

2017-12-14 - Fedora Tech 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)
- Dial-in Number: (712) 775-7035
 - Participant Code: 479307#
 - International numbers: [Conference Call Information](#)
 - Web Access: <https://www.freeconferencecallhd.com/wp-content/themes/responsive/flashphone/flash-phone.php>
- IRC:
 - [Join the #fcrepo chat room via Freenode Web IRC](#) (enter a unique nick)
 - Or point your IRC client to #fcrepo on irc.freenode.net

Attendees

- [Danny Bernstein](#)
- [Esmé Cowles](#)
- [Ben Pennell](#) ★
- [Michael Durbin](#)
- [Peter Eichman](#)
- [Joshua Westgard](#)
- [Bethany Seeger](#)
- (Please add your name)

Agenda

1. Announcements
 - a. 4.7.5 Release Process begins
2. 4.7.5 release - Planning for week of January, 15th 2018
 - a. Release manager - [Osman Din](#)
 - b. Volunteers
 - i. Testers
 1. [Danny Bernstein](#)
 2. [Joshua Westgard](#)
 3. [Peter Eichman](#)
 4. ?
 - ii. Someone to review 4.7.5 commit message for signs of missing documentation?
 - iii. Preparers of Module Release Candidates

Module	Who?
fcrepo-module-auth-rbac	Osman Din
fcrepo-module-auth-xacml	Osman Din
fcrepo-module-auth-webac	Osman Din
fcrepo-mint	Osman Din
fcrepo-audit	Osman Din
fcrepo-webapp-plus	Osman Din
fcrepo4-vagrant	Osman Din

- c. Resources:
 - i. component release process tracker: https://docs.google.com/spreadsheets/d/1I_zTMxh2l2rf2wpafoTwhSTR5GZuEoaTcZmTKCI3xT4/edit#gid=1769378986
 - ii. [Release Testing - 4.7.5](#)
3. [Fedora API Test Suite](#)... needing:
 - a. Try the tool against an API implementation
 - b. Code reviewing the tool... lots of low-hanging fruit
4. Simple, synchronous query in Fedora
 - a. What will it take to make this happen?

- b. [Prior art](#)
 - c. Queries to support
 - i. select ?s where {?s ?p ?o}
 - ii. select ?s where {?s <some-pred> ?o}
 - iii. select ?s where {?s <some-pred> <some-object>}
5. Tickets re
- a.
 - b.
 - c.
 - d. Unable to locate Jira server for this macro. It may be due to Application Link configuration.
6. 5.0.0 rele
- a. API Alignment
 - b. Pairtrees?
7. Beyond 5.0.0 - Areas of improvement
- a. Persistence?
 - b. Journaling?
 - c. Simple, synchronous query?
8. ...
9. Tickets In-Review

key	summary	type	created	updated	due	assignee	reporter	priority	status	resolution
Unable to locate Jira server for this macro. It may be due to Application Link configuration.										

Ticket Summaries

1. Please squash a bug!

key	summary	type	created	updated	due	assignee	reporter	priority	status	resolution
Unable to locate Jira server for this macro. It may be due to Application Link configuration.										

2. Tickets resolved this week:

key	summary	type	created	updated	due	assignee	reporter	priority	status	resolution
Unable to locate Jira server for this macro. It may be due to Application Link configuration.										

3. Tickets created this week:

key	summary	type	created	updated	due	assignee	reporter	priority	status	resolution
Unable to locate Jira server for this macro. It may be due to Application Link configuration.										

Minutes

2. Release

- Esme can help with Samvera testing
- Osman will perform builds of module release candidates, unless he wishes to delegate
- May be a challenging release given holidays

3. API test suite been running for a while. Has anyone had a chance to look at it?

- Danny hasn't looked at it yet, but seems easy to use. Going to try it out this week
- Interested in people trying it against other implementations, like Trellis and Cavendish.
- Bethany: Is this considered done as far as the contractors are concerned?
 - Danny: Good question, the repo doesn't say how complete it is. Will check in with Andrew
 - Since the API alignment isn't complete, it shouldn't pass against modeshape impl yet
 - Code reviewing the tool?

4. Simple querying in Fedora?

- Previously, cbeer had added this functionality, but it had been later removed
- Is data structured for this in modeshape to be reasonably performant?
- Mike: Was one of the agitators for this, opposed it being cut
 - Some stuff is inferred, some not directly searchable
 - What types of queries do we want to support?
 - What exceptions are we willing to tolerate?
 - Extension spec?
 - Is it okay if it doesn't work consistently on server managed triples, like date fields?
 - Just wants to be able to search dc:identifier. This would work, modeshape has an index that can be searched.
- Esme: Valkyie, making a list of queries that the repository needed.
 - Needed queries
 - all objects of given type
 - Doing a search for dc:identifiers
 - They will come up with a list of queries they need
- Danny: Would it be helpful at this point to fill out the list
 - Discuss some of the known limitations of modeshape's internal indices
 - Mike: Last modified date is across two fields. Might need to normalize way stored in fedora. Need to work out if this is needed
 - Esme: Types and containment triples are harder to make searchable
 - Search for non-server managed triples that are directly assigned are easy.
 - RDF type are not stored in the index modeshape maintains. That is inserted into responses.
 - Use case: find all objects of a type in order to do bulk object on it
 - Can't search on fcr namespace and ldp namespace. Might be okay to not support those, but it would be weird to have an LDP server that didn't support it
 - Could add support for this in after if there is demand for it
- Mike will start document to gather first pass at known limitations of implementation and requirements

5. Tickets requiring attention

- 2520
 - Bethany would like more feedback on what expectations are for mimetype
 - She will take another look at it to try to work through what the validation issue is
- 2650
 - Bethany will take a look
- 2544
 - there was a work around for that, using a different accept type. No one has strong feelings that it shouldn't be closed, so will make a note on ticket
 - Josh - as a larger strategy, this is something we will need to address
 - Paging mechanism is problematic in RDF rest api, but something we will need to deal with
 - Work around okay for now, but many members issue needs to be addressed in future implementations

6. 5.0.0 release?

- Need to wrap up creation of mementos, one of the last main things to bring into alignment with spec
- Pairtrees - Do we want to remove them?
 - Peter (?): In favor of removing them
 - Significant bit of internal work to hide them at fedora level, while still might need them in jcr
 - Danny: In doing away with pair trees, they would still be around internally
 - Peter: Need them, otherwise performance tanks after about 1000 children
 - Esme: Might want to look at Aaron Coburn's Appletree implementation. Takes checksum, makes path based on that. Includes hiding internal paths
 - Esme: would involve renaming everything in your repository, so it would need to take place as part of a major version change
 - Esme: Would either need migration tooling, or tooling for enabling/disabling the feature
 - Danny: how hard would a migration tool be to created?
 - Esme: Would be complex, but possible. If you have been using auto-generated UUIDS, could go through repo and remove pairtree.
 - Danny: interesting proposal, do we need community feedback?
 - Yes, more feedback would be good.
 - Esme: to write up brief description of proposal for fedora-tech
 - For discussion in new year

Action Items

Action: Check in with Andrew about completeness of the test suite

Action: Mike will put together a document with first pass at the feature set.

Action: Esme to write up brief description of proposal to remove pair trees for fedora tech