

2018-01-11 - 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.freeconferencelhd.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


- [Peter Eichman](#)
- [Aaron Birkland](#)
- [Esmé Cowles](#)
- [Andrew Woods](#)
- [David Wilcox](#)
- [Bethany Seeger](#)
- [Ben Pennell](#) ★
- [Daniel Lamb](#)
- [Jared Whiklo](#)

Agenda

1. 4.7.5 release - Planning for week of January, 15th 2018
2. Fedora 5.0.0
 - a. Release date target
 - b. PairTrees / AppleTrees
3. Tickets requiring attention

- a.  [FCREPO-2520](#) - Improve mimetype validation CLOSED - [Bethany Seeger](#)

- i. Any more discussion needed here? If a mimetype goes in, it should come out at the very least.
- ii. [Bethany Seeger](#) suggested an alternative implementation: <https://github.com/fcrepo4/fcrepo4/pull/1272#issuecomment-353173508> that would move the check closer to the HTTP layer
- iii. [Ralf Claussnitzer](#): Sounds reasonable. HTTP servers should respond with BAD REQUEST if given an unparseable mime type string. Should the repository layer check again?

b. T	Key	Summary	Assignee	Reporter	P	Status	Resolution	Created	Updated	Di
	 FCREPO O-2740	Float value not updated when current value rounds to the new one	Yinlin Chen	Kevin Ford	✓	IN REVIEW	Fixed	Mar 27, 2018	May 30, 2018	

[1 issue](#)

4. ...

Ticket Summaries

1. Please squash a bug!

Key	Summary	T	Created	Updated	Due	Assignee	Reporter	P	Status	Resolution
-----	---------	---	---------	---------	-----	----------	----------	---	--------	------------

FCREP O-2988	Creating versions of containers containing tombstones issue		Feb 08, 2019	Feb 13, 2019	Unassigned	Ben Pennell		RECEIVED	Unresolved
FCREP O-2841	Vulnerability on the commons-httpclient-3.1 dependency of fcrepo-integration-rdf		Aug 06, 2018	Aug 06, 2018	Unassigned	Peter Eichman		RECEIVED	Unresolved
FCREP O-2800	Issues running Fedora multi-threaded		May 14, 2018	May 14, 2018	Unassigned	Harsha Ummerpillai		RECEIVED	Unresolved
FCREP O-2790	Using a hash URI as a ldp:membership Resource doesn't work		May 04, 2018	May 04, 2018	Unassigned	Esmé Cowles		OPEN	Unresolved
FCREP O-2780	GET inbound refs incomplete response		Apr 26, 2018	Apr 30, 2018	Unassigned	Kevin Ford		OPEN	Unresolved
FCREP O-2746	Failed to reset the input stream when not accessing s3 directly or s3 compatible gateway		Apr 06, 2018	Apr 06, 2018	Unassigned	Thomas Schneider		OPEN	Unresolved
FCREP O-2739	Server returns an error after inserting a particular triple syntax		Mar 17, 2018	Mar 19, 2018	Unassigned	Stefano Cossu		OPEN	Unresolved
FCREP O-2712	GET with Prefer=InboundRefs on /fcr:metadata does not include InboundRefs		Mar 09, 2018	Mar 09, 2018	Unassigned	Kevin Ford		RECEIVED	Unresolved
FCREP O-2670	baseURI incorrect for audit JMS messages		Feb 02, 2018	Feb 05, 2018	Unassigned	Bethany Seeger		OPEN	Unresolved
FCREP O-2666	HTML UI does not update properly after binary resource reversion		Jan 23, 2018	Feb 01, 2018	Unassigned	Danny Bernstein		OPEN	Unresolved
FCREP O-2665	Integration tests flap on non-virtualized windows 10 machines.		Jan 23, 2018	Jan 24, 2018	Unassigned	Danny Bernstein		OPEN	Unresolved
FCREP O-2663	MOVE - Cannot delete Solr index		Jan 16, 2018	Jan 17, 2018	Unassigned	Ferdinando Terranova		OPEN	Unresolved

