

# 2013-11-18 to 11-29 FF Stand-up Mtgs

- Attendees
- General
- Minutes
  - Monday 2013-11-18
  - Tuesday 2013-11-19
  - Wednesday 2013-11-20
  - Thursday 2013-11-21
  - Friday 2013-11-22
  - Monday 2013-11-25
  - Tuesday 2013-11-26
  - Wednesday 2013-11-27
  - Thursday 2013-11-28 Meetings Cancelled (Thanksgiving in US)

## Attendees

- Unknown User (escowles@ucsd.edu)
- Michael Durbin
- frank asseg
- Eric James
- A. Soroka
- Osman Din
- Chris Beer
- Benjamin Armitor
- Andrew Woods
- Scott Prater
- Greg Jansen ★
- Ye Cao

## General

- Daily at 11am ET
- Call-in: Google-hangout at:
  - [https://plus.google.com/hangouts/\\_/calendar/eW91cm1IZGIhc2hIbGYuY29tXzVIYzdpNXQ2Z282dTdidHI4aTVrbGJxOTUwQGdyb3VwLmNhbGVuZGFyLmdvb2dsZS5jb20.opn3ai06k1ge0no877ei3cku4o](https://plus.google.com/hangouts/_/calendar/eW91cm1IZGIhc2hIbGYuY29tXzVIYzdpNXQ2Z282dTdidHI4aTVrbGJxOTUwQGdyb3VwLmNhbGVuZGFyLmdvb2dsZS5jb20.opn3ai06k1ge0no877ei3cku4o) (participant limit is 15)
- ReadyTalk (in case we reach hangout limit)
  - 866-740-1260, participant code: 2257295
  - internationally: <http://www.readytalk.com/intl>

## Minutes

### Monday 2013-11-18

[Sprint planning](#)

### Tuesday 2013-11-19

- **Esme**
  - Will create ticket for "First Steps" document
  - Getting UCSD test servers
- **Mike**
  - Has created tickets, are in backlog
  - Willing to off-load versioning tickets, if others are interested
- **Frank**
  - Working on AWS testing
  - May be able to get access from Yale's cluster
    - needs tickets
- **Eric**
  - Working on Solr integration
  - Has questions about approach for customizing index mapping
- **Adam**
  - Will be working RDF iteration
  - Created test profiles
- **Osman**
  - Working on AuthZ testing

- **Chris**
  - Clear-out API tickets
  - Moving onto F3 testing benchmarks
  - Will be making tickets
- **Ben**
  - Need to talk about CDI and scope
    - needs tickets
- **Scott**
  - Set-up latest version of F4
  - Set-up workstation for AuthZ tests
- **Greg**
  - ??
- **Ye**
  - Not on call
  - Via email
    - Today I planned the steps to finish solr task
    - <https://www.pivotaltracker.com/s/projects/684825/stories/57453388>
    - Solr indexer will read a config file for fields definitions, which includes standard node attributes, e.g., dateTime, mixinTypes, hasChild ...and custom fields, which user can define.
    - Create query samples that make use of standard fields and custom fields in
    - <https://wiki.duraspace.org/display/FF/SolrQueryRecipes>

## Wednesday 2013-11-20

- **Esme**
  - Performance testing on UCSD servers
    - Roughly the same performance as seen locally
  - Adding list/update/delete operations to benchtool
- **Frank**
  - Testing ingest of large files
  - Only able to add large files (> 1-gb) by using "file" configuration
- **Eric**
  - Working on JMS indexer
  - Interested in documenting modules
- **Adam**
  - Continuing with RDF iteration
- **Osman**
  - Finished testing AuthZ
  - Will be picking up backlog tickets
- **Chris**
  - Continuing with performance testing
    - Likely will finish F3 today
  - Will be looking into MODE-2098
- **Ben**
  - Started running YourKit, against benchtool run
    - Many exceptions were thrown
- **Scott**
  - Looking at platform profiles
  - Set-up a profile from UWisc production collection
  - Will be closing up workflow profile
  - Looking forward to AuthZ performance testing
- **Mike**
  - Not on call
- **Greg**
  - 3pm ET for documentation call
  - Mapping UNC content models to F4 node-definitions
  - BenchTool patch
- **Ye**
  - Not on call
  - Via email
    - The doc( public void update(final String pid, final String doc) {}) pass by JMS contains records of the node(with pid) and its parent and child(s). \*\*\* I wonder if it is complete or update only regarding parent/child node? These nodes will be re-indexed in solr in this case, so I hope information is complete in each JMS received.
    - I get error when run the latest fcrepo4 (commit 760b84277b0c75142aeea8325717dddc81451b9c)as below. I am not sure if it is just port conflict on my local machine?  
Failed to start Apache ActiveMQ (w306-60833-1384963117147-0:1, java.io.IOException: Transport Connector could not be registered in JMX: Failed to bind to server socket: tcp://0.0.0.0:61616?maximumConnections=1000&wireformat.maxFrameSize=104857600 due to: java.net.BindException: Address already in use: JVM\_Bind)  
WARN 16:58:37.664 (ResourceXmlApplicationContext) Exception thrown from LifecycleProcessor on context close

