/browse/DROOLS-1250
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...	Mario Fusco	0	0	1	Hi, Mario, Yes, correct.	Osamu Kuniyasu	https://issues.iboss.org/browse/DROOLS-1225
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...	Tom Pijl	0	1	0			https://issues.iboss.org/browse/DROOLS-1221
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	Sante Stanisci	0	1	0			https://issues.iboss.org/browse/DROOLS-1168
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>	Mario Fusco	0	0	1	I tried it, but is a trivial solution, because Drools can set multiple fields in multiple class that extends EntityBase in this example case rule "completeCodDivisaContAndCambioDivisaCont" when eval(checkRulesEnabled == 0) (JtbRLavt(codDiviCont == null cambioDiviCont == null, \$entityFind : this) or JtbRLavr(codDiviCont == null cambioDiviCont == null, \$entityFind : this) or GtbSpes(codDiviCont == null cambioDiviCont == null, \$entityFind : this) or GtbTipiPaga(codDiviCont == null cambioDiviCont == null, \$entityFind : this) or PartitaMag(codDiviCont == null cambioDiviCont == null, \$entityFind : this) or CtbMovt(codDiviCont == null cambioDiviCont == null, \$entityFind : this) or CtbMovr(codDiviCont == null cambioDiviCont == null, \$entityFind : this) or CtbMovrCoan(codDiviCont == null cambioDiviCont == null, \$entityFind : this) or CtbMovi(codDiviCont == null cambioDiviCont == null, \$entityFind : this) or CtbParr(codDiviCont == null cambioDiviCont == null, \$entityFind : this) or VtbArti(codDiviCont == null cambioDiviCont == null, \$entityFind : this) or MtbSpes(codDiviCont == null cambioDiviCont == null, \$entityFind : this)) \$entity : EntityBase() from \$entityFind then java.util.Map() from \$entityFind.completeCodDivi(codDiviCont == null cambioDiviCont == null, \$entityFind : this) end	Sante Stanisci	https://issues.iboss.org/browse/DROOLS-1168
DROOLS-1168	This is a workaround found 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 //...</code>	Sante Stanisci	0	1	0			https://issues.iboss.org/browse/DROOLS-1168
DROOLS-1168	Solution proposed in last post is not possible because \$entity : EntityBase() from \$entityFind \$entity is EntityBase and setCodDiviCont(), setCambioDiviCont() methods are in one of class listed (JtbRLavt, JtbRLavr, GtbSpes,) In compile phase dro...	Sante Stanisci	0	1	0			https://issues.iboss.org/browse/DROOLS-1168

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
DROOLS-1168	I would like a solution, I can invoke the mask calculation in some way?	Sante Stanisci	0	1	0			https://issues.iboss.org/browse/DROOLS-1168
DROOLS-1168	May I understand why you're insisting using property reactivity for this use case? Which advantages do you see?	Mario Fusco	0	0	1	the advantages are that I want to manage the dependencies of the class fields and drools allows me to trigger cascading rules.	Sante Stanisci	https://issues.iboss.org/browse/DROOLS-1168
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.	Mario Fusco	0	0	1	No, because I want PropertyReactive. With no property reactive rules have cascade firing?	Sante Stanisci	https://issues.iboss.org/browse/DROOLS-1168
DROOLS-1168	No, because I want PropertyReactive. With no property reactive rules have cascade firing?	Sante Stanisci	0	1	0			https://issues.iboss.org/browse/DROOLS-1168
DROOLS-1168	And how to Drools know that property was changed?	Sante Stanisci	0	1	0			https://issues.iboss.org/browse/DROOLS-1168
DROOLS-1168	Ok, it is just for efficiency. But you consider my request for enhancement next version? Best Regards	Sante Stanisci	1	0	0			https://issues.iboss.org/browse/DROOLS-1168
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-pare...	Geoffrey De Smet	0	0	1	Yes, ultimately kie-parent-with-dependencies should go. I will keep it around until all the repos are migrated (I would like to do the migration one repo at a time). About 'kie-platform-bom', I am not sure what the value is. Our projects won't use it, but it might be useful for our users in case they want single POM which will import all the versions they need. This initial proposal does not include it, but we may decide to add it back later on.	Petr Šíroký	https://issues.iboss.org/browse/DROOLS-1137
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 L...	Geoffrey De Smet	0	0	1	I found the cause for this on my end, there were some null properties in the entities optaplanner was operating on.	Ido Flax	https://issues.iboss.org/browse/DROOLS-1071
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.	Geoffrey De Smet	0	0	1	yes, it was line 86/87	Ido Flax	https://issues.iboss.org/browse/DROOLS-1071
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/marshall...	Maciej Swiderski	0	0	1	Hello Maciej, I was not aware of this system parameter, thanks! It seems enough to me, I was just afraid that it could be a problem when generating requests to the server. Regarding changing the actual pattern, I have prepared a PR for this (to save you some time), it is not a problem at all if you want to refuse it: https://github.com/droolsjbpm/droolsjbpm-integration/pull/289 Thanks again and sorry for not checking upstream, I had only checked 6.3.x	William Siqueira	https://issues.iboss.org/browse/DROOLS-1053
DROOLS-1052	[~psiroky] only for master or for other branches too ?	Michael Biarnes Kiefer	0	0	1	I believe it's enough to do it just on master.	Petr Šíroký	https://issues.iboss.org/browse/DROOLS-1052
DROOLS-1052	Should I only look in the java files or in the drl files too? What should I do in this case: /** * Finds currently available server ports. * * @author Apache Directory Project * @see <a href="http...	Michael Biarnes Kiefer	0	0	1	I am not a lawyer, so I am not quite sure. I would say, to be safe, we should keep the @author tags in sources which were directly copy&pasted from somewhere else (e.g. we are not the original authors of that code).	Petr Šíroký	https://issues.iboss.org/browse/DROOLS-1052
DROOLS-972	Can you please send a PR adding a new example in drools-examples-cdi so I could start working from there?	Mario Fusco	0	0	1	I've added some tests in linked PR that shows how the @KSession annotation fails to follow the bean scope as the specification dictates. Each bean also includes an injection of MyBean which is used to check the correct behaviour. In the PR is also included a new scope definition, and I'm working on trying to understand if we can create our own scope to wrap a KieSession lifecycle. This scope should be similar to the @ConversationScoped but not linked at all to the HttpSession, but from the behaviour point of view should be extremely similar	Mario Fusco	https://issues.iboss.org/browse/DROOLS-972
DROOLS-963	This should be fixed for 6.4.0.Final, does it sound possible?	Mauricio Salatino	0	0	1	This will be fixed in 6.4.0.Final. It got resolved as part of Maven upgrade from 3.2.2 to 3.2.5.	Petr Šíroký	https://issues.iboss.org/browse/DROOLS-963
DROOLS-905	Thanks for the pointers [~martenscs]. Is there any other prefix we need to worry about? Or does startsWith("bundle") cover all the needed cases?	Petr Šíroký	0	1	0			https://issues.iboss.org/browse/DROOLS-905

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
DROOLS-901	any news here [-mfusco)?	David virgil naranjo	1	0	0			https://issues.jboss.org/browse/DROOLS-901
DROOLS-823	Hi Geoffrey, I don't know what should be the expected behaviour? i.e. I run in drools-jdeps -jdkinternals ./drools-verifier/src/test/resources/org/drools/verifier/model.jar only for test purposes. It says: Picked up JAVA_TOOL_OPTIONS: -Dfile.encoding=...	Michael Biarnes Kiefer	0	1	0			https://issues.jboss.org/browse/DROOLS-823
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 sun.misc.Unsafe usages from our depended libraries (com.google.common l...	Geoffrey De Smet	1	0	0			https://issues.jboss.org/browse/DROOLS-823
DROOLS-821	Last week [-geoffrey] added new feature for optaplanner-engine (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 ...	Petr Šíroký	0	0	1	Note that drools-osgi contains jbpm too (and now optaplanner too), hence the rename to kie-osgi proposal.	Geoffrey De Smet	https://issues.jboss.org/browse/DROOLS-821
DROOLS-760	There were some changes related to formatting done by [-kgaevskij]. I believe the formatters are now up-to-date. [-kgaevskij] can you please confirm that?	Petr Šíroký	0	0	1	Hi Petr Šíroký, we decided to use actual formatter in Stunner. I saw some tries in other repositories like Uberfire, but mine goal was Stunner only, so I can't guarantee another repositories use it.	Kirill Gaevskii	https://issues.jboss.org/browse/DROOLS-760
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 _everything_ gets fixed, but that may very well never happen.	Petr Šíroký	1	0	0			https://issues.jboss.org/browse/DROOLS-760
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	Mark Proctor	1	0	0			https://issues.jboss.org/browse/DROOLS-729
DROOLS-686	Petr in that SNAPSHOT setup, after building successfully once, did you try plugging out your network cable and building again? The second time it must be able build offline (and that originally didn't work with pressgang 3).	Geoffrey De Smet	0	0	1	I did not plug out the the network cable physically, but rather disabled the networking interfaces in the system (which is effectively the same, I was not able to connect to any resource on the network/internet). I tried the build multiple times, so I think we should be ok.	Petr Šíroký	https://issues.jboss.org/browse/DROOLS-686
DROOLS-660	public Class<?> fastFindClass(final String name) { lock.lock(); Class<?> cls = findLoadedClass(name); if (cls == null) { final byte[] clazzBytes = this.store.read(convertClassToResourceP...	chi chi	0	0	1	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	Davide Sottara	https://issues.jboss.org/browse/DROOLS-660

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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...	Davide Sottara	0	0	1	I'm sorry, this is the first time I report a question, I think this is a bug in 5.6.0.Final; this occurs when we down rule and at the same time we get 100 requests per seconds; then we got a deadlock problem Found one Java-level deadlock: ===== "http-8108-21": waiting for ownable synchronizer 0x000000042fd7a3c8, (a java.util.concurrent.locks.ReentrantLock\$NonfairSync), which is held by "http-8108-2" "http-8108-2": waiting for ownable synchronizer 0x000000042fde6b70, (a java.util.concurrent.locks.ReentrantLock\$NonfairSync), which is held by "http-8108-21" Java stack information for the threads listed above: ===== ==== "http-8108-21": at sun.misc.Unsafe.park(Native Method) parking to wait for <0x000000042fd7a3c8> (a java.util.concurrent.locks.ReentrantLock\$NonfairSync) at java.util.concurrent.locks.LockSupport.park(LockSupport.java:156) at java.util.concurrent.locks.AbstractQueuedSynchronizer.parkAndCheckInterrupt(AbstractQueuedSynchronizer.java:811) at java.util.concurrent.locks.AbstractQueuedSynchronizer.acquireQueued(AbstractQueuedSynchronizer.java:842) at java.util.concurrent.locks.AbstractQueuedSynchronizer.acquire...	chi chi	https://issues.jboss.org/browse/DROOLS-660
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 w...	Davide Sottara	1	0	0			https://issues.jboss.org/browse/DROOLS-660
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.	Michael Biarnes Kiefer	0	0	1	Yes, James Dennes (Github Staff) told me it's done on all repo's of https://github.com/droolsjbpm on all branches (but not forks of course) on 20-JAN-2015. No idea if it automatically applies to new repo's in that same github organization...		https://issues.jboss.org/browse/DROOLS-656
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 op...	Geoffrey De Smet	0	0	1	yeah - i tried with a rep I will remove right now. I tried from two sides - local as mbiarnes and on the server as kiereleaseuser. kiereleaseuser modified some files and committed and pushed them. mbiarnes modified the same files and tried ti do a git push -f origin master and get: Counting objects: 17, done. Delta compression using up to 4 threads. Compressing objects: 100% (12/12), done. Writing objects: 100% (17/17), 5.12 KiB 0 bytes/s, done. Total 17 (delta 1), reused 3 (delta 0) remote: error: GH003: Sorry, force-pushing to master is not allowed. To git@github.com:droolsjbpm/testRepo.git ! [remote rejected] master -> master (pre-receive hook declined) error: failed to push some refs to 'git@github.com:droolsjbpm/testRepo.git' I will close this JIRA then	Michael Biarnes Kiefer	https://issues.jboss.org/browse/DROOLS-656
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 differen...	Edson Tirelli	0	0	1	Yes, but making "declare" vs "define" explicit.	Davide Sottara	https://issues.jboss.org/browse/DROOLS-652

