

# 2013-08-22 - 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](http://irc.freenode.net)

## Attendees

- [Andrew Woods](#)
- [A. Soroka](#)
- [Ed Fugikawa](#)
- [Jonathan Green](#)
- [Unknown User \(escowles@ucsd.edu\)](#)
- [Nigel Banks](#)
- [Osman Din](#)
- [Benjamin Armintor](#)

## Agenda

- Fedora 3.7 [Test Plan](#)
- Fedora 4 Current Work
  - [Sprint B2 Summary](#)
    - [Triplestore](#)
    - [Large files](#)
    - [Policy-driven storage](#)
  - [Async API strawman](#)

## Previous Actions

- [Scott Prater](#) to add example RI queries to [wiki](#)
- [Jonathan Green](#) to add example RI queries to [wiki](#)
- [A. Soroka](#) and [Michael Durbin](#) to add example RI queries to [wiki](#)
- [Greg Jansen](#) to add example RI queries to [wiki](#)
- [Unknown User \(escowles@ucsd.edu\)](#) to send email requesting RI query examples, RDF schema examples, etc be added to [wiki](#)
- [Andrew Woods](#) to investigate F4 RDF responses and if arbitrary properties are included

## Minutes

### Getting a Triple store implemented (Esme)

- Generic identifier to fixity nodes, now have public URI.
- Jena and sesame working.
- We have a SPARQL indexer working, including removes.
- Soon will be configurable rather than hard coded.
- Documentation on the way.

### Open ended questions/statements

- Blank nodes (where the RDF doesn't have a proper identifier), causes problems when deleting or selecting.
- Our API allows for the creation of blank nodes
- Fixity nodes cause blank nodes.
  - Added generic id to fixity nodes (addresses issue of blank nodes)
  - Fixity nodes will likely not be in the index anyways
- Attaching metadata directly to an object properties rather than data-stream?
- Should we disallow blank nodes?
  - Some standards require blank nodes to be constructed.
- Could use uuid to give them an identifier?
- If attached to object rather than data stream than blank nodes will be hard to avoid.

- Possible solution custom indexer class?

How is it working now? Anything remain

- Repository created/deleted/modified events via JMS.
- Indexer listens to messages and indexes them in either JENA/SESAME triple store.
- Deployment is not done but most likely will be a deployable WAR file.
- Currently a JAR
- To test now use the included integration tests.
- Maybe this afternoon, make it load some data from a fixture file?
- Action: Create ticket for adding configuration element of SPARQL endpoint to be exposed in repository:describe
- Action: Create a simple deployment wrapper for new message-consumer/indexer project

### Large file support federation and events (Nigel).

- SHA-1 is generated for every file
  - Takes a long time to compute checksum for large files
  - Up to an hour for ~60-gb files
- It would be nice if modeshape made the alg pluggable
  - We need to look into the ability to swap checksum alg
- Outcomes: Talk with Mode-shape about why SHA1 was chosen and if something else could suffice?
- Do documentation discussing how to enable/use the file-connector.
- Do write a how to guide about what one with large files should do.
- Will document process/config for of storing checksum in sidecar

### Wrap up work on policy driven storage (Osman)

- Backup currently configuration issues are causing problems
- Filesystem binary store working
- Post a question on Modeshape forum to clarify? Need to know if we can back up a repository if it's using a inifinispan binary store.
- Run time storage policy creation now works with the Rest API, and spring bean file.
- Wants to add more functionality, maybe next week.
  - Before expanding the policy rule definitions we will wait for specific use cases.

### Fedora 3.7 (Ben)

- Needs Windows Testing
- JUnit tests assumed the tests would be run in a predictable order (Not the case on Windows). Patch on the way.
- Java 6 tests should be done today.
- Go over manual tests, see if we can get someone from 3.6 to do 3.7 manual tests.
- Lots of tickets around performance. XML processing in particular.
- Do a 3.7.1 soon after? (To sort the performance issues).
- Does anyone run windows?

## Actions