

2017-09-07 - 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](#)
- [Aaron Birkland](#) (will be late)
- [Jared Whiklo](#)
- [Bethany Seeger](#) ★
- [Unknown User \(acoburn\)](#)
- [Christopher Johnson](#)
- [Peter Eichman](#)
- [Carrick Rogers](#)

Agenda

1. Upcoming "Alignment to Spec" sprint
2. Progress on: [Possible messaging memory leak?](#)
3. Need volunteer for: [FCREPO-2544](#) - Time to first byte of response is too long for containers with many contained resources CLOSED
 - a. Solution is outlined in ticket
 - b. <https://github.com/fcrepo4/fcrepo4/pull/1223>
4. [LTS policy proposal](#)
5. Wrapping up fcrepo-import-export-verify: <https://github.com/fcrepo4-labs/fcrepo-import-export-verify/pull/48>
6. [Report](#) on 4.7.4 Benchmarks and Testing ([Carrick Rogers](#))
7. Status of "in-flight" tickets

Ticket Summaries

1. Please squash a bug!

Key	Summary	T	Created	Updated	Due	Assignee	Reporter	P	Status	Resolution
FCREP O-3022	Intermittent build failure	🔴	Jul 15, 2019	Jul 16, 2019		Unassigned	Aaron Birkland	✔	OPEN	Unresolved
FCREP O-3005	GET and PUT on resource/fcr: metadata/fcr: acl and resource/fcr: versions/fcr: acl return 400	🔴	Jun 06, 2019	Jun 06, 2019		Danny Bernstein	Danny Bernstein	⬆	OPEN	Unresolved
FCREP O-3003	JSON-LD representation of rdf:Lists is incorrect	🔴	Jun 03, 2019	Jun 03, 2019		Unassigned	Jared Whiklo	✔	OPEN	Unresolved

FCREP O-2998	Potential issue with namespace registry		May 22, 2019	Jun 21, 2019	Bethany Seeger	Anna Dabrowski		OPEN	Unresolved
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 failes 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

Showing 20 out of 50 issues

2. Tickets resolved this week:

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

No issues found

3. Tickets created this week:

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

No issues found

4. Avalon R6 4.7.4 Testing



Minutes

Action Items

- Sprint starting next week -
 - Danny Bernstein will not be available for the first day of sprint next week
 - Danny Lamb switching to second sprint
 - Aaron Coburn asking about the consideration of current features that are necessarily API spec features? (like webac, versioning subtrees, pairtree)
 - responses from [Aaron Birkland](#):
 - There may be some changes that might break the existing behaviors of fedora that are not part of the API spec.
 - We do not have a policy that we'll be maintaining some of these hard to maintain behaviors
 - the consensus from last weeks call was that backwards compat is not maintainable
 - no concrete decisions about taking things out yet.
 - based on discussion at start of sprint, things may end up getting taken out.
 - some of the old behavior that's in the way or complicates things may be subject to removal
 - bug regarding [pairtree](#) - [Unknown User \(acoburn\)](#) encourages discussion around this at start of sprint (taking out pairtree behavior)
 - versioning - most concerning regarding implementing momento that's consistent with the spec - it may be the case that we have to change the way that versioning is reflected in modeshape.
 - There was discussion between [Aaron Birkland](#) and [Andrew Woods](#) about sprinters working on issues other than spec alignment. There is probably room to have folks working on other forms of fixes/improvements to fedora that aren't part of the API spec.
- Update on memory leak issue from [Peter Eichman](#)
 - modeshape bug - fixed in modeshape 5.4, he has been unable to reliably get the system to fail and hang in test case.
 - With avoiding transaction rollbacks they were able to run the remainder of the load and it completed successfully.
 - tested with batch loader? Not gotten back to that yet
 - ran into another issue which may or may not be related - there were some binaries where the metadata was there, but the binaries themselves were not in the system and may never have been.
 - binaries were probably large enough to not hit the 4k minimum size (repository.json config)
 - related to modeshape upgrade? no, this was happening in the modeshape 5.3 system causing the failed transactions he was seeing.
 - these files may never have actually gotten written to the file system (based on looking at backups)
 - next steps: might write files directly to the file system to replace them, though knows that's a bit dangerous
 - one file was a xml file - one jp2 (and a few other types) – maybe one of the directories where these files were sitting got deleted? It's two pages in a row that are missing.
- JIRA issue <https://jira.duraspace.org/browse/FCREPO-2544>
 - Do we use the sprint to remove this type of behavior?
 - If we are going to break things, lets do it now
 - Is this a reference implementation or just another impl of the spec? If reference then any out-of-spec features should be carefully considered. Good question for the fedora leaders group.
 - Why do we have pairtrees?
 - technique to control tree structure. JCR ism or specifically Modeshape-ism? 1000 direct children causes performance to slow quite a bit.
 - It's a mode shape issue – it has to do with Fedora's use of same-name siblings, which MODE specifically warns against using
 - it's how nodes are constructed and how we place properties on these nodes – you should be able to have a node with lots and lots of children, but the same name siblings causes issues.
 - ldp-contains uses this method getChildren() - causes this to be recursive and collects lots of nodes as a result and is slow. This is how fedora is using modeshape.
 - Any tickets around this? Yes, there have been a few.
 - Part of the root of the whole fedora code base - it's a huge issue. It's a limit that exists in fedora. Pairtrees is how fedora is trying to go around that limitation, but it's half-baked.
 - It's complicated