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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...	Matteo Mortari	1	0	0	using just "event X ..." is very similar to "class X ...", which imho would be more appropriate for a declaration, or a joint declaration/definition.		https://issues.jboss.org/browse/DROOLS-599
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/droolsj...	Matteo Mortari	1	0	0			https://issues.jboss.org/browse/DROOLS-599
DROOLS-566	Thanks, what about the other way around, eg DROOLS SNAPSHOT with mvel 2.1.9 - does that work?	Claus Ibsen	0	0	1	No. The problem is that between mvel 2.1.x and 2.2.x we migrated from asm 3.3 to asm 5. Both mvel and drools use asm directly but unfortunately in asm 5 they changed some interfaces to abstract classes, causing this incompatibility.	Mario Fusco	https://issues.jboss.org/browse/DROOLS-566
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?	Mario Fusco	0	1	0			https://issues.jboss.org/browse/DROOLS-552
DROOLS-529	Could you please submit the rules you used to benchmark the engine? Thanks!	Davide Sottara	0	0	1	Attaching a zipped maven project. Can be imported in eclipse to run. It has a runnable class. More details inside the README at top level. Please let me know if there are any questions. Tune some number based on the machine you want to run this on. I hooked up the JVM to jvisualvm and the thread view shows information about thread contention. You could notice the difference between the time spent on monitor clearly in that view between RETEEO vs PHREAK execution. I also have included a patched SegmentUtilities class (from 6.0.1-Final), it is named as .jav to avoid being used automatically. You can rename to .java and then rebuild/re-run to see the difference again.	Ravi Sanwal	https://issues.jboss.org/browse/DROOLS-529
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?	Mark Proctor	0	0	1	Attaching a zipped maven project. Can be imported in eclipse to run. It has a runnable class. More details inside the README at top level. Please let me know if there are any questions. Tune some number based on the machine you want to run this on. I hooked up the JVM to jvisualvm and the thread view shows information about thread contention. You could notice the difference between the time spent on monitor clearly in that view between RETEEO vs PHREAK execution. I also have included a patched SegmentUtilities class (from 6.0.1-Final), it is named as .jav to avoid being used automatically. You can rename to .java and then rebuild/re-run to see the difference again.	Ravi Sanwal	https://issues.jboss.org/browse/DROOLS-529
DROOLS-507	Is there any update on this? i'm facing the same issue..	Dushman Elvitigala	1	0	0			https://issues.jboss.org/browse/DROOLS-507
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...	Davide Sottara	0	1	0			https://issues.jboss.org/browse/DROOLS-507
DROOLS-492	[~cmoullard] 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?	Marco Rietveld	1	0	0			https://issues.jboss.org/browse/DROOLS-492
DROOLS-483	Thanks for reporting this, can you submit a PR? It will be easier to discuss the details Thanks! Davide	Davide Sottara	0	0	1	Here's the first draft of the pull request.	Frank Pavageau	https://issues.jboss.org/browse/DROOLS-483

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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 modifc...	Frank Pavageau	0	0	1	"Hi and sorry for the huge delay in dealing with this. The improvement you're suggesting is ok in principle, but as you anticipated, it doesn't work well in an OSGi environment and this is why we didn't merge your pull request yet. I believe this further issue could be fixable by making our OSGi Activators to recognize this situation and deal with it, but this needs a bit of time to be done correctly. That is why we chose the safest (and slowest) path for now. Anyway I think that what you're proposing makes completely sense and I'll try to extend what you did to make it work also on OSGi asap. You don't need to rebase your PR, I'll do it when I'll have time to work on it. Thanks for your help and patience, Mario" & "Fixed by https://github.com/droolsjbpm/drools/commit/67ba90c392027fd199b9f2141fd58163d850ccc2"	Mario Fusco	https://issues.jboss.org/browse/DROOLS-483
DROOLS-470	This issue still hurts. Any chance we can get it for 6.4?	Geoffrey De Smet	1	0	0			https://issues.jboss.org/browse/DROOLS-470
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?	Geoffrey De Smet	1	0	0			https://issues.jboss.org/browse/DROOLS-470
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:att...	Damon Horrell	1	0	0			https://issues.jboss.org/browse/DROOLS-424
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 b...	Davide Sottara	0	0	1	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	Adar Dembo	https://issues.jboss.org/browse/DROOLS-365
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 asse...	Adar Dembo	0	1	0			https://issues.jboss.org/browse/DROOLS-365
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 runni...	Henryk Konsek	1	0	0			https://issues.jboss.org/browse/DROOLS-355

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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!	Petr Šíroky	0	0	1	Hi, I took over from Henryk Konsek to provide a reproducer. I pushed this to GitHub: https://github.com/ctron/DROOLS-355 Simply run 'ant -f build.ant' and it should download, copy and run Equinox, showing the issue.	Jens Reimann	https://issues.jboss.org/browse/DROOLS-355
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 Const...	Jens Reimann	0	0	1	Jens Reimann that is indeed an issue. kie-remote-ws-common should be OSGi bundle. The reason why we are not seeing this on Karaf is because we use the wrap protocol (which is of course far from ideal).	Petr Šíroky	https://issues.jboss.org/browse/DROOLS-355
DROOLS-355	And which is not available on Equinox ;-) So should I open a new issue for that? Or can you do that. I guess you can provide better information than I can.	Jens Reimann	0	0	1	Jens Reimann I've created https://issues.jboss.org/browse/JPBM-5250 and will look into this (hopefully soon). However, we may have bigger issue. Looking at https://github.com/droolsjbpm/droolsjbpm-integration/blob/6.5.x/drools-osgi/drools-karaf-features/src/main/filtered-resources/repository/features-core.xml#L210 the kie-remote-client seems to depend also on jsoup and resteasy-jaxrs-api and it also uses the wrap protocol. The thing is that we don't really have any direct control over those, so getting them OSGi-fied may be much more difficult.	Petr Šíroky	https://issues.jboss.org/browse/DROOLS-355
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 ?	Thierry Guerin	0	0	1	Thanks. Regarding the Pull Request on drools part; our code repository is hosted on github: The drools-decisiontables module is within the drools Repository @ https://github.com/droolsjbpm/drools . If you fork this repository, make your changes in your fork and commit to github you will be able to create a Pull Request for your changes to be reviewed and included in the codebase. Think more "git workflow" than "JIRA workflow". Thanks, Mike	Michael Anstis	https://issues.jboss.org/browse/DROOLS-337
DROOLS-315	Edson - what state is the patch in - has it made it in - will it make it?	Michael Neale	1	0	0			https://issues.jboss.org/browse/DROOLS-315
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 ...	Geoffrey De Smet	1	0	0			https://issues.jboss.org/browse/DROOLS-315
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 deadl...	Davide Sottara	1	0	0			https://issues.jboss.org/browse/DROOLS-300
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.	Shane Bryzak	1	0	0			https://issues.jboss.org/browse/DROOLS-299
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.	Geoffrey De Smet	0	0	1	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?	Michael Biarnes Kiefer	https://issues.jboss.org/browse/DROOLS-297
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 proje...	Michael Biarnes Kiefer	0	1	0			https://issues.jboss.org/browse/DROOLS-297

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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"?	Geoffrey De Smet	0	0	1	Since during a release all projects are examined by a "find" if there are still -SNAPSHOT [not all -SNAPSHOTS are changed by the script update-versions.sh] an automatic way to control SNAPSHOTS, after this has done manually in my opinion is dispensable. This JIRA will be closed.	Michael Biarnes Kiefer	https://issues.jboss.org/browse/DROOLS-297
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 t...	Davide Sottara	1	0	0			https://issues.jboss.org/browse/DROOLS-290
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 ...	Petr Široký	1	0	0			https://issues.jboss.org/browse/DROOLS-166
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 knowledgepack...	Prashant Hariharan	0	1	0			https://issues.jboss.org/browse/DROOLS-147
DROOLS-113	Is some work being done on this feature?	Guillaume Taglang	0	1	0			https://issues.jboss.org/browse/DROOLS-113
DROOLS-49	So what do you suggest to do ?	Charles Moulliard	0	0	1	"I just saw that (at least) some of the pull requests, still use the hack <manifestFile>META-INF/MANIFEST.MF</manifestFile>. Thanks to cvgaviao's pull request in December 2012, we don't use that any more. I presume you were working on an old clone and some of the issues might have already been fixed by cvgaviao's work. Please reopen the pulls requests: with proper styling (so the changes are clear) based on top of the <packaging>bundle</packaging> work of cvgaviao, instead of undoing that work. preferably isolated per issue ("bug: improving osgi handling" versus "new feature: osgi-fying more modules"), so non-controversial changes are merged early instead of being delayed with the controversial changes (making them even more complex to understand). I'll handle the "bug: improving osgi handling" pull requests."	Geoffrey De Smet	https://issues.jboss.org/browse/DROOLS-49
DROOLS-1385	Friends,Please comment your ideas	Mylammal AV	0	0	1	Just build the kjar by reading the source file from wherever it is located. E.g.: https://github.com/droolsjbpm/drools/blob/master/drools-compiler/src/test/java/org/drools/compiler/CommonTestMethodBase.java#L407	Edson Tirelli	https://issues.jboss.org/browse/DROOLS-1385
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 I...	Mario Fusco	0	0	1	I updated the description. Please let me know if it still needs clarification.	Andrew Collins	https://issues.jboss.org/browse/DROOLS-1008
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 ...	Justin Goldsmith		1				https://issues.jboss.org/browse/DROOLS-1008
DROOLS-651	The documentation uses the term _binding variable_ 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 sho...	Marc Dzaebel	0	0	1	caching would not be cheap, it would require an array of Objects for each variable in each tuple. So it would probably add another 25% to over all memory usage. This would be over kill when typically most people are just doing person.age, so the additional cache is just wasted space with no performance gains.	Mark Proctor	https://issues.jboss.org/browse/DROOLS-651

Requirement	Question	Asker	Not answered	Unsuccessfully answered	Successfully Answered	Answer	Answered by	URL
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 rec...	Matteo Mortari	1			We could in the future look at allowing optional cached bindings, but right now it's lower priority and would add in a lot of complexity both on the engine and the language.		https://issues.jboss.org/browse/DROOLS-166

JBTM

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Michael Musgrove	0	1	0			https://issues.jboss.org/browse/JBTM-2905
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!	Ondra Chaloupka	0	0	1	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?	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-2905
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?	Tom Jenkinson	0	0	1	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	Ondra Chaloupka	https://issues.jboss.org/browse/JBTM-2905
JBTM-2835	What does the error look like with your patch? What is the effect of the NPE?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-2835
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	Ondra Chaloupka	0	0	1	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.	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-2822
JBTM-2822	Sorry Ondra - I don't quite follow. Which branch does not have this commit in?	Tom Jenkinson	0	0	1	Sorry Tom Jenkinson, wrong wording. I mean 'if branch' in 'doPrepare' method that refers about heuristics which happened. I think there is missing handling of suppressed exception - it's not added to exception which will be thrown. My change could adjust it https://github.com/ochaloup/narayana/commit/cb2ab4252a6c2e1106751d28103daefea42e86f6	Ondra Chaloupka	https://issues.jboss.org/browse/JBTM-2822
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.	Tom Jenkinson	0	0	1	Ok, I see. I will do so.	Ondra Chaloupka	https://issues.jboss.org/browse/JBTM-2822