FCREP O-2662	Sparql Delete works inconsistently in transaction		Jan 03, 2018	Jan 04, 2018	Unassigned	Ben Pennell		OPEN	Unresolved
FCREP O-2659	Maven Enforcer Plugin `bannedDependencies` rule fails to resolve dependencies correctly		Dec 05, 2017	Nov 16, 2018	Ralf Claussnitzer	Ralf Claussnitzer		OPEN	Unresolved
FCREP O-2658	Invalid ebucore: hasMimeType value inhibits access to content		Dec 01, 2017	Dec 01, 2017	Unassigned	Kevin Ford		RECEIVED	Unresolved
FCREP O-2650	MD5s not being stored in transaction		Oct 24, 2017	Aug 02, 2018	Unassigned	Ben Pennell		OPEN	Unresolved
FCREP O-2648	Tomcat stuck threads during object ingest into Fedora with fcrepo-audit module enabled		Oct 19, 2017	Nov 29, 2017	Unassigned	Peter Eichman		OPEN	Unresolved
FCREP O-2635	binaries accepted by fcrepo did not arrive at final Modeshape storage location		Sep 29, 2017	Sep 29, 2017	Unassigned	Joshua Westgard		RECEIVED	Unresolved
FCREP O-2607	Incorrect ldp: contains reference on versioned resources		Sep 14, 2017	Sep 14, 2017	Unassigned	Ben Pennell		RECEIVED	Unresolved
FCREP O-2573	Namespaces not created correctly via PUT and POST		Aug 28, 2017	Aug 29, 2017	Unassigned	Bethany Seeger		RECEIVED	Unresolved

Showing 20 out of 46 issues

2. Tickets resolved this week:

Key	Summary	T	Created	Updated	Due	Assignee	Reporter	P	Status	Resolution
FCREP O-2993	Add Support for X-If-State-Token on PUTs		Apr 22, 2019	May 15, 2019		Danny Bernstein	Danny Bernstein		CLOSED	Fixed
FCREP O-2782	Adopt DuraSpace Checkstyle Configuration		Apr 28, 2018	May 16, 2019		Andrew Woods	Andrew Woods		CLOSED	Fixed

2 issues

3. Tickets created this week:

Key	Summary	T	Created	Updated	Due	Assignee	Reporter	P	Status	Resolution
-----	---------	---	---------	---------	-----	----------	----------	---	--------	------------

No issues found

Minutes

1. 4.7.5 release

- Branches have been pulled together, entering release candidate review period, targeting February 5 for release
- Release candidate is out, ready for release testing, looking for signups
- Reaching out for someone to do windows testing
- This should be a backwards compatible release, so it would be helpful if someone could drop the new war file on top of existing data.

2. 5.0 release

- Primarily focused on API alignment
- Conversations within committer group to work out what exactly is in scope for this release.

Release date

- Keeping in mind the calendar that the API specification itself is working around
- Back and forth in terms of implementations of spec and writing of spec.
- Want to have two implementations of the specification as part of releasing it. One is the modeshape 5.0 implementation.
- Intent is to finalize API by April, which would imply that there are implementations at that time.

What are we lacking?

- Memento
 - Jared - would like more people to look at this.
 - DannyB will be helping review this and get it over the hump
- Webac is fairly close.
- Page with all the test results from comparison: <https://wiki.duraspace.org/display/FEDORAAPI/Fedora+API+Spec+and+Delta+Document+Verification>
- Compatibility test suite - which is intended to be completed around April. Roughly a bit over halfway done, but needs to be verified.
- CRUD operations are pretty close, just a few things remaining

Looking for point people to verify alignment

- Memento - Jared and DannyB
- CRUD - Jared
- Webac - AaronB and Peter. There are testing scripts, including Authz tests.
- Fixity - Bethany
- Messaging - AaronB

Sprints?

