

2013-11-21 - Fedora Committer Meeting

Time/Place

This meeting is a hybrid teleconference and IRC chat. Anyone is welcome to join...here's the info:

- Time: [11:00am Eastern Daylight Time US \(UTC-4\)](#)
- ReadyTalk:
 - U.S.A/Canada toll free: 866-740-1260, participant code: 2257295
 - International toll free:
 - <http://www.readytalk.com/intl>
 - Use the above link and input 2257295 and the country you are calling from to get your country's toll-free dial-in number
 - Once on the call, enter participant code 2257295
- IRC:
 - [Join the #duraspace-ff chat room via Freenode Web IRC](#) (enter a unique nick)
 - Or point your IRC client to #duraspace-ff on irc.freenode.net

Attendees

- [A. Soroka](#)
- [Ed Fugikawa](#)
- [Andrew Woods](#) ★
- [Benjamin Armintor](#)
- [Osman Din](#)
- [Eric James](#)
- [Michael Durbin](#)
- [Unknown User \(escowles@ucsd.edu\)](#)
- [Scott Prater](#)
- [Greg Jansen](#)

Agenda

1. Performance status and next-steps
2. 30 seconds on Fedora 3. Maybe 1 minute.
3. Versioning design questions and issues
 - a. fcr:versions endpoint
 - b. default behavior
 - c. scope for impending release
4. Admin search: expectations and current capabilities
5. Ingesting files larger than RAM (works in some Infinispan configs but not others)

Previous Actions

-

Status

[Osman](#)

[Esme](#)

- Updating [benchtool](#) to add support for performing read/update/delete operations, and allowing selecting random number generation impl. at runtime (some people report much faster performance with XORShiftRNG, I see much faster performance with `java.util.Random`).
- Re-testing large-scale ingest with updated benchtool, to be able to test read/update/delete operations on large-scale repo.
- Testing ingest of large files with different Infinispan backends to diagnose problems.

[Mike](#)

- Spent some time determining the best way to show that versioning works in fedora 4 and how users should be exposed to it in the pre-beta release.
- Making small fixes to enable demonstration of the feature
- Doing some preliminary design work on the future of this functionality

[Eric](#)

- began implementing file watcher to provide dynamic filesystem federation
- started looking at messaging module and solr indexing schemes
- code review and documentation

Frank

Adam

Chris

Ben

- Completed a chore to reorganize JSON configs to prevent build errors
- Ran benchtool (1000/1024/15) against F4 w/ YourKit profiler attached
- Started looking into <https://www.pivotaltracker.com/story/show/61119020> possible memory leaks

Scott

- Finished documenting some test profiles – still need to move cluster platform specs to platform profiles
- Read up on Modeshape, Infinispan
- More documentation tasks (About page), acceptance testing (versioning and search)
- AuthZ/no Authz performance testing (slow progress)

Greg

- Working on a content model examples project, with models verified via node types API
- Converted UNC preserved object model to a node type..
- UNC cluster is now ready for benchmarks, test platform documented
- Running minimal config test (single node) and setting up cluster config (4, then 8 nodes)

Ye

Andrew

- Fedora Sponsors meeting today (x2)
 - Will be soliciting module configuration management needs/expectations
- All fedora-dev-on-sourceforge emails have been migrated, finishing member migration
- Move "Fedora Futures" wiki-space this weekend?
- Received UI [feedback](#) from Robin Dean at Colorado Alliance of Research Libraries
- Fedora website redesign, should be released next week

Minutes

Performance status and next steps

1. Concern, having consistent approaches
2. What are people doing
 - Esme: working on benchtool, adds read/update/delete capabilities
 - Repeating 500-GB ingest in F3 and F4
 - Looking at testing profiles in wiki
 - Scott: Has created wiki testing workflows
 - Make sure benchtool readme is complete
 - Final product: big table with profiles
 - Need to surface results and issues as soon as possible
 - Ben: ran benchtool with YourKit listening
 - Looking into addressing memory-leaks
 - Will want to share profiling results, need a download location

Fedora 3

1. Bug report related to F3
 - There is a ticket
 - Very surprised the issue was not caught in release/integration testing

Versioning

1. Cleaning up implementation for the next release
2. Version endpoint currently returns RDF that is questionably useful
 - Would prefer to return an RDF list of versions
 - Need to turn what ModeShape provides into a useful response
3. Short-term goals
 - Can see a previous version
 - Can get a previous version - FedoraContent will not currently work with versions
 - Toggle versioning
 - Can rollback - later

4. Currently versioning is a performance hit
 - Need to note if F3 versioning is turned on in testing
 - The current F4 versioning is the "worst case scenario", maximum versioning
5. Future: need to enable default of versioning off...

Admin Search

1. Need to clarify current status
 - If properties are not defined in node-types, they will not be indexed
 - Default indexing dterms
2. Expectations
 - identifiers
 - enable indexing, but limiting access to certain search results

Ingesting files larger than RAM

1. "file" config enables ingesting content of size larger than RAM
 - Other configs fall down
2. Esme, Frank, and Ben working it

Status Review

1. Greg is working on mapping F3 to F4 content models
 - Scott is also interested in this
 - Greg to create ticket

New Actions

.