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-2733	Hi [~tomjenkinson], I was wondering on your comment https://issues.jboss.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.	Ondra Chaloupka	0	0	1	Sorry, you are right that you can OR them.	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-2733
JBTM-2686	Where is the link to the PR?	Michael Musgrove	1	0	0			https://issues.jboss.org/browse/JBTM-2686
JBTM-2660	Good work Amos, I wonder if we can upgrade the SPI <code>org.jboss.spec.javaee.resource:jboss-connector-api_1.5_spec:jar:1.0.0.Final</code> dependency to a later one that includes JTA 1.2?	Tom Jenkinson	0	0	1	yeah, it looks like we need to upgrade to <code>jboss-connector-api_1.7_spec:1.0.0.Final</code> which includes JTA 1.2	Amos Feng	https://issues.jboss.org/browse/JBTM-2660
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	Tom Jenkinson	0	0	1	"I've already done some code to try it. I've added a <code>MessageListenerProxy</code> for that. It has an impact on the <code>ConnectionProxy</code> class because it has to instantiate a <code>SessionProxy</code> even if the transaction has not started yet." & "Hi Tom, It's not related to Artemis at all but to the <code>org.jboss.narayana.jta.jms</code> package, and to this previous ticket : https://issues.jboss.org/browse/JBTM-2616 OK, I'll start a discussion on the forum."	Alexis Hassler	https://issues.jboss.org/browse/JBTM-2658
JBTM-2634	Can you ask for feedback on the forum now it seems complete.	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-2634
JBTM-2617	Feel free to discuss in the forum if you would like some more feedback - not sure which class you are referring to.	Tom Jenkinson	0	0	1	I've just pushed the fix for this. Instead of holding connection in the recovery helper forever, I've implemented <code>XAResource</code> interface for it. Recovery helper acts as a proxy for JMS <code>XAResource</code> during recovery. It then creates and closes connections when needed.	Gytis Trikleris	https://issues.jboss.org/browse/JBTM-2617
JBTM-2606	Can we get a link to your branch here please?	Tom Jenkinson	0	1	0	https://github.com/gytis/narayana/commit/781869d9c46ca994efe31f8acc773c4658e326e0		https://issues.jboss.org/browse/JBTM-2606
JBTM-2605	Is this done?	Tom Jenkinson	0	0	1	I think that it is. Unless there were other requirements for just Spring integration other than a quickstart.	Gytis Trikleris	https://issues.jboss.org/browse/JBTM-2605

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Tom Jenkinson	0	0	1	Tom Jenkinson WildFly ships with it's own Artemis (including the journal) and Narayana reuses it. One of the last changes that went in WF10-CR5 was the upgrade of Artemis and among other things the JounalImpl now has an extra feature, making it backwards incompatible. Sure the fix could be made on Artemis side, but having it this way is the right way to do it and allow us to be able to explore that feature in the future.	Luis Barreiro	https://issues.jboss.org/browse/JBTM-2599
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	Tom Jenkinson	1	0	0			https://issues.jboss.org/browse/JBTM-2599
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.	Mark Little	0	0	1	It was added for JBTM-2168 within commit https://github.com/jbosstm/narayana/commit/e0cd867ad9763c511a7878a2faddc46d5f87898	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-2593
JBTM-2564	So do I need to add those to all the quickstart poms?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-2564
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? [~hvhovsepy] - 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?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-2564
JBTM-2564	Hi [~hvhovsepy] - 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.	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-2564
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.	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-2564
JBTM-2564	Did you need any fixes to Narayana to make it work on IPv6 because I think [~hvhovsepy] 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.	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-2564
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?	Hayk Hovsepyan	0	1	0			https://issues.jboss.org/browse/JBTM-2564

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-2553	Any more details on this?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-2553
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 be 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/adaptor/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 ?	Amos Feng	0	0	1	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.	Andrew Dinn	https://issues.iboss.org/browse/JBTM-2553
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.	Andrew Dinn	0	0	1	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.	Andrew Dinn	https://issues.iboss.org/browse/JBTM-2553

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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 ?	Amos Feng	0	0	1	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.	Andrew Dinn	https://issues.iboss.org/browse/JBTM-2553
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.	Andrew Dinn	0	0	1	Thanks adinn to look into it. I think we use the Jacoco 0.7.4.201502262128. I'm trying update to the 0.7.5.201505241946 but it still fails.	Amos Feng	https://issues.iboss.org/browse/JBTM-2553

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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!	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-2553
JBTM-2532	Did you put back a commit for the unit tests for now?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-2532
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?	Gytis Trikleris	1	0	0			https://issues.iboss.org/browse/JBTM-2522
JBTM-2496	Thanks Amos - what is missing to make it work with WildFly 10, please?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-2496
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 ?	Amos Feng	1	0	0			https://issues.iboss.org/browse/JBTM-2495
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	Tom Jenkinson	0	0	1	Hi Tom, I have created this task after agreement with Mike via IRC. He agreed to handle this ASAP. Thanks, Hayk	Hayk Hovsepyan	https://issues.iboss.org/browse/JBTM-2490
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?	Tom Jenkinson	0	0	1	Are we not planning to make the release completely automatic and incorporating it into our CI?	Gytis Trikleris	https://issues.iboss.org/browse/JBTM-2481
JBTM-2481	Are we not planning to make the release completely automatic and incorporating it into our CI?	Gytis Trikleris	0	1	0			https://issues.iboss.org/browse/JBTM-2481
JBTM-2464	Can this be closed?	Tom Jenkinson	0	0	1	Michael Musgrove made changes - 04/Aug/15 5:30 AM Status Pull Request Sent [10011] Resolved [5] Resolution Done [1]	Michael Musgrove	https://issues.iboss.org/browse/JBTM-2464
JBTM-2458	Is this done?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-2458

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-2458	What was the resolution?	Tom Jenkinson	0	0	1	It was implemented in the subtask	Gytis Trikleris	https://issues.iboss.org/browse/JBTM-2458
JBTM-2440	Do you know why these only seemed to start failing when we upgraded to JDK8?	Tom Jenkinson	0	0	1	"I not certain but diff'ing the config between now and the last successful build reveals that the tests were disabled NARAYANA_TESTS=0 on the successful build "&"I do recall - we disable the tests because byteman is known not to work with the IBM J9 compiler. So the fix is just to go back to not running the tests on J9 and figure out how to make byteman work with J9 as an optional task."&"I am marking the jira as minor with the fix to add a new profile that disables all of the byteman tests."	Michael Musgrove	https://issues.iboss.org/browse/JBTM-2440
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	Tom Jenkinson	0	0	1	The IBM manual says The Attach API can be enabled by setting the com.ibm.tools.attach.enable system property to the value yes; I tried this but tests marked up with the @BMUnitRunner annotation still fail with the error JVMTI_ERROR_NOT_AVAILABLE I tried enabling it in the pom via: systemPropertyVariables; systemProperties; and <argLine>-Dcom.ibm.tools.attach.enable=yes</argLine> none of which work. I am going to work around the issue using an ibm jdk profile that disables byteman tests.	Michael Musgrove	https://issues.iboss.org/browse/JBTM-2440
JBTM-2432	Is there anything in Narayana that requires Java8 though? If not then why not still allow that at build time?	Mark Little	0	1	0			https://issues.iboss.org/browse/JBTM-2432
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?	Mark Little	0	1	0			https://issues.iboss.org/browse/JBTM-2432
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	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-2429
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)	Michael Musgrove	0	1	0			https://issues.iboss.org/browse/JBTM-2414
JBTM-2412	Thanks Tom, Which component would that be. Tom	Tom Ross	1	0	0			https://issues.iboss.org/browse/JBTM-2412

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-2396
JBTM-2393	[~tomjenkinson], do we have analytics id etc.?	Gytis Trikleris	1	0	0			https://issues.iboss.org/browse/JBTM-2393
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 iirc? 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?	Tom Jenkinson	0	0	1	OK using horneq with (linux) asynchronous IO enabled and some other tweaks brings it to something more comparable (6 fold difference which when we factor in the extra http requests means a difference of 2-3 times), thanks: Benchmark Mode Samples Score Score error Units o.j.n.r.EmptyTxnTest.testEmptyTxn thrpt 1 1494.553 NaN ops/s o.j.n.r.NoTxnTest.testNoTxn thrpt 1 5041.441 NaN ops/s o.j.n.r.TxnTest.testTxn thrpt 1 215.502 NaN ops/s Note, I cannot try it out with the volatile store since RTS needs to list all uids for recovery purposes (and the volatile store doesn't support this op).	Michael Musgrove	https://issues.iboss.org/browse/JBTM-2390
JBTM-2379	Is it OK to raise this one now that 2377 is merged thanks?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-2379
JBTM-2379	Great, please can you let me know when the PR is merged so I can start the 5.1.1 release?	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-2379
JBTM-2375	I think a link to github is probably OK. Maybe do it by tag?	Tom Jenkinson	0	0	1	Gytis Trikleris made changes - 06/May/15 9:05 AM Status Open [1] Pull Request Sent [10011] Git Pull Request https://github.com/jbosstm/narayana.io/pull/2	Gytis Trikleris	https://issues.iboss.org/browse/JBTM-2375
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?	Tom Jenkinson	0	0	1	"When I click the first link I see empty boxes and arrows. Not really possible to comment or even compare and contrast with the original unless those boxes get names. " & "Looks like a drawing tool issue. I cannot see text in preview mode as well. I'm attaching the exported png file. "	Mark Little & Gytis Trikleris	https://issues.iboss.org/browse/JBTM-2366
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?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-2358

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Tom Jenkinson	0	0	1	Maybe we could support both default file names - ie we look for both before failing.	Michael Musgrove	https://issues.jboss.org/browse/JBTM-2357
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?	Michael Musgrove	1	0	0			https://issues.jboss.org/browse/JBTM-2345
JBTM-2343	How is this different to the current class that allows thread-to-transaction association and disassociation?	Mark Little	0	0	1	This is an interface that clients (such as JPA and II) implement and register with us (via <code>javax.transaction.transaction#registerSynchronization(Synchronization sync)</code>). We call them back when the thread-to-transaction association changes. So it is different. I have updated the forum thread so that we can track the API discussion. Perhaps I should have not have raised a separate JIRA for it (and subsumed it into JBTM-2342). The reason I chose a separate JIRA is because I though the interface belongs to the SPI (https://github.com/jbosstm/jboss-transaction-spi) whereas JBTM-2342 is for the implementation of the feature.	Michael Musgrove	https://issues.jboss.org/browse/JBTM-2343
JBTM-2334	Does that address the same use case though? If not then we should re-examine what we're doing.	Mark Little	1	0	0			https://issues.jboss.org/browse/JBTM-2334
JBTM-2318	Agreed. Maybe create a separate JIRA for the tooling?	Mark Little	1	0	0			https://issues.jboss.org/browse/JBTM-2318

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Tom Jenkinson	0	0	1	I think we should first remove the bits before each INFO. None of that information seems important and is duplicated (pretty much) by the text that follows the INFO block anyway. Then I like your idea of logging INFO (for what remains) only if the default ports are changed.	Mark Little	https://issues.iboss.org/browse/JBTM-2318
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. ARJUN...	Michael Musgrove	0	1	0			https://issues.iboss.org/browse/JBTM-2318
JBTM-2299	It looks like we don't need it as we has removed the TAO dependency in JBTM-1665 and used the http rest methods to access the TM [~tomjenkinson], Does it make sense ?	Amos Feng	0	0	1	That sounds right, it looks like a relic and should be pruned.	Tom Jenkinson	https://issues.iboss.org/browse/JBTM-2299
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?	Gytis Trikleris	0	1	0			https://issues.iboss.org/browse/JBTM-2296
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?	Gytis Trikleris	0	0	1	They looked good to me, thanks.	Gytis Trikleris	https://issues.iboss.org/browse/JBTM-2296
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?	Mark Little	0	0	1	Hi Mark, I clarified the description a bit. Its an addition, not a replacement of the existing standalone set. Its for ease of publishing on the site. Hope you agree it would be beneficial for that. Thanks, Tom	Tom Jenkinson	https://issues.iboss.org/browse/JBTM-2286

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-2280	Suggestimate suggests that this is similar to JBTM-2225, which appears to be true, no?	Stephane Epardaud	0	0	1	Well, I guess it's similar, but not quite the same since that other issue is about the CDI 'BeanManager', not about 'DataSource's.	Gavin King	https://issues.jboss.org/browse/JBTM-2280
JBTM-2280	[~mmusgrove]: would this be enough to get rid of JNDI for us entirely then?	Stephane Epardaud	0	0	1	I can look into providing an implementation using this approach. Our docs do say "This is not recommended, but provides a fallback for environments where use of JNDI is not feasible" but I do not know why it's not recommended - I will find out. But given that we have ceylon.transaction fully working with the JNDI approach why do we need to jettison its use?	Michael Musgrove	https://issues.jboss.org/browse/JBTM-2280
JBTM-2280	So this was closed and then fixed?	Stephane Epardaud	0	0	1	Hi Stephane, I added a feature over here JBTM-2325 it will be in the next version of Narayana - hope it works for you. I created a forum entry over here if you would like to discuss it further: https://developer.jboss.org/message/909177#909177 Thanks, Tom	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-2280
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());	Mark Little	0	0	1	I agree Mark. The way I had it was more like the ability to set a default and switch on whether it can be overridden. I think it is much clearer to just allow/suppress setting the single reference of a CAF via that property. I have sent a new PR so when that passes I will merge it in.	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-2279
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.	Michael Musgrove	0	0	1	Hi Mike, From my point of view the fixed version 6.0.0 is ok as we do not have any customer request for this functionality. The expose-all-logs functionality just came from fixes of https://bugzilla.redhat.com/show_bug.cgi?id=900289 . I put Gytis as assignee just because he was working on bz#1151322. Thank you Ondra	Ondra Chaloupka	https://issues.jboss.org/browse/JBTM-2270
JBTM-2264	Tom, great write up. I suggest changing this issue from a bug to something else (feature request?)	Mark Little	0	0	1	Tom Jenkinson made changes - 03/Oct/14 10:23 AM Issue Type Bug [1] Feature Request [2]	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-2264
JBTM-2262	BTW what did these scripts do in the old days? Did they add the jacorb jar explicitly?	Mark Little	0	0	1	Yes they always assumed JacORB, this script was modified during our move to ensure JDK orb was suitably high profile.	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-2262
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?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-2259

