

2013-12-02 to 2013-12-13 FF Stand-up Mtgs

•

Monday December 2nd

Tuesday December 3rd

Wednesday December 4th

- New Issues
- Status Reports
- Thursday December 5th
- Friday December 6th
- Monday December 9th
- Tuesday December 10th
- Wednesday December 11th
- Thursday December 12th

Monday December 2nd

Sprint Planning Meeting

Tuesday December 3rd

- [A. Soroka](#) got clarity, made a pull request that will be merged soon and should be able to proceed.
- [Andrew Woods](#) completed and sent out summary, is preparing for advisors' call where the one-click will be presented, working on mapping of use cases (not just candidate use cases, but road map). Will determine which features are without use cases. The new fedora 4 website is out there if anyone wants to provide feedback (). Would like to start user docs for release.
- [Chris Beer](#) spent more time struggling with clustering, is trying new approaches and produced three outstanding pull requests while waiting.
- [Eric James](#) is catching up with JMS messaging module work, waiting for index pluggable code to be merged to begin work to support solr. [A. Soroka](#) made an issue and passed it off.
- [Unknown User \(escowles@ucsd.edu\)](#) is looking at the JMS pull request and is running performance tests comparing minimal config file backed store with level DB store. Also testing different disk speeds. Attempting to make clustering work based on wiki instructions but has been unable to get it to work. Will make tickets for clustering and feedback features. Volunteered to help review PRs.
- [frank asse](#) updated benchtool, rewrote a large part of it to make it pluggable and added feature improvements. Had two questions, method to list all the objects in fedora, but isn't sure whether they're used anymore. Requested that [Unknown User \(escowles@ucsd.edu\)](#) could review the benchtool changes on a Pull Request. Will start on the fixity great.
- [Greg Jansen](#) created a project and POM that starts fedora 4 in cargo and runs benchtool as ant tasks (profiles and properties) for cluster setup. Afterwards will go back to other tickets.
 - <https://github.com/gregjan/fcrepo-test-profiles/blob/master/pom.xml>
- [Nigel Banks](#) working on documentation to convert fedora 3 islandora content models into fedora 4 CNDs (like book solution pack). This is almost complete.
- [Michael Durbin](#) write up meeting minutes and almost completed PR update for versioning (sparql update to add mixin type that autocreates a versioning policy property)

Wednesday December 4th

New Issues

- assign someone to fix build on linux (details in [IRC logs](#), test that fails is JQLConverterIT.testIt:206)

Status Reports

- [Unknown User \(escowles@ucsd.edu\)](#), [A. Soroka](#) and [Andrew Woods](#) couldn't make the meeting but provided details on what they were working on.
- [Greg Jansen](#) is working on a test starter for clusters, github project, profiles
- [Osman Din](#) was out yesterday but is working on inspecting modeshape output on the filesystem, not successful with tools, but is writing a script to analyze
- [Nigel Banks](#) doing content modeling documentation for islandora, doing some wiki gardening and cleaning up design - content modeling section of wiki and moving it to admin configuration. There are some people now at DGI who are doing human-level testing of tuque code which may be ready to be merged in coming weeks.
- [Michael Durbin](#) finishing up versioning tickets, rebased,
- [Eric James](#) updated code to support changes in interface. Will coordinate with Adam. Working on getting standalone solr to work.
- [frank asse](#) finished new benchtool work. Esme merged into master. Looking at fixity errors. Needs feedback about best solution.
- [Chris Beer](#) managed to get clustering started, but there's an issue with versioning. Outstanding pull requests are closed.
- [Benjamin Armintor](#) spending time addressing the biggest waste of cycles on ingest (enumerable node lookups and session creations). [A. Soroka](#) indicated you could register an event listener to listen to events of certain types which can prevent many of these node lookups.
- [Scott Prater](#) spent yesterday doing local work, discovered JQLFcrepoTransform project, is working on content model work. Will look into benchtool improvements.