## Thursday 2013-11-21

[Committers meeting](#)

## Friday 2013-11-22

- **Esme Cowles**

- f3 impossibly fast - discovered he was just reading profile and not content
  - make f4 store in the right place
  - examples of other work flow scripts for tests
- **Michael Durbin**
  - finishing up versioning work
  - created better response, inc resolvable URIs
  - PRs today
- **Frank Asseg**
  - not much new stuff, worked on a SCAPE bug
  - cluster ingest w/large file
  - problems on physically distinct nodes, gets timeouts
  - will tweak configuration settings
- **Eric James**
  - not on call
- **A. Soroka**
  - PR for Ben on spring cleanup
  - RDF iteration - additional endpoints that were not initially high priority.
- **Osman Din** - not on call, sent status
  - continuing to experiment with different parsers to get rid of the intermittent AtomJMSIT errors thrown.
- **Chris Beer**
  - not on call
- **Benjamin Armintor**
  - hasn't seen any large file issues yet
  - Looking for error in mode/ispn, on store DS bigger than heap. (works in file store) Should not be an issue in our code.
  - downstream git project that runs ITs with large files
    - try to create a test to isolate issue against just mode or ispn (frank will help and try to get bug to MODE tomorrow)
  - When trying benchtool tests, getting error: "This operation cannot be performed when the session has pending changes"
- **Andrew Woods**
  - Had two sponsor calls yesterday (~20 participants total) with only complements on progress
  - Suggestion for creating screencasts, ala: [http://www.youtube.com/playlist?list=PLJFitFaE9AY\\_DDlhl3Kq\\_vFeX27F1yt6l](http://www.youtube.com/playlist?list=PLJFitFaE9AY_DDlhl3Kq_vFeX27F1yt6l)
  - Tech mailing list merge/transition should be complete this weekend (new: [fedora-tech@googlegroups.com](mailto:fedora-tech@googlegroups.com))
  - "Fedora Futures" wiki space is now "Fedora Repository", links not changed
  - Will be reviewing code and testing functionality today
- **Scott Prater** - not on call, sent status
  - finished up the test profiles
  - downloaded benchtool and working on AuthZ/no Authz test, running the "Trees" Setup profile with the four CRUD workflow profiles.
- **Greg Jansen**
  - working on tests and content model examples
  - trying a test harness based on cargo jetty run of fcrepo-webapp. has anyone tried this? (will consult Adam)
- **Ye Cao**
  - not on call

## Monday 2013-11-25

- **Esme**
  - Wednesday
  - Fixed jms-pluggable bug
  - Finishing performance testing
  - We be revisiting property updates in UI
- **Mike**
  - Wednesday
  - Will be working versioning tickets
- **Frank**
  - Working ISPN/LevelDB bug,
  - Issue with cluster ingest
  - Will be looking at FileSystemConnector this week
- **Eric**
  - gone Thurs
  - Not been work F4 much lately, local responsibilities
  - Thinking about property indexing mapping plan
  - plan meeting for search mapping
- **Adam**
  - gone Wed
  - PR RDF iteration
  - Created some clean-up tickets
- **Osman**
  - gone Thurs
  - Finished ATOM test failure ticket
- **Chris**
  - gone thurs
  - Performance testing
    - Heap issues
- **Ben**
  - gone thurs
  - Finished yourkit profiling

- **Scott**
  - around all week
  - Working AuthZ test
  - Needs help configuring test environment
  - Questions for CND/node validations
- **Greg**
  - Working content modeling wiki
  - Will be working AuthZ Osman

## Tuesday 2013-11-26

- **Esme Cowles**
  - indexer IT improvements, taking a more direct approach to verifying msg receipt
  - made existing test less brittle - still activemq port conflict on linux - almost there
  - posting test results -
  - Esme will start guiding folks on testing work flows (testing Czar)
- **Michael Durbin**
  - polishing versioning
    - HTML endpoint
    - versioning policies, beyond "always version"
    - you can turn off in CND by removing versionable mixin type
- **Frank Asseg**
  - talked w/Randall re: large files - accepted issue as blocker
    - <https://issues.jboss.org/browse/MODE-2103>
    - is more widely scoped than just the leveldb store
    - Randall has observed this in file system cache
  - will check if mode filesystem connector is read-write
- **Eric James**
  - pondering indexing of objects and DS
    - how to replicate f3-like indexing operations
    - working on JMS and text indexer
- **A. Soroka**
  - getting tests to run on indexer
- **Osman Din**
  - changing authZ IT - hardcoded values
  - working more tickets in backlog, inc ATOMJMS issues
- **Chris Beer**
  - tracked down performance test failures w/real world data set
    - blowing heap at 100,000 objects - observer session is too long lived, holding references
    - works with observer commented out
  - posting graphs - weird bumps/breaks in performance
- **Benjamin Armintor**
  - created tickets based on initial performance numbers and analysis
  - most of CPU and non-trivial memory for node lookup
  - observer for JMS involves node lookups - expensive - looking to reduce
- **Andrew Woods**
  - admin search works well, full text on properties and SPARQL queries
  - there is an IT - properties not defined in CND are also searchable
  - redeploying Fedora site - review of content/structure
- **Scott Prater**
  - setting up AuthZ tests based on Esme's work flow profiles
  - working on content models - how to use node types - how to perform validation
  - wiki gardening
- **Greg Jansen**
  - completed single node tests
  - running cluster tests today
  - have local tickets today
- **Ye Cao**

## Wednesday 2013-11-27

- **Unknown User ([escowles@ucsd.edu](mailto:escowles@ucsd.edu))**
  - worked on editable views, get all predicates working
  - haven't been able to get page to reload reliably - advice is to check ETAG is updated when props change
  - working on running r/w tests with Ben's branch
  - Action: will be creating new page to incrementally test PRs as they roll into master for performance improvements
- **Michael Durbin**
  - finishing versioning tickets - default policy is now to not version (in PR)
  - html template to browse/view fcr:versions
- **frank asseg**
  - success federation over 500 GB
  - first refresh will take 20 -30 secs, after that all is well
  - on subnodes as well, every first refresh will take a while
- **Eric James**
  - not on call, emailed status
  - working on the solr indexer
- **A. Soroka**

- not on call
- [Osman Din](#)
  - trying to make sense of binary data modeshape persists
    - will document findings
  - Will be submitting PR for <https://www.pivotaltracker.com/story/show/61152842>
- [Chris Beer](#)
  - performance testing
  - filling gaps with other tickets
  - CARL feedback - new tickets
- [Benjamin Armintor](#)
  - needs to respond to PR comments
  - try to find next ticket in performance - places we can avoid doing lookups of nodes
- [Andrew Woods](#)
  - reviewing PRs and tickets
  - pulling together user docs - from the UI level, features and what they mean
  - fedora-repository.org, migrating user mailing list
  - mapping use cases to key feature set for 4.0
- [Scott Prater](#)
  - saw a slow test results: ingest of 10,000 objects of 1mb - 9hrs
  - will try f3 comparison with f4, to establish a f3 baseline
  - will spec out administrator page about node types
    - creating CND, ingesting types
    - how to create object of a node type
- [Greg Jansen](#)
  - have the cluster running at UNC
  - working on cluster deployment and testing scripts - trying to make them generally useful
- [Ye Cao](#)
  - not on call

**Thursday 2013-11-28 Meetings Cancelled (Thanksgiving in US)**