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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.	Hayk Hovsepian	0	1	0			https://issues.jboss.org/browse/JBTM-2245
JBTM-2231	Hi Jesper, Please can you give me some ideas of what this for? Is it the JI tracer? I have no problem with adding it, please just provide some context as there could be more I can give you. Thanks, Tom	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-2231
JBTM-2231	Hi Jesper, Is the attached PR what you were thinking? Thanks, Tom	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-2231
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.	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-2230
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.	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-2224
JBTM-2222	Forum discussion? Could be around, I just can't find it.	Mark Little	0	0	1	For discussion: https://community.jboss.org/message/881063#881063	Paul Robinson	https://issues.jboss.org/browse/JBTM-2222
JBTM-2213	I think this should be 5.next version?	Tom Jenkinson	0	0	1	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?	Gytis Trikleris	https://issues.jboss.org/browse/JBTM-2213
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?	Gytis Trikleris	0	0	1	Yeah, 5.next just stops me chasing around with the version number that might change. 5.later are issues that will be in 5 but not necessarily next version and can be brought forward to 5.next prior to release.	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-2213
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?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-2210
JBTM-2210	So all of these warnings have been there for 7 months?!	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-2210

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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.	Tom Jenkinson	0	0	1	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).	mark little	https://issues.iboss.org/browse/JBTM-2210
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).	Mark Little	0	0	1	The IDE warnings I usually can't make use of as they are full of "ArrayList is a raw type. References to generic type ArrayList<E> should be parameterized". I will look into whether I can suppress specific types of warnings as I doubt anyone has the appetite to fix those any time soon.	Tom Jenkinson	https://issues.iboss.org/browse/JBTM-2210
JBTM-2209	Hi Jesper, Please can you add a few more details to the Jira on the reasons for the change? Tom	Tom Jenkinson	0	1	0	As to checking the cloud build logs, this is really only done when a build fails.		https://issues.iboss.org/browse/JBTM-2209
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.	Mark Little	0	1	0	I just had a look but I can't see a way to fail javac just for when using deprecated APIs alone. Its worth pointing out that showDeprecation is off by default and we specifically enable it in the org.jboss.jboss-parent pom.		https://issues.iboss.org/browse/JBTM-2209
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?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-2206
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.	Tom Jenkinson	0	0	1	The perf profile looks for anything more than a 10% change (regression or improvement). The run did not show any "significant" change. The figures for each performance test were: com.hp.mwtests.ts.arjuna.performance.Performance1.onePhase: 895697.232857 vrs 902345.326032 (change: 0) com.hp.mwtests.ts.arjuna.atomicaction.CheckedActionTest.testThreadActionData: 437838.597325 vrs 430096.956384 (change: 0) com.hp.mwtests.ts.arjuna.atomicaction.CheckedActionTest.testCheckedAction: 32320.338816 vrs 33034.983201 (change: 0) com.arjuna.ats.jta.xa.performance.JTAStoreTests.jtaTest: 90531.981753 vrs 89787.577741 (change: 0) com.hp.mwtests.ts.arjuna.performance.Performance1.twoPhase: 817996.341431 vrs 790490.246540 (change: 0) where the first figure is the benchmark for the code being tested in the PR and the second figure is the figure for the code in the master branch	Michael Musgrove	https://issues.iboss.org/browse/JBTM-2206

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-2206	So no perceivable difference at all?	Mark Little	0	1	0			https://issues.iboss.org/browse/JBTM-2206
JBTM-2206	Can you include a link to the forum posting here for future reference?	Mark Little	0	0	1	I have added a link to a forum post where we discuss how we check PRs for performance improvements and regressions	Michael Musgrove	https://issues.iboss.org/browse/JBTM-2206
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	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-2203
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?	Tom Jenkinson	0	0	1	"See other related issue(s), but agree with Tom. Again, we can deprecate but that presents a maintenance issue. We need to understand what the cost/benefit trade-off looks like."&"Maybe delay creation of lock until first use? Might help the performance issue, whatever that is."	Mark Little	https://issues.iboss.org/browse/JBTM-2203
JBTM-2203	Maybe delay creation of lock until first use? Might help the performance issue, whatever that is.	Mark Little	0	1	0			https://issues.iboss.org/browse/JBTM-2203
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?	Mark Little	0	1	0			https://issues.iboss.org/browse/JBTM-2203
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?	Mark Little	0	1	0			https://issues.iboss.org/browse/JBTM-2202
JBTM-2202	What does this have to do with REST transactions? If it's local-only, then distribution over REST is not an option.	Paul Robinson	1	0	0			https://issues.iboss.org/browse/JBTM-2202
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.	Mark Little	0	1	0			https://issues.iboss.org/browse/JBTM-2189

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-2189
JBTM-2184	What's the procedure for ensuring our artifacts make it into maven central? How are you determining our missing dependencies?	Michael Musgrove	0	1	0			https://issues.jboss.org/browse/JBTM-2184
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.	Andrew Rubinger	0	1	0			https://issues.jboss.org/browse/JBTM-2184
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?	Michael Musgrove	0	1	0			https://issues.jboss.org/browse/JBTM-2184

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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: <pre>[tom@dhcppc1 quickstart](JBTM-2182) (15:40:29) \$ grep -n pom transaction-spi ./jca-and-tomcat/pom.xml ./jta-and-hibernate/pom.xml ./jca-and-hibernate/pom.xml</pre> 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...	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-2184
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.	Mark Little	0	1	0			https://issues.iboss.org/browse/JBTM-2183
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.	Tom Jenkinson	0	0	1	After consultation with Jonathan I have downgraded the priority to optional. The recommendation is make sure that the recovery credentials on a connection are configured correctly: <pre>set the recovery-username and recovery-password attributes to have sufficient privileges to access the oracle transaction tables</pre> if multiple datasources are configured against the same database make sure that only one of them is enabled for recovery (by setting the boolean attribute no-recovery="true false").	Michael Musgrove	https://issues.iboss.org/browse/JBTM-2183

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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	Michael Musgrove	0	0	1	"So I suggest stepping back from testing the collections within Narayana initially. It'd be good to look at what we use and determine what we could use as replacements. Then write some standalone tests that compare and contrast how these various collections perform independently. Once we have that data we should get together and review. Therefore, we're not looking at testing all of these collections implementations against each other, only a subset of what we think we need. "&"I'd also start with the native collections class (which is known to be suboptimal in a number of places) and the GS implementation."	Mark Little	https://issues.iboss.org/browse/JBTM-2180
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	Tom Jenkinson	0	0	1	Hi Tom Jenkinson, 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	Amos Feng	https://issues.iboss.org/browse/JBTM-2163
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	Amos Feng	0	0	1	Great, we should look at exposing some of the configuration options natively via http://www.osgi.org/javadoc/r4v42/org/osgi/service/cm/ManagedService.html	Tom Jenkinson	https://issues.iboss.org/browse/JBTM-2163
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.	Mark Little	0	1	0	For now, if you are sure the recovery service is running feel free to close this issue. We can open another one later to backport it if necessary.		https://issues.iboss.org/browse/JBTM-2159
JBTM-2151	Hi, The QS is merged is this "resolved"? Thanks, Tom	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-2151
JBTM-2151	Should I remove the fix version number as I don't believe this is work that is currently scheduled?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-2151

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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	Tom Jenkinson	1	0	0			https://issues.jboss.org/browse/JBTM-2119
JBTM-2116	Is this something that needs to be meaningful in a distributed environment? What about a heterogeneous environment?	Mark Little	1	0	0			https://issues.jboss.org/browse/JBTM-2116
JBTM-2116	Hi Jesper, it would be great if you could provide a couple of details on the intended usage of the request, please?	Tom Jenkinson	1	0	0			https://issues.jboss.org/browse/JBTM-2116
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?	Tom Jenkinson	1	0	0			https://issues.jboss.org/browse/JBTM-2116
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?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-2081
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	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-2081
JBTM-2081	Or if there's no property file default to JTA (or JTS)?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-2081
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?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-2081
JBTM-2081	What values are in the default properties file these days?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-2081

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Mark Little	0	0	1	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	Tom Jenkinson	https://issues.iboss.org/browse/JBTM-2081
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?	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-2079
JBTM-2040	For some reason I don't spot the attachment? It does seem a little extreme but could the level be dropped?	Tom Jenkinson	0	0	1	"I forgot to attach the log. It's attached now. "&"TRACE logging is needed by the Transaction Analyser to gather its data, so the level can't be changed. The Transaction Analyser could filter out these lines programmatically (and it does IIRC). However, this tool is just doing what a user would have to if they were diagnosing such issues without tooling support. It's also a pain for developers working on this tool, as they have to work a lot with the logs to automate the analysis process."	Paul Robinson	https://issues.iboss.org/browse/JBTM-2040
JBTM-1972	JIRA is not the appropriate place to ask this question. Please can you raise it via the forums at http://www.iboss.org/ejb3	Michael Musgrove	1	0	0			https://issues.iboss.org/browse/JBTM-1972
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?	Paul Robinson	1	0	0			https://issues.iboss.org/browse/JBTM-1955
JBTM-1955	I've read through the XTS docs. I've applied the minor changes (mostly typos) here: https://github.com/jbosstm/documentation/commit/238a06a9e50ff03144c7f542c08f1b36a9e6b93 . 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.	Paul Robinson	1	0	0			https://issues.iboss.org/browse/JBTM-1955
JBTM-1942	Is this issues resolved now? /cc [~paul.robinson]	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1942
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?	Tom Jenkinson	0	0	1	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?	Michael Musgrove	https://issues.iboss.org/browse/JBTM-1930

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Michael Musgrove	0	0	1	I agree that we shouldn't reword the xsd in text docs, but I just grepped what I think is the schema and I can't find that probe operation: https://github.com/wildfly/wildfly/blob/master/build/src/main/resources/docs/schema/jboss-as-txn_2_0.xsd (maybe its not called probe, the operation to resync the view of the object store) Also we should make sure our docs are still somewhat grepable: CLI, tools tooling should get a link to wherever the documentation for the CLI resides (as you suggest).	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-1930
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?	Tom Jenkinson	0	0	1	JBTM-1919 job is using my branch. I was waiting to make sure this worked before applying to all our arquillian.xml files. If someone else has time to do this before me, then you can see the change I made by looking at my branch, pulled in by the CI job.	Paul Robinson	https://issues.jboss.org/browse/JBTM-1925
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.	Paul Robinson	1	0	0			https://issues.jboss.org/browse/JBTM-1925
JBTM-1915	Is this still likely for tomorrow?	Tom Jenkinson	0	0	1	Paul Robinson made changes - 30/Sep/13 2:42 PM Status Reopened [4] Resolved [5] Resolution Done [1]	Paul Robinson	https://issues.jboss.org/browse/JBTM-1915
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?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-1909
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.	Mark Little	0	0	1	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?	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-1908
JBTM-1908	Did we upgrade the surefire version recently?	Mark Little	0	0	1	Hi Mark, It was upgraded from 2.12.4 to 2.16 as part of: https://issues.jboss.org/browse/JBTM-1875 Tom	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-1908
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?	Tom Jenkinson	0	0	1	Yes, that sounds like a good idea. Then we just need to look at the related issue I created and wait for the surefire team to (hopefully) address the enhancement request.	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-1908

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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 ?	Amos Feng	0	1	0			https://issues.iboss.org/browse/JBTM-1855
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?	Tom Jenkinson	0	0	1	Tom Jenkinson, I looked through the commit and it seems good to me.	Gytis Trikleris	https://issues.iboss.org/browse/JBTM-1855
JBTM-1855	Hi Amos, I think this is resolved? At least, the pull was merged? Tom	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1855
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.	Michael Musgrove	0	1	0			https://issues.iboss.org/browse/JBTM-1852
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}	Michael Musgrove	0	0	1	You first need to build ArjunaJTS submodules with idlj enabled: ./build.sh -f ArjunaJTS/pom.xml clean install -DskipTests -Didlj-enabled. Then to run that single test use: ./build.sh -f ArjunaJTS/jts/pom.xml test -Dtest=JavalDLRecoveryUnitTest -Didlj-enabled However, for me it only fails when I run tests like this: ./build.sh -f ArjunaJTS/jts/pom.xml -Didlj-enabled. But before running this you need to remove this line "<exclude>**/JavalDLRecoveryUnitTest.java</exclude>" from jts-idlj profile configuration in ArjunaJTS/jts/pom.xml.	Gytis Trikleris	https://issues.iboss.org/browse/JBTM-1816
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.	Tom Jenkinson	0	0	1	Once the remote tests are integrated then RCTest will test the success condition.	mark Little	https://issues.iboss.org/browse/JBTM-1816
JBTM-1809	Did this get resolved? I saw some commits on the 4.17 branch	Tom Jenkinson	0	0	1	changes pushed to both target branches	Michael Musgrove	https://issues.iboss.org/browse/JBTM-1809

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Paul Robinson	0	1	0			https://issues.iboss.org/browse/JBTM-1808
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?	Paul Robinson	0	1	0			https://issues.iboss.org/browse/JBTM-1808
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	Tom Jenkinson	0	0	1	Narayana: ArjunaCore arjunacore - The contents are the same as the individual jars.	Paul Gier	https://issues.iboss.org/browse/JBTM-1789
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	Tom Jenkinson	1	0	0	Narayana: ArjunaJTA narayana-jta - Adds jbossts-properties.xml		https://issues.iboss.org/browse/JBTM-1788
JBTM-1772	Is this not something that the build can automate?	Mark Little	0	1	0	Narayana: ArjunaJTS narayana-jts-jacorb - Adds jbossts-properties.xml and jbossts-idlj-properties.xml		https://issues.iboss.org/browse/JBTM-1772
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	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1770
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	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1754
JBTM-1754	Tom, I have a question: Shall a rename the prefix from "Narayana: " to "JBossTS" for branch 4.17?	Weinan Li	0	0	1	Hi Weinan, I think that would be best, if you didn't mind? Thanks, Tom	Tom Jenkinson	https://issues.iboss.org/browse/JBTM-1754
JBTM-1754	Hi Weinan, I think that would be best, if you didn't mind? Thanks, Tom	Tom Jenkinson	0	0	1	Not a problem at all Then I'll rename the prefix to "JBossTS: " for branch 4.17.	Weinan Li	https://issues.iboss.org/browse/JBTM-1754

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-1742	In what way did it break the build? I thought I had built everything when I changed this yesterday.	Mark Little	0	0	1	I put some comments on JBTM-1746 (which was raised when CI broke). I will "back port" those comments here.. Comment 1: Note that it is in the JDK orb IDL processing, rather than jacorb. A default build does not build the JDK orb but CI will which is where this issue was first discovered. You can reproduce this issue by: <pre>./build.sh clean install -f ArjunaJTS/idl/idlj/pom.xml</pre> Comment 2: The problem seems to be the following bit of the commit which breaks the JDK ORB idl processing: These statements were removed: <pre>#ifdef ORB_CURRENT_SUPPORT interface Current : CORBA::Current #else interface Current #endif</pre> And then replaced with just: <pre>interface Current : CORBA::Current</pre>	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-1742
JBTM-1721	Is there a pull request for this on WildFly or does it require M5?	Tom Jenkinson	1	0	0			https://issues.jboss.org/browse/JBTM-1721
JBTM-1707	Hi, Could you please confirm, can we use JBossTS 4.3.GA release for above mentioned problem. Thanks, Sube	Sube Singh	1	0	0			https://issues.jboss.org/browse/JBTM-1707
JBTM-1699	Should we link this to JBTM-1970 as well?	Mark Little	1	0	0			https://issues.jboss.org/browse/JBTM-1699
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	Tom Jenkinson	0	0	1	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?	Weinan Li	https://issues.jboss.org/browse/JBTM-1695
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?	Weinan Li	0	1	0			https://issues.jboss.org/browse/JBTM-1695
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.jboss.org/wiki/MavenDeployingaThirdpartyRelease	Paul Gier	1	0	0			https://issues.jboss.org/browse/JBTM-1694

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-1689	Which profiles do your team regularly use? I guess we should keep just the most commonly used ones?	Paul Gier	0	1	0			https://issues.iboss.org/browse/JBTM-1689
JBTM-1680	Please can you add an estimate for this?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1680
JBTM-1679	[~tomjenkinson], maybe this would be worth pushing to 5.next?	Gytis Trikleris	0	0	1	Sure, go for it.	Tom Jenkinson	https://issues.iboss.org/browse/JBTM-1679
JBTM-1679	How is this different to JBTM-1107? Should we close this as a duplicate?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1679
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?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1679
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?	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-1672
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?	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-1672
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	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-1664
JBTM-1658	Assign this to me, as I think I know what the problem is?	Mark Little	0	0	1	Tom Jenkinson made changes - 28/May/13 10:29 AM Assignee Michael Musgrove [mmusgrov] Mark Little [marklittle]	Tom Jenkinson	https://issues.iboss.org/browse/JBTM-1658

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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.	Paul Robinson	0	0	1	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 :)	Tom Jenkinson	https://issues.iboss.org/browse/JBTM-1645
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 :)	Tom Jenkinson	0	0	1	Yeah, I guess there will be situations like that. Maybe we just have one skip command per sub-job: !MAIN !XTS !QA_JTA !QA_JTS_JACORB !QA_JTS_JDKORB Less magic then.	Paul Robinson	https://issues.iboss.org/browse/JBTM-1645
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?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1645
JBTM-1643	Why focus solely on JTS? Why not support something like this for each component?	Mark Little	0	1	0			https://issues.iboss.org/browse/JBTM-1643
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	Tom Jenkinson	0	0	1	Hi Tom, yeah, there is no OR - so either the int or the string value would be perfect	Jesper Pedersen	https://issues.iboss.org/browse/JBTM-1635

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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	Tom Jenkinson	0	0	1	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?	Weinan Li	https://issues.jboss.org/browse/JBTM-1598
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?	Weinan Li	0	1	0			https://issues.jboss.org/browse/JBTM-1598
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	Tom Jenkinson	0	0	1	Cool patches! Besides setting up the components versions, I see the dependency management sections are cleaned up and the version tags are used instead. Thank you very much for help Tom	Weinan Li	https://issues.jboss.org/browse/JBTM-1598
JBTM-1591	ESB = Enterprise Service Bus?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-1591
JBTM-1584	Should this be done as part of the BLACKTIE-381 / BLACKTIE-202 package?	Tom Jenkinson	0	0	1	This JIRA is meant to replace the underlying transport used to implement XATMI calls whereas 381 is about starting transactions using JAB (ie it is independent). 202 is different from both of these.	Michael Musgrove	https://issues.jboss.org/browse/JBTM-1584
JBTM-1583	I think it would make sense for the implementer of this to also work on BLACKTIE-381?	Tom Jenkinson	1	0	0			https://issues.jboss.org/browse/JBTM-1583
JBTM-1583	How about using rest rather than SOAP?	Tom Jenkinson	1	0	0			https://issues.jboss.org/browse/JBTM-1583
JBTM-1577	Hi Mike, you means quickstart/blacktie/integration1 ?	Amos Feng	0	1	0			https://issues.jboss.org/browse/JBTM-1577
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 ?	Amos Feng	0	0	1	Hi Amos, Its actually a bit more complex than that. What I am trying to do is have most of the code for BlackTie in Java and just be able to call the users C code as a JNI method. This will eliminate the need for socket communication between the java and C code when it is on the same node. Also, I have updated the description as we no longer support .sar in AS7. Tom	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-1576

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-1576	Thanks Tom, 1. we need re-use MDBBlackTieService like following: {code} public class CPPXAtmiService extends MDBBlackTieService { native boolean tpService(byte[] data); public Response tpService(TPServiceInfo 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 nativeLibName = System.mapLibraryName("name of so"); System.load("/path/to/so" + nativeLibName); tpService(data); /* but how to get the response from jni ? */ } } {code} 2. the C code to implement tpService {code} JNIEXPORT jboolean JNICALL Java_CPPXAtmiService_tpService(...) { // 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	Amos Feng	0	0	1	Hi Amos, Broadly speaking that is right. In terms of code what you would try to do is re-use the Java blacktie library <code>jatmibroker-xatmi/src/main/java/org/jboss/narayana/blacktie/jatmibroker/xatmi/mdb/MDBBlackTieService</code> -> <code>jatmibroker-xatmi/src/main/java/org/jboss/narayana/blacktie/jatmibroker/xatmi/BlackTieService.java</code>. If you look in there, we call <code>tpService</code>. I expect that would become a JNI call. Then, in the C code, we would implement <code>tpService</code> to lookup the users actual service implementation and call there code. The pieces I haven't prototyped are when a users C service needs to call something like <code>tpcall</code> itself to call a second service or return data via <code>tpsend</code>. These would need to call back into the Java piece. To be clear, the C side would not use socket code at all. Hope that clarifies, Tom	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-1576
JBTM-1576	if we have this artifact for the C++ XATMI services, the <code>tpcall</code> in JBoss container don't need to send the message through the HornetQ and just only to load class and call <code>tpService</code>. This is your means *the benefits of performance (not requiring socket comms between java and C)* ?	Amos Feng	0	0	1	Hi Amos, > the user could compile his service with our c++ xatmi implementation and run as a standalone server (currently in CORBA Container and I'm working to remove the corba to use the raw socket), also could compile with jni xatmi implementation and deploy in JBossAS. 1 The user can use the same .so containing their service to deploy into an external C+ server or the JBossAS > if we have this artifact for the C++ XATMI services, the <code>tpcall</code> in JBoss container don't need to send the message through the HornetQ and just only to load class and call <code>tpService</code>. This is your means the benefits of performance (not requiring socket comms between java and C) This is true too. Tom	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-1576
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.	Michael Musgrove	0	1	0			https://issues.jboss.org/browse/JBTM-1576
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 <code>atmibroker-*</code> to parse this. I wonder how much more performant <code>ByteBuffers</code> are? Perhaps Amos as we move forward you can evaluate both and update us via the Jira?	Tom Jenkinson	0	0	1	yeah, no problem. I'll write a performance test after finishing Java socket server in BLACKTIE-399.	Amos Feng	https://issues.jboss.org/browse/JBTM-1576
JBTM-1576	Please can you add an estimate for this issue?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1576

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Michael Musgrove	0	0	1	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) - txx_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 txx_ttl time=0	Tom Jenkinson	https://issues.iboss.org/browse/JBTM-1573
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) - txx_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 txx_ttl time=0	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1573
JBTM-1568	Please can you add an estimate for this issue?	Tom Jenkinson	0	0	1	Amos Feng made changes - 14/Jan/13 11:28 PM Field Original Value New Value Original Estimate 4 hours [14400] Remaining Estimate 4 hours [14400]	Amos Feng	https://issues.iboss.org/browse/JBTM-1568
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	Amos Feng	0	0	1	Hi Amos, I think the ear-as-module feature is not available yet so I think we need to do this in two stages. pre-M2 (as in now), if you could merge the two into a single ear (i.e. the 4hr) that would be great. Then, leave the Jira open (or create a new one) to track the feature in upstream AS and then work on that part when it is available. Thanks!	Tom Jenkinson	https://issues.iboss.org/browse/JBTM-1568
JBTM-1568	I can not access ncl openvpn as my crt has been expired. can you mail me these txt ?	Amos Feng	0	0	1	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	Tom Jenkinson	https://issues.iboss.org/browse/JBTM-1568

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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	Tom Jenkinson	0	0	1	Thanks Tom, It looks like the upstream wildfly update to 8.0.0.Alpha4-SNAPSHOT and blacktie-wildfly is still Alpha3 as it does not been merged into the upstream. so the simple fix is to update the blacktie/pom.xml with parent version 8.0.0.Alpha4-SNAPSHOT Amos	Amos Feng	https://issues.jboss.org/browse/JBTM-1568
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.	Michael Musgrove	0	0	1	I've commit the fix and hope to work.	Amos Feng	https://issues.jboss.org/browse/JBTM-1568
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	Tom Jenkinson	0	0	1	resubmitted the pull req: https://github.com/jbosstm/narayana/pull/387	Crispin Boylan	https://issues.jboss.org/browse/JBTM-1565
JBTM-1558	Can you provide more detail to this task, since I thought we already had an SPI?	Mark Little	0	0	1	I have revised the description to apply more accurately to the statement of work	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-1558
JBTM-1558	There are a few SPI related issues at the moment. Maybe cross-link them?	Mark Little	1	0	0			https://issues.jboss.org/browse/JBTM-1558
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.	Mark Little	0	0	1	Regarding the question about what I meant by the "remote thread" case: <tomjenkinson> smarlow: AS1:EJB1 enlists JPA(XAR1) then calls AS2:EJB2 over JTS <tomjenkinson> smarlow: AS2:EJB2 enlists JPA(XAR2) this is done by a jacob request processor thread at random <tomjenkinson> smarlow: AS1:EJB1 calls commit on the transaction, JPA(XAR1) afterCompletion is handled in the same thread as before - this is fine - but then the call to AS2 to commit the transaction goes over a different jacob request processor, now afterCompletion in AS2 is done by a different thread The above doesn't raise any concurrency concerns since there is no (active) application thread on AS2 at the time that AS2 invokes the Synchronization.afterCompletion. Another case, would be similar to the above but with one difference. AS1:EJB1 enlists JPA(XAR1) then loops on calling AS2:EJB2 over JTS AS2:EJB2 enlists JPA(XAR2) and does processing and returns Looping continues until the transaction times out while current application execution is in AS2:EJB2 or just in AS1:EJB1. I don't think that only checking the reaper thread name covers when the transaction times out when the application thread of execution, is in AS2:EJB2. Question 1: Does the root coordinator (AS1) reaper always handle transaction timeout or can the subordinate	Scott Marlow	https://issues.jboss.org/browse/JBTM-1556

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Scott Marlow	0	1	0			https://issues.iboss.org/browse/JBTM-1556
JBTM-1526	Please can you itemize the dependencies that you are worried about - thanks	Tom Jenkinson	0	0	1	Hi Tom, sorry for the late delay. I've just come back from two weeks PTO I'm not so familiar with the jbossts usage in AS7, so I've done a grep and here is the result: ./org/jboss/as/xts/main/jboss-as-xts-7.1.1.Final.jar ./org/jboss/jts/integration/main/jbossjts-integration-4.16.2.Final.jar ./org/jboss/jts/main/jbossjts-4.16.2.Final.jar ./org/jboss/xts/main/jbosstxbridge-4.16.2.Final.jar ./org/jboss/xts/main/jbossxts-4.16.2.Final.jar ./org/jboss/xts/main/jbossxts-api-4.16.2.Final.jar Please correct me if the above list is not complete.	Weinan Li	https://issues.iboss.org/browse/JBTM-1526
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	Tom Jenkinson	0	0	1	Hi Tom, sorry for the late delay. I've just come back from two weeks PTO I'm not so familiar with the jbossts usage in AS7, so I've done a grep and here is the result: ./org/jboss/as/xts/main/jboss-as-xts-7.1.1.Final.jar ./org/jboss/jts/integration/main/jbossjts-integration-4.16.2.Final.jar ./org/jboss/jts/main/jbossjts-4.16.2.Final.jar ./org/jboss/xts/main/jbosstxbridge-4.16.2.Final.jar ./org/jboss/xts/main/jbossxts-4.16.2.Final.jar ./org/jboss/xts/main/jbossxts-api-4.16.2.Final.jar Please correct me if the above list is not complete.	Weinan Li	https://issues.iboss.org/browse/JBTM-1526
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	Tom Jenkinson	0	0	1	I'll use mvn dependency:tree to generate it and let's see what does these packages needed and whether they could be removed.	Weinan Li	https://issues.iboss.org/browse/JBTM-1526
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 [~mmusgrove] Please can you raise a pull request for the first one? Thanks! Tom	Tom Jenkinson	0	0	1	sure, I think the missing version tags is also an independent issue, I'll create a JIRA for that also. Let's make this issue be the parents of all the relevant issues I've found Now I need a profile that just contains the components that needed by AS7, and I'll use this profile to do the building against EAP6 environment and create jira issues as I make progress.	Weinan Li	https://issues.iboss.org/browse/JBTM-1526
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.	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1503
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"	Michael Musgrove	0	0	1	Mike, It's still relevant to XTS. It will become relevant to REST when we get REST-JDI. Therefore I've dropped the rest component.	Paul Robinson	https://issues.iboss.org/browse/JBTM-1498

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-1490	Elevated to critical, please do downgrade (optional?) if not required.	Tom Jenkinson	1	0	0			https://issues.jboss.org/browse/JBTM-1490
JBTM-1485	Does this change represent a possible backward incompatibility?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-1485
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?	Marc Richards	0	1	0			https://issues.jboss.org/browse/JBTM-1479
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	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1479
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?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1463
JBTM-1463	Please can you add an estimate for this?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1463
JBTM-1454	Why? Also, will you be testing HTTP versus IOP initially to get a base?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-1454
JBTM-1424	Please can you add a component for this?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1424
JBTM-1424	Shouldn't this also be fixed for 4.17?	Paul Robinson	1	0	0			https://issues.jboss.org/browse/JBTM-1424
JBTM-1423	Please can you add a component for this?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1423

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1407
JBTM-1398	Should org.jboss.as.weld.deployment.WeldAttachments be lifted into its own subsystem that doesn't depend on EJB3?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1398
JBTM-1398	I think this is resolved for now. Do you agree Tom?	Paul Robinson	0	1	0			https://issues.jboss.org/browse/JBTM-1398
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	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1382
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.	Michael Musgrove	0	0	1	"If I understand you correctly you are saying that the corresponding TEST*.txt are in the zip, if so that is fine, I looked at the pull that passed with your fix in and it didn't seem to be the case. " & "Yeah, that builds zip look fine to me, as far as I am concerned this can be marked as resolved - thanks"	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-1382
JBTM-1369	Can we have an update on this?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-1369
JBTM-1369	Can you paste in the commands you used to get those numbers, please?	Tom Jenkinson	0	0	1	Tom, I updated my previous comment to indicate how to run the commands manually	Michael Musgrove	https://issues.jboss.org/browse/JBTM-1369
JBTM-1369	What version of JacORB was used please?	Nick Cross	0	1	0			https://issues.jboss.org/browse/JBTM-1369

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-1369	What hardware architecture is your CI cluster?	Jason Greene	0	1	0			https://issues.jboss.org/browse/JBTM-1369
JBTM-1366	Just wondering if this should be in the Arquillian Transaction Extension Jira instance rather than here?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1366
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.	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-1359
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.	Mark Little	1	0	0			https://issues.jboss.org/browse/JBTM-1359
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?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1357
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.	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1340
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	Tom Jenkinson	1	0	0			https://issues.jboss.org/browse/JBTM-1340
JBTM-1339	Isn't this out of date now?	Mark Little	0	0	1	I think Mike is still working with Tomasz on this. I have assigned it to Mike, I am not sure why it was unassigned.	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-1339
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.	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-1339
JBTM-1338	If there is commitAsync, there should also be rollbackAsync?	Ales Justin	0	1	0			https://issues.jboss.org/browse/JBTM-1338
JBTM-1338	How do I get access to BaseTransaction in AS7? As JNDI lookup (java:jboss/UserTransaction) returns ServerVMClientUserTransaction.	Ales Justin	0	1	0			https://issues.jboss.org/browse/JBTM-1338

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-1328	Was this a one-off?	Tom Jenkinson	0	0	1	The failure is due to a transaction timing out unexpectedly. This is the only time this test has been seen to fail so closing as a spurious failure.	Michael Musgrove	https://issues.iboss.org/browse/JBTM-1328
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	Tom Jenkinson	0	0	1	Reopening to change the resolution status	Michael Musgrove	https://issues.iboss.org/browse/JBTM-1328
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?	Tom Jenkinson	0	0	1	Done (moved into the All tab)	Michael Musgrove	https://issues.iboss.org/browse/JBTM-1313
JBTM-1310	Does this affect all versions or 4.17.1+ in particular?	Tom Jenkinson	0	0	1	I have removed the affects versions field since it adds new functionality (and therefore affects all versions).	Michael Musgrove	https://issues.iboss.org/browse/JBTM-1310
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	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-1277
JBTM-1271	Hi Paul, this is a duplicate of JBTM-1072 (assigned to you) I think?	Tom Jenkinson	0	0	1	Paul Robinson made changes - 27/Sep/12 4:33 AM Field Original Value New Value Link This issue duplicates JBTM-1072 [JBTM-1072]	Paul Robinson	https://issues.iboss.org/browse/JBTM-1271
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.	Paul Robinson	1	0	0			https://issues.iboss.org/browse/JBTM-1266
JBTM-1254	Mike, Does that mean that you agree with me that we should resolve this as "won't fix"?	Paul Robinson	0	0	1	I don't think I said that. I agreed (in my first comment) that your rephrasing of the test is much clearer than what we currently have.	Michael Musgrove	https://issues.iboss.org/browse/JBTM-1254
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)?	Tom Jenkinson	0	0	1	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.	Paul Robinson	https://issues.iboss.org/browse/JBTM-1251

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Paul Robinson	0	1	0	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.		https://issues.jboss.org/browse/JBTM-1251
JBTM-1247	What is the use case for this?	Mark Little	0	1	0	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?		https://issues.jboss.org/browse/JBTM-1247
JBTM-1227	Did this make it into PropertiesFactory?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-1227
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?	Paul Robinson	0	0	1	Paul, the Hudson jobs failed due to the same failure again, so I retested it with your simple example against the latest AS7 on vmg36 and vmg35 and it passed. It seems that the simple example from arq.zip cannot reproduce the failure from narayana-java6-ipv6-WSTX11-interop and narayana-java6-ipv6-WSTFSC07-interop jobs." & "My guess is that the surefire VM is not being started with IPv6 settings. I ran it locally in supposed IPV6 mode and no -Djava.net.preferIPv4Stack=false -Djava.net.preferIPv6Addresses=true was on the command line to the spawned surefire process."	Ivo Studensky & Tom Jenkinson	https://issues.jboss.org/browse/JBTM-1217
JBTM-1212	Amos, Has this been done for the docs yet?	Paul Robinson	0	1	0			https://issues.jboss.org/browse/JBTM-1212
JBTM-1212	Can we mark this as resolved now?	Tom Jenkinson	1	0	0			https://issues.jboss.org/browse/JBTM-1212
JBTM-1208	Paul, I wonder that the contents of recovery records have only the text or the binary characters ?	Amos Feng	0	1	0			https://issues.jboss.org/browse/JBTM-1208
JBTM-1199	http://172.17.131.2/view/Narayana+BlackTie/job/narayana-java6/1237/ Seemed to start happenening for no obvious reason - I guess it is an underlying change in AS/arq?	Tom Jenkinson	1	0	0			https://issues.jboss.org/browse/JBTM-1199

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-1199	Amos, please can you investigate why https://github.com/jbossas/jboss-as/commit/435c862b9514bfc0a036bb2b916ac3dc63b4f4a 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?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1199
JBTM-1188	Where did the requirement for a per-tx override come from? Is the existing global config option not sufficient?	Jonathan Halliday	0	1	0			https://issues.jboss.org/browse/JBTM-1188
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?	Paul Robinson	0	1	0			https://issues.jboss.org/browse/JBTM-1187
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 ?	Amos Feng	1	0	0			https://issues.jboss.org/browse/JBTM-1187
JBTM-1184	Status update please?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-1184
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?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-1182
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.pro...	Ivo Studensky	1	0	0			https://issues.jboss.org/browse/JBTM-1173
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	Tom Jenkinson	0	0	1	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?	Paul Robinson	https://issues.jboss.org/browse/JBTM-1173
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?	Paul Robinson	0	0	1	Paul, the txbridge/pom.xml tests are passing on http://hudson.qa.jboss.com/hudson/job/narayana-java6-ipv6 so the bridge demo may be testable. As to the demo app, probably not (see my comment on https://issues.jboss.org/browse/JBTM-1156)	Michael Musgrove	https://issues.jboss.org/browse/JBTM-1173
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.	Paul Robinson	0	0	1	I don't think JBTM-1156 is being pushed to .Final is it?	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-1173

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1166
JBTM-1156	Paul, do you have any experience with setting up jobs on the QE cluster?	Tom Jenkinson	0	0	1	No I don't. We should ask Ivo if there are any limitations/quirks that we need to be aware of.	Paul Robinson	https://issues.iboss.org/browse/JBTM-1156
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	Tom Jenkinson	0	0	1	Current status: the job only fails the two standalone XTS arquillian tests. Note that I disabled the JDBC QA tests since it is not a supported feature in EAP6 (and because we only have oracle listening on an IPv6 endpoint).	Michael Musgrove	https://issues.iboss.org/browse/JBTM-1156
JBTM-1156	Great, are you working with Paul on resolving the XTS issues? +1 on the JDBC tests	Tom Jenkinson	0	0	1	The XTS failures also happen for simple (not transactional) JAX-WS services running in standalone (not XTS) configuration. I suspect it's either a bug in JBossWS, or a problem with our configuration of JBossAS when making it run in IPv6 mode. See forum posting: https://community.jboss.org/thread/200927	Paul Robinson	https://issues.iboss.org/browse/JBTM-1156
JBTM-1147	Please can you add an estimate for this issue?	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-1147
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?	Ondra Chaloupka	0	0	1	Come off it, the emma jdk6 binaries are fine. java goes to great lengths to ensure that binaries built with older jdks still run on newer ones. The error trace is not from emma code. It's from code that emma has instrumented. The changes emma is making to the arjuna code don't pass the newer verifier. http://www.oracle.com/technetwork/java/javase/jdk7-relnotes-418459.html It's not worth patching emma, which is largely defunct, to emit the more recent bytecode. Move to jacoco instead.	Jonathan Halliday	https://issues.iboss.org/browse/JBTM-1082
JBTM-1073	Worth a link to where this was done?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1073
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?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1072

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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" -u jbosstm-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?	Paul Robinson	0	0	1	"I am thinking that btny-pulls-blacktie should be able to be ran on any node so we could pin it to linux64, let me check... I think the quickstarts is on windows as that allows headless running for the XTS tests" & "re-pinned the blacktie pulls job: http://172.17.131.2/view/Narayana+BlackTie-pulls/job/btny-pulls-blacktie/30/console - it's running" & "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."	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-1072
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.	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1072
JBTM-1063	Hi Amos, Please can you backport this to the branch? Thanks! Tom	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-1063
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?	Paul Robinson	0	0	1	I need to change pom.xml jboss-as-parent dependency version to 7.1.2.SNAPSHOT and XTS/localjunit/pom.xml jboss-as-arquillian-container-managed dependency version	Paul Robinson	https://issues.jboss.org/browse/JBTM-1063
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	Tom Jenkinson	1	0	0	1. (Slightly unrelated) can we move the XTS demo to quickstarts/?		https://issues.jboss.org/browse/JBTM-1063
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	Amos Feng	0	1	0	Way ahead of you		https://issues.jboss.org/browse/JBTM-1063

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-1063	Amos, Can you give us an update on when this will be ported to the 4.16 branch?	Paul Robinson	1	0	0	JBTM-1056		https://issues.iboss.org/browse/JBTM-1063
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 ?	Amos Feng	0	1	0	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.		https://issues.iboss.org/browse/JBTM-1063
JBTM-1063	I think that was fixed by https://issues.iboss.org/browse/JBTM-1086 Have you done an 'svn up' on the 4.16 branch?	Paul Robinson	1	0	0	I see it being used for "hack/simplicity" (not sure which). However, this is referenced in the TXBridge documentation. I agree, I think we need to remove the dependency on the XTS demo. I think we just need to copy&simplify the Restaurant service into the TXBridge demo and update the docs.		https://issues.iboss.org/browse/JBTM-1063
JBTM-1063	I think that the branch version needs to hook in the txbridge tests?	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-1063
JBTM-1063	how about commands on the txbridge tests?	Amos Feng	1	0	0			https://issues.iboss.org/browse/JBTM-1063
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	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1061

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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.	Tom Jenkinson	0	0	1	Tom,	Paul Robinson	https://issues.iboss.org/browse/JBTM-1057
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".	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-1052
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?	Paul Robinson	0	1	0			https://issues.iboss.org/browse/JBTM-1051
JBTM-1051	PS, shouldn't this issue be moved to an AS7 Jira?	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-1051
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?	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-1051
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.	Paul Robinson	0	0	1	You are right, Paul. A minor change in configuration does not have to be identified by these tests. That is why I asked Rajesh for general testing of the configuration files. I am still waiting for response.	Ivo Studensky	https://issues.iboss.org/browse/JBTM-1051
JBTM-1051	Can this be moved to an AS7 Jira?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-1051
JBTM-1051	Sure, but can we re-open issues like this in AS7 project do you think?	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-1051
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?	Paul Robinson	1	0	0			https://issues.iboss.org/browse/JBTM-1031
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	Tom Jenkinson	0	0	1	Amos Feng made changes - 30/Mar/12 8:26 AM Link This issue depends BYTEMAN-190 [BYTEMAN-190]	Amos Feng	https://issues.iboss.org/browse/JBTM-1016
JBTM-1016	Please can you add an estimate for this issue?	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-1016

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-1016	Paul, can I delete XTS/localjunit/crash-recovery-tests/check.sh as now we can use maven plugin to check rule ?	Amos Feng	1	0	0			https://issues.iboss.org/browse/JBTM-1016
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?	Mark Little	0	1	0			https://issues.iboss.org/browse/JBTM-989
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?	Tom Jenkinson	0	0	1	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? There's nothing wrong with the AS7 style (IMO). I am the one with the problem as I am not familiar with that codestyle. I'm sure that over time I would get used to the AS7 style and no longer like my current style.	Paul Robinson	https://issues.iboss.org/browse/JBTM-989
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...	Paul Robinson	1	0	0			https://issues.iboss.org/browse/JBTM-989
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}	Paul Robinson	0	1	0			https://issues.iboss.org/browse/JBTM-968
JBTM-956	Hi Amos, is this closed?	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-956

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Paul Robinson	0	0	1	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.jboss.org/browse/JBTM-948?focusedCommentId=12661667&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-12661667	ivo Studensky	https://issues.jboss.org/browse/JBTM-948
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.jboss.org/browse/JBTM-948?focusedCommentId=12661667&page=com.atlassian.jira.plugin.system.issuetabpanels:comment-tabpanel#comment-12661667	Ivo Studensky	0	1	0			https://issues.jboss.org/browse/JBTM-948
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	Tom Jenkinson	0	0	1	Tom Jenkinson made changes - 06/Jun/12 4:41 AM Fix Version/s 5.0.0.Final [12318121] Fix Version/s 5.0.0.M2 [12316286]	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-940
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	Tom Jenkinson	0	0	1	Thanks to Tom for pointing out that we must use the version of orb.idl supplied with the compiler (it's allowed to be non portable). The other issue is that idlj (the compiler distributed with the standard JDK) does not support the local keyword. This keyword is used by ArjunaJTS/jts/idl/omg/XA.idl (interface Connector and CurrentConnection which we don't use, and if we did they would simply have to be treated as non-local). The svn commit for the pom changes removes the local keyword from these two interfaces. To use the idlj compiler build with the profile -P idlj The jacorb profile is active by default.	Michael Musgrove	https://issues.jboss.org/browse/JBTM-934

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-934
JBTM-934	Will it work on OpenJDK?	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-934
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?	Tom Jenkinson	0	0	1	1) No. Put it in the objectstore. 2) YAGNI. Don't make life difficult for yourself.	Tom Jenkinson	https://issues.iboss.org/browse/JBTM-933
JBTM-932	Yeah, I considered the tx reaper, but what about if the timeout is 0?	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-932
JBTM-928	[~tomjenkinson], these test still don't work. Shall I investigate further or keep them ignored?	Gytis Trikleris	0	1	0			https://issues.iboss.org/browse/JBTM-928
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	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-914
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?	David Lloyd	0	1	0			https://issues.iboss.org/browse/JBTM-909
JBTM-883	Does this work? Are you ready to create a pull request?	Paul Robinson	0	1	0			https://issues.iboss.org/browse/JBTM-883
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?	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-877

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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	Michael Musgrove	0	0	1	"XTS writes both coordinator and participant recovery records to the recovery store obtained by calling <code>StoreManager.getRecoveryStore()</code>. I don't believe changing the config to use an in-memory store raises any special difficulties as far as XTS is concerned – other than the fact that the coordinator and participant web service may not reside in the same JVM/on the same host (i.e. there may be more than one store in play which you might need to reconfigure)." " & "It's actually three stores rather than four, since the recovery store is not a separately configurable entity, it's just an internal api alias for obtaining the action store. The problems here stem from trying to retrofit a 'multiple instances of any given environment bean type' model onto a config system that previously assumed each was a singleton. The current semantic of the unnamed getters <code>xPropertyManager.getFooBean()</code> is 'get the instance called 'default', NOT 'get an abstract instance that contains the defaults for all other instances'. Indeed such defaults are not globally configurable at all - they are hardcoded. Thus it's not currently possible to e.g. change the store dir for all three stores with a single config option. Rather, you have to set it explicitly for each of the three store instances. This is by design and is not a bug. Ultimately the problems stem from not wanting to depend on a full DI framework, which limits the ability to inject named bean instances cleanly. The wider issue will need to be rethought as part of the 'removing statics / allowing dynamic reconfig / multiple cohabiting TMs' work. Meanwhile I'd	Andrew Dinn & Jonathan Halliday	https://issues.jboss.org/browse/JBTM-852
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	Tom Jenkinson	0	0	1	The resolution includes a test: <code>com.hp.mwtests.ts.arjuna.objectstore.AltStoreTest</code> and a doc update: https://github.com/jbosstm/documentation/blob/master/docs/development_guide/en-US/Configuration_Options.xml	Michael Musgrove	https://issues.jboss.org/browse/JBTM-852
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?	Mark Little	0	0	1	A static field can be thought of as a class-scoped variable, and thus a class loader-scoped variable. Given that there is presently a 1:1 mapping between transaction managers and class loaders, all static configuration and state today is by definition applicable per transaction manager. Therefore, if all currently static state is migrated to per-TM state, all present functionality should be preserved. From this perspective, adding new, global-to-all-TM state would be new functionality, and not doing so is preserving the status quo. And thus migrating currently per-TM state to be global state would be a significant behavior change. So I guess what I'm getting at is, I'd recommend that all the currently global state should be associated with the TM state, and not kept global, to preserve current behavior	David Lloyd	https://issues.jboss.org/browse/JBTM-846

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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)	Tom Jenkinson	0	0	1	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.	Mark Little	https://issues.iboss.org/browse/JBTM-846
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.	Mark Little	0	0	1	In regards to recovery - from a user perspective, I think that different people may be responsible/authorized for recovering transactions in the management system versus in the end-user application. That doesn't necessarily imply two recovery managers I guess, but I think that might be an important use case to address one way or another.	David Lloyd	https://issues.iboss.org/browse/JBTM-846
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.	Mark Little	0	0	1	Lets go for the "refocus this solely on the WF use case"	Tom Jenkinson	https://issues.iboss.org/browse/JBTM-846
JBTM-845	Isn't this duplicated by JBTM-1451?	Mark Little	0	1	0			https://issues.iboss.org/browse/JBTM-845
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 maven/pom-template.xml where they have a corresponding version number in the AS7 pom tree? Ta	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-829

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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 testsuite (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?	Paul Robinson	0	1	0			https://issues.jboss.org/browse/JBTM-817
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?	Paul Robinson	0	1	0			https://issues.jboss.org/browse/JBTM-817
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.	Misty Stanley-Jones	1	0	0			https://issues.jboss.org/browse/JBTM-813

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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.	Misty Stanley-Jones	1	0	0			https://issues.jboss.org/browse/JBTM-813
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.	Misty Stanley-Jones	1	0	0			https://issues.jboss.org/browse/JBTM-811
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?	Michael Musgrove	0	0	1	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.	Tom Jenkinson	https://issues.jboss.org/browse/JBTM-810
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.	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-810
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?	Mauro Molinari	0	1	0			https://issues.jboss.org/browse/JBTM-809

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-785	Hi Paul, Is there any update on this one? Is it covered by any QS? Tom	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-785
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	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-785
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?	Gytis Trikleris	0	1	0			https://issues.iboss.org/browse/JBTM-741
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>	Jonathan Halliday	0	0	1	Good point. I assumed they were running. So this issue then becomes: add them to the running suite of tests	Mark Little	https://issues.iboss.org/browse/JBTM-728
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}	Paul Robinson	1	0	0			https://issues.iboss.org/browse/JBTM-713
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	Mark Little	1	0	0	I'll move it to the next release version though. Feel free to move it back if you think you'll have time to get to it. In the meantime I'll continue with local tests.		https://issues.iboss.org/browse/JBTM-700
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	Tom Jenkinson	0	1	0			https://issues.iboss.org/browse/JBTM-659
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)?	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-659

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-659	Is this bug resolved in 4.6.1.CP13?	Tom Jenkinson	0	0	1	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.	Michael Musgrove	https://issues.jboss.org/browse/JBTM-659
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.	Michael Musgrove	0	1	0			https://issues.jboss.org/browse/JBTM-659
JBTM-659	Can we mark this as complete?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-659
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.	Jonathan Halliday	0	1	0			https://issues.jboss.org/browse/JBTM-636
JBTM-595	LogFactoryInterface yields a AbstractLogInterface - how does one get from LogI18n to AbstractLogInterface? I confess this logging infrastructure is pretty baffling to me	David Lloyd	0	1	0			https://issues.jboss.org/browse/JBTM-595
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.	David Lloyd	0	1	0			https://issues.jboss.org/browse/JBTM-595

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Jonathan Halliday	0	0	1	Hi, Jonathan, I'm pretty certain that something is amiss. In 4.2.2.GA I used to get the real causes of the exceptions from JBoss in my Swing client applications in case of constraint violations for instance, now in 5.1.0.GA I only get a <code>java.lang.reflect.UndeclaredThrowableException</code> ... Caused by: <code>javax.transaction.RollbackException: [com.arjuna.ats.internal.jta.transaction.arjunacore.commitwh enaborted]</code> ... You have to agree that this is now particularly helpful. I don't know which of the referred issues are fixed in TS 4.6.1 or not, but I'd like to ask you to verify this behavior.	Bozhidar Batsov	https://issues.jboss.org/browse/JBTM-575
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.throwPersisten ceException(AbstractEntityManagerImpl.java:576) at org.hibernate.ejb.AbstractEntityManagerImpl.throwPersisten ceException(AbstractEntityManagerImpl.java:614) at org.hibernate.ejb.AbstractEntityManagerImpl\$1.beforeCompl etion(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?	Jonathan Halliday	0	0	1	"I think so." & "thats a viable solution for us. We should be able to get the root cause of the problem and with your suggestion we could get it"&"Jonathan, Brilliant. As long as I can navigate the stack trace from the <code>javax.ejb.EJBTransactionRolledbackException</code> to the <code>MySQLIntegrityConstraintViolationException</code> (or at least the <code>ConstraintViolationException</code>) that would be awesome. I am happy to be the guinea pig and test this if you can work up a JAR for JBoss AS 5.1.0. Regards, Richard."	Bozhidar Batsov & Armin Haaf & Richard Kennard	https://issues.jboss.org/browse/JBTM-575
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.	Richard Kennard	0	1	0			https://issues.jboss.org/browse/JBTM-575

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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?	Bozhidar Batsov	1	0	0			https://issues.jboss.org/browse/JBTM-575
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.	Richard Kennard	0	1	0			https://issues.jboss.org/browse/JBTM-575
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>.	Misty Stanley-Jones	0	1	0			https://issues.jboss.org/browse/JBTM-575
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?	Misty Stanley-Jones	0	0	1	It's LogBrowser and LogEditor in the ArjunaCore package. If it was WONTFIX then the Resolution would be WONTFIX. Done means it has been completed.	Mark Little	https://issues.jboss.org/browse/JBTM-573
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?	Mauro Molinari	0	1	0			https://issues.jboss.org/browse/JBTM-535
JBTM-535	Could this be targetted for 4.7.0? Resolution should be quite straightforward... Thanks in advance.	Mauro Molinari	0	1	0			https://issues.jboss.org/browse/JBTM-535
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.	Mauro Molinari	0	1	0			https://issues.jboss.org/browse/JBTM-535

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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	Tom Jenkinson	1	0	0			https://issues.iboss.org/browse/JBTM-527
JBTM-503	Hi Andrew, any news on this? Could it be considered for 4.7.0?	Mauro Molinari	0	1	0			https://issues.iboss.org/browse/JBTM-503
JBTM-465	Please can you review this in context of JBTM-1230, I think the changes we made there might have addressed your requirements	Tom Jenkinson	0	0	1	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.	Mauro Molinari	https://issues.iboss.org/browse/JBTM-465
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.	Mauro Molinari	0	1	0			https://issues.iboss.org/browse/JBTM-465

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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	Tom Jenkinson	0	0	1	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.	Mauro Molinari	https://issues.jboss.org/browse/JBTM-465
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.	Mauro Molinari	0	1	0			https://issues.jboss.org/browse/JBTM-465
JBTM-437	Jonathan, can you take this?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-437
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.	John Mazzitelli	1	0	0			https://issues.jboss.org/browse/JBTM-435

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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.	Paul Robinson	0	0	1	The idea is that compensation transactions are supported in a local (in JVM) manner and also via ESB "transports", such as JMS, FTP, file, JDBC etc. Since this issue was created (4 years ago), we've added REST-TX now and I don't recall seeing much push for this within SOA-P, JBossESB or Switchyard, I think we could close the issue with the assumption that we've now done enough to support the original intention. But it's definitely not a duplicate of 1472.	Mark Little	https://issues.jboss.org/browse/JBTM-400
JBTM-389	Which svn tag are your src changes based on?	Jonathan Halliday	0	0	1	I didn't check it out from svn, just downloaded from JBoss Transactions Downloads version 4.2.3GA.	Dmitri Voronov	https://issues.jboss.org/browse/JBTM-389
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 ?	Jonathan Halliday	0	1	0			https://issues.jboss.org/browse/JBTM-369
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	Sacha Labourey	0	1	0			https://issues.jboss.org/browse/JBTM-369
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?	Jonathan Halliday	0	0	1	Hi Jonathan, actually I agree with Sacha, that every component should be integrated into the JBoss in a standard or JBoss manner: jar, sar, rar or -service.xml. That would simplify customization. Another advantage is that an application requiring any of these services is able to deploy its own customized version of the service.	Dmitri Voronov	https://issues.jboss.org/browse/JBTM-369

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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.	Andrew Dinn	0	1	0			https://issues.iboss.org/browse/JBTM-355
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 <code>TransportTimer.setTransportTimeout(long)</code> from your application code.	Andrew Dinn	0	0	1	i.) I am using Oracle Database 10g. i.) In stateless session bean I create entity manager (per session bean call) by entity manager factory. Then I persist entities (250 entities per one stateless session bean call). In finally block I ONLY close entity manager. There will no problem for me to flush, clear and finally close entity manager instead of only closing. You are correct, prepare phase is MUCH more faster (in milliseconds) with flushing done during the business process and not only at the end of the process (at the commit of the transaction) Thank you very much for your help, time and great description of the possible solutions. I think that this feature request can be rejected.	Pavel Kadlec	https://issues.iboss.org/browse/JBTM-355
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?	Jonathan Halliday	1	0	0			https://issues.iboss.org/browse/JBTM-345
JBTM-334	Just a clarification: this is not about deprecating what NORMAL does, just changing the name from NORMAL to PERIODIC. Correct?	Mark Little	1	0	0			https://issues.iboss.org/browse/JBTM-334

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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.	Jonathan Halliday	0	0	1	I don't disagree - I also think that the level ought to be INFO. But in this particular case, the type of message that is logged - a periodic success message - is more appropriate to TRACE in terms of the current logging "standard", as noted by Mark here - http://jira.jboss.com/jira/browse/JBTM-20#action_12346766 The guidelines I'd adhere to for JBossAS are: 1) An idle JBossAS should produce no logs (that's obviously the thrust of this issue) 2) Any logs that are generated on a per-request basis should be logged at level TRACE 3) Any one-time "debug" status messages (that are not repetitive in nature and do not grow in proportion to the number of requests) can be at DEBUG 4) INFO is OK for very low-volume, one-time messages like a greeting message "JBoss Foo version 2.0.0" or whatever	David Lloyd	https://issues.jboss.org/browse/JBTM-331
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?	Mark Little	1	0	0			https://issues.jboss.org/browse/JBTM-319
JBTM-319	Or the ConnectionFactory?	Mark Little	1	0	0			https://issues.jboss.org/browse/JBTM-319
JBTM-319	Can you take any queries and discussion to the forums, where it's easier to track? Thanks.	Mark Little	1	0	0			https://issues.jboss.org/browse/JBTM-319
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 coordinator (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?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-312
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 ... 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?	Toshiya Kobayashi	1	0	0			https://issues.jboss.org/browse/JBTM-300

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
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 '/lsvn/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	Daniel Bevenius	0	1	0			https://issues.jboss.org/browse/JBTM-280
JBTM-264	Is this resolved as part of the CLI work possibly? e.g. can the AS7 CLI be deployed externally?	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-264
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	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-264
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?	Andrew Dinn	0	0	1	Hi, sorry about not specifying which documents. The reason for this is that we have had this issue in the ESB documentation and I believe that the its the same template which was used with the transacations documentation. I've seen this in the AjunaCore Programmers Guide. I'll be going through the docs in the next week or so, and I'll let you know which other docs are affected. Thanks, Daniel	Daneil Bevenius	https://issues.jboss.org/browse/JBTM-252
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	Daniel Bevenius	0	1	0			https://issues.jboss.org/browse/JBTM-252
JBTM-250	What about JTX BaseTransaction.java?	Mark Little	0	0	1	Hmm, that will be the JTS version, eh? Ok, I'll patch that too.	Andrew Dinn	https://issues.jboss.org/browse/JBTM-250
JBTM-204	http://www.nabble.com/Spring---JBossTS-integration-t3901668.html You might want to mark this as Started ;-)?	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-204
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?	Jonathan Halliday	0	0	1	Yes please.	mark Little	https://issues.jboss.org/browse/JBTM-144

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-116	Which lower level? JBossTS understands -1. It means "transactions never time out".	Mark Little	0	0	1	"Okay, just looked at the code again and I missed a step. Setting the default timeout to 0 means that it defers to the default timeout in the property file. So -1 = don't call the reaper (no timeout) 0 = don't replace the property file value (i.e. use the property default) >0 = timeout So the code in the mbean should read public void setTransactionTimeout(int timeout) throws javax.transaction.SystemException { if (timeout != 0) { TxControl.setDefaultTimeout(timeout); } } Ignore 0 and set the others." & "BTW The lower level was the transaction reaper code "	Kevin Conner	https://issues.jboss.org/browse/JBTM-116
JBTM-103	Mark I'm not clear on this one. The effective default timeout appears to be 0 regardless of which TransactionManagerImpl is used. Am I missing something?	Jonathan Halliday	0	1	0			https://issues.jboss.org/browse/JBTM-103
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 }	Moises Quinto	0	1	0			https://issues.jboss.org/browse/JBTM-98
JBTM-95	We have such a thing in place now as far as I can see. Let's close this?	Mark Little	0	0	1	I made a similar proposal as part of the ESB project, we should adopt the same structure. tags/XXX Read only copies of each release, never updated branches/XXX Modifications (patches) to previous releases as a 'work in progress'. trunk Current onward development of the codebase workspace/XXX Playground for developers to try ideas not directly associated with any release/onward development. We should put this structure up on the wiki and then this can be closed.	Kevin Conner	https://issues.jboss.org/browse/JBTM-95
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.	Davide Cavestro	1	0	0			https://issues.jboss.org/browse/JBTM-86

Requirement	Question	Asker	Not Answered	Unsuccessfully Answered	Successfully Answered	Answer	Answered by	URL
JBTM-86	Hi Andrew, hi Jonathan, what is the fix version that includes the whole fix? 4.14? Thanks in advance!	Mauro Molinari	0	1	0			https://issues.jboss.org/browse/JBTM-86
JBTM-49	Does this still make sense? We should probably ensure our tooling is consistent with AS7 as much as anything.	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-49
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	Tom Jenkinson	0	1	0			https://issues.jboss.org/browse/JBTM-49
JBTM-46	FYI. I am going to mark TxManager as deprecated. See the linked task. Any reason why you didn't use links?	Adrian Brock	0	1	0			https://issues.jboss.org/browse/JBTM-46
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.	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-20
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 distribution and let them choose between the two implementations."	Mark Little	0	1	0			https://issues.jboss.org/browse/JBTM-5