

2016-09-08 - 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

- [Longshou Situ](#)
- [Esmé Cowles](#)
- [Yinlin Chen](#)
- [Bethany Seeger](#)
- [Aaron Birkland](#)
- [Jared Whiklo](#)
- [Jim Coble](#)
- [Andrew Woods](#)
- [Jennifer Lindner](#) ★
- [Michael Durbin](#)
- [Doron Shalvi](#)
- [A. Soroka](#)
- [Daniel Lamb](#)

Agenda

1. Updates and Announcements
 - a. [Archivematica Camp](#) - updates? outcomes?
 - b. Import/Export [sprint update](#)
 - c. [Bill Branan](#) and [Danny Bernstein](#) will likely be giving the "many members" issue a look
 - i. <https://github.com/fcrepo4/fcrepo4/commits/rwp-mode-queries>
 - ii. <https://github.com/fcrepo4/fcrepo4/pull/1086>
 - iii.  [FCREPO-2105](#) - Retrieving the description of an object with many reference properties is slow CLOSED
 - iv.  [FCREPO-2111](#) - Refactor URI translator chain to allow use without Jcr Node OPEN
2. Revisit default behavior on PUT requests containing RDF bodies:
 -  [FCREPO-1928](#) - PUT request should default to Prefer:handling=lenient CLOSED
 - a. For reference, [tamsin johnson's comment](#)
3. Fedora 4.7.0 planning
 - a. Depends on community verification of `/fcr:backup` and `/fcr:restore`
4. Import/Export: Clarify  [FCREPO-2164](#) - Add support for config file CLOSED
5. gradle in fcrepo-exts?
6. ...
7. Status of "in-flight" tickets

T	Key	Summary	Assignee	Reporter	P	Status	Resolution	Created	Updated	Due
	FCREP O-3698	Java client improperly encodes filename field	Ben Pennell	Ben Pennell		IN REVIEW	Fixed	Apr 15, 2021	Apr 15, 2021	
	FCREP O-3695	Single-request transactions should not use db transaction tables	Unassigned	Peter Winckles		IN REVIEW	Ready for Test	Apr 12, 2021	Apr 29, 2021	

	FCREP O-3692	Add metrics to migration-utils	Michael Ritter	Danny Bernstein		IN REVIEW	Fixed	Apr 08, 2021	Apr 29, 2021
	FCREP O-3691	Improve performance of search queries that return rdf types	Danny Bernstein	Danny Bernstein		IN REVIEW	Fixed	Apr 08, 2021	Apr 29, 2021
	FCREP O-3690	Add a configuration option for setting the S3 endpoint	Peter Winckles	Peter