- The fix - exposing pairtrees even more as a logic concept that clients could use to control traversing repo. Would further ingrain them into certain client/server workflows.
- [Unknown User \(acoburn\)](#) has code that hides the pairtrees - might make sense to have something like that become part of main codebase?
- [Jared Whiklo](#): Do folks feel like we should hold off on this PR going into master and look at this some more?
- [Aaron Birkland](#) - what are the considerations for import/export (migrations in general) if we remove pair trees?
 - consideration of how the currently exposed web resources are resolved
 - [Danny Bernstein](#) - could imagine a way get beyond this issue via the import mechanism
 - with a little perl code run over an import/export dataset, that should be something one could easily deal with
- [Unknown User \(acoburn\)](#) This PR code is masking the problem - real problem is time to first byte – what is actually causing the buffer? That's what should be fixed.
 - Solving this problem is more impactful to the client
 - if this is time to first byte can we send chunks of data, so they get something quicker?
 - in principal the JAXRS framework responses should be streamed - we shouldn't have the collect everything and then send pattern, but it's what fedora currently does
- LTS - not going to discuss it now – read through and add comments to the page if you have any and we can dicuss it more next week
 - [Unknown User \(acoburn\)](#) - the JVM support policy will need to change in order to support an LTS policy
 - there was some talk about this at an earier tech meeting;
 - java 9 will be released in about 2-3 weeks - java 8 will reach EOL in Sept. 2018 - if the policy doesn't support EOL JVM's - we would need to change the policy if we want to still support that.
 - "the java version that was in currently in use when LTS version went into effect is the version that is targeted"
- Import / Export PR for the verify tool - can someone from the sprint take a look at this?
- From irc: [Aaron Birkland](#) - by the way, would somebody be willing to review/merge <https://jira.duraspace.org/browse/FCREPO-2578>. Our IT department flagged Fedora re: vulnerability with jquery 1.9.x
 - [Jared Whiklo](#) will take a look at jquery at some point
 - there was some discussion of getting rid of jquery on irc - someone should create a JIRA ticket for this work. (separate from issue above)
- Report on 4.7.4 Benchmarks and Testing ([Carrick Rogers](#))
 - avalon with active fedora
 - setting up test to simulate end user experience
 - anyone looking through the Avalon R6 4.7.4 Testing document (document is above) that notices some test missing - email him so it can potentially be added