Approach was to see how far we got before the holidays. Not a lot got done in the meantime.

Andrew: at a minimum, set up at least one sprint to finalize things

Would another earlier sprint be useful?

Jared: just pick a time, can probably get time allocated. It is an effective way to clean up small chunks of work, get testing done

Peter suggesting a february/march sprint is a possibility

Most people would prefer a sprint versus doing adhoc work

Aaron: With a date for a sprint he can at least ask.

Andrew: inclination is to do two sprints:

- First sprint: Finalize alignment of various components
- Second sprint: tie up loose ends prior to release

b. PairTrees / AppleTrees

Andrew: our current approach is not buying us anything in terms of performance (post creates interim nodes). The logic in fedora basically traverses over pair trees to find the children, so that eliminates benefits. Since the pair tree elements are there, you could write logic that doesn't automatically traverse

Peter: ended up going with requesting N-triples to get all the objects for a container, since the time-to-first byte was sufficient here to not cause timeouts. Didn't feel comfortable relying on pair trees as a retrieval/paging feature

Andrew: original intent was a partitioning mechanism, so that when making a request you wouldn't get more than 256 children to keep responses quick. Very quickly got feedback from people in the user interface that it was virtually impossible to find specific children when needing to drill down.

Esme: adding a Prefer header to manipulate that behavior would be nice. Effectively, header to turn off retrieving pair tree children

Andrew: haven't tried that yet, but is only a small piece of code to remove the pair tree children. What is the ideal behavior for 5.0?

Esme: going to make a wiki page to outline the three options:

- Keep configuration how it is
- Use Apple Trees instead
- Simplify config so it doesn't create pair trees, but smooths the path to creating pair or apple trees by the client

Bethany: Aaron C. mentioned blank nodes don't work well with Apple Trees, would want to check to make sure this is compatible

Andrew: Copy and move operations will not work with Apple Trees. Messages as they are currently emitted would not suppress the pair tree structure

Bethany: When 5.0 is put out there are ripple effects for the camel-toolbox since the messaging will have changed.

Aaron B: Does the apple/pair tree question have relevance to the import/export machinery that people might use to go from 4 to 5?

It sounds like we have one scenario where if we don't have pair trees in fcrepo5, would the intermediate nodes be 404s?

Is there a scenario where pair tree uris might be illegal in fcrepo5?

Esme: Thinks they would migrate over, but creating new uris might be different. Definitely are questions about how behavior would need to be different in fcrepo5 when migrating from 4.

Aaron: Along those lines, what about someone migrating from fcrepo4 to a different fedora impl?

Andrew: import/export tool should help with that sort of migration. Decision of 5.0 should be independent.

Aaron: impls without concepts of pair trees could be an issue.

Esme will put together that [wiki page to contain the discussion](#)

3. Tickets requiring attention:

b. Tickets in review

- 2656 - local-file uris. Needs some more review and the conversation on the mailing list is part of it
- 2636 - Can disable versioning of children through a modeshape. Andrews inclination is to close the ticket as no-fix.
- 2604 - Audit of messages being emitted, no need for a PR. Keeping it around for when doing the audit of messaging implementation. Leave in review and close when messaging is where we want it to be.
- 2591 - Changing interaction models for containers. Andrew is suggesting that current behavior stay the same, aka if you don't specify a model then you get a basic container. Cannot change the model after creation. Reopen ticket and assign it to API spec.
- 2544 - Been back and forth about time to first byte, Maryland has had the same issue. Can it be closed based on decisions Peter has made? Peter will take a look.
- 2520 - Reopen and assign to Bethany. Move check up to the http layer. Where did we get agreement from Ralph from? Going to circle back with him.
- 2459 - import of verisons. put on hold for more discussion
- 1786 - inclined to reopen and put in API alignment epic