Thursday December 5th

Fedora Committer's Call

Friday December 6th

- [Benjamin Armintor](#) working on events, needs a good idea to solve the issues.
- [Chris Beer](#) still working on clustering. Ran into some wierd transaction errors. May talk to infinispn people.
- [A. Soroka](#) continuing to bang around on indexer.
- [frank assege](#) working on fixity issue with projections. Has some questions, we'll talk on the hangout after standup. Subclassed filesystem projection to add a mixin that the fixity checker can query.
- [Eric James](#) looking into transforms that may be wired into Adam's indexer work. Not sure if errors he's encountering are unique to his setup or not.
- [Michael Durbin](#) will be working on HtmlUnit integration test failures.
- [Unknown User \(escowles@ucsd.edu\)](#) working on testing with large files. Is seeing really long times get properties as the datastream includes. Has meetings but will work to set up a 2 node cluster next.
- [Nigel Banks](#) working on single click tickets including OSX icon ticket which is causing problem, may take a different approach. Working on the restart issue, possibly a Spring issues.
- [Osman Din](#) is running frank's tool a single node fedora 3 in response to scott's request to re-run tests. Working on documentation for modeshape data layout.
- [Greg Jansen](#) handed off file testing to Esme, and is working on a number of authZ tickets to move basic policy enforcement to fcrepo-webapp. Moving auth-roles-common and auth-roles-basic in fcrepo-webapp.
- [Andrew Woods](#) cleaned up and fixed kitchen sink and pluggable projects. Working on some pull requests (batch updates).

Monday December 9th

- [Scott Prater](#) not present, but reported that he will work to finish up the AuthZ performance tests, finish the second half of the [content-model-to-node-types-page](#).
- [frank assege](#) not presented today or tomorrow, but reported that he worked on generalizing a CacheEntry interface to solve the fixity check on federated content issue.
- [Andrew Woods](#) reviewed some things, wanted to point out that Chris put in an update to batch operations API to create fedora:batch which replaces fedora:datastreams and works well to group API calls. Wants to create a top-level performance page (in user docs or admin docs) that summarizes in a very clear way the single node tests for file and leveldb and different operations. Also, create another set of pages (like admin docs) that have the features with a short description with a link off to design docs (or graduated design docs) with Nigel's UI walkthrough and Eric's configuration guide. Andrew would like to set the structure up and then have others finish out the pages. [Unknown User \(escowles@ucsd.edu\)](#) volunteered to take the lead on some or all of that.
- [Greg Jansen](#) has a few AuthZ tickets remaining, is looking to add an action form into out HTML rendering. Has some clustering testing to do.
- [Osman Din](#) working on documentation for storage layout, waiting for review and feedback. Also developed a tool to view the bson files. [Andrew Woods](#) suggested he update old test reports with his new tests.
- [Nigel Banks](#) wasn't present.
- [Unknown User \(escowles@ucsd.edu\)](#) finished up testing with ingest large files up to 1TB and access. Did a little bit of testing with federation that listing up to 1000 objects in a directory was fine, 10K files seemed to be a problem. Testing Greg's cluster project, but hasn't yet gotten the two node cluster to work yet. There was a question about using fedora with a cloud storage back-end. JCR doesn't support asynchronous access.
- [Michael Durbin](#)
- [Eric James](#) was working on FcrTransform to make sure it was working. Reviewing Adam's code. Would like to get indexing code to work on standalone solr (vs embedded solr).
- [A. Soroka](#) the pull request brings indexing aparatus up such that it does index to solr and triplestores, but it needs better test coverage and there are problems with the transform endpoint. Using anything other than the default transform creates a NPE. This needs to be addressed, and Adam will create a ticket. With that straightened out, it'll be ready for documentation and inclusion.
- [Chris Beer](#) performance testing and knocking off some tickets but is running short on low-hanging API improvements. Has the goal to beat fedora 3 in this stanford dataset by the end of the week.
- [Benjamin Armintor](#) uploaded a branch from which he was going to send a pull request, but is testing a rebase now. The profiling tests show this branch reduces overhead by a 6th or quarter. Made a compromise to improve messaging. For eventing: node removed events may come from another path. Perhaps we'll have a conversation immediately after standup.

Tuesday December 10th

- [frank assege](#) will be out today
- [Benjamin Armintor](#) not present today
- [Chris Beer](#) performance testing with new configurations, like asynchronous options. These look promising and have possibly 4-5fold increase on laptop, but much more modest gains on server. Will try increasing CPUs.
- [A. Soroka](#) is trying to marshall a bunch of pull requests (Ben's, Osmi's, Eric's).
- [Eric James](#) looking through Adam's code, has some other ideas about indexing but may not apply to this release. Revisit wiki pages to explain and document this approach.
- [Michael Durbin](#) documentation and such
- [Unknown User \(escowles@ucsd.edu\)](#) did some clustering work, got it working briefly, started writing it up and then it stopped working. Looking at performance and test results documentation.
- [Nigel Banks](#) not present today
- [Osman Din](#) spent time documenting modeshape artifact layout. Worked on Atom parsing problem. Will work with Adam to get PR reviewed. [Andrew Woods](#) asked why we still have messages with bodies. [A. Soroka](#) indicated that we always retained them to support a legacy use case. This affects Osman's PR and one from Ben. [A. Soroka](#) will try to merge the code without the Atom serialization.
- [Greg Jansen](#) is helping Scott set up authorization testing. Got a pull request in that adds access roles and basic policy enforcement point. Has another that adds access roles to http-webapp and is working on a PR for velocity templates. Is working on clustering tool as well.
- [Andrew Woods](#) is planning on spending the next hour going through pivotal tickets to merge and close low hanging fruit. Frank has a fixity bug in, others should look at it.

Wednesday December 11th

- [Andrew Woods](#) not much to report since yesterday, wanted to review Greg's AuthZ work. Has an open question about finished tickets... please let him know which (if any he should do first).
- [Chris Beer](#) doing more performance testing on Stanford data. Has asked to increase CPU's on those VMs. A couple outstanding bugs, one of which was merged in relating to namespace preferred URI predicates. Changed namespace back to info:fedora for translated URIs.
- [Eric James](#) working on documenting the external search. Not sure about Adam's PR. Would like to get it merged to move forward with standalone Solr instance.
 - The group confirmed that posting a CND is preferable to asking users to update their configuration.
- [Unknown User \(escowles@ucsd.edu\)](#) doing testing on file system federation. Working on documentation: created features page tasks, brought in pages to admin guide.
- [frank asseg](#) saw that a fix for Modeshape large file issue was pushed to master. Ran the benchmarks again, and the fix does indeed work. Responded to Stephano on tech list, but wondered why we have a fedora nodetypes endpoint and what restrictions apply. Frank also completed CacheEntry feature, but wants feedback before writing unit tests.
 - Its unclear where this documentation (limitations of node types) should live. [Andrew Woods](#) suggested a content modeling page should include this. This could live in an admin guide, but links should exist elsewhere (like properties CRUD).
- [Greg Jansen](#) is finishing up feature work with AuthZ, then can do cluster work and documentation.
- [Osman Din](#) worked on documentation for Level DB file store. Working on tool to analyze these files.
- [Scott Prater](#) resumed authZ testing yesterday. Is running performance tests now. Will move to content modeling documentation in the next days. [Andrew Woods](#) suggested we consolidate several of these pages on content modeling into a single place.
- [Michael Durbin](#) is still working on documentation
- [Benjamin Armintor](#) not present
- [Nigel Banks](#) not present
- [A. Soroka](#) out sick

Thursday December 12th

- [2013-12-12 - Fedora Committer Meeting](#)
- [Sprint Wrap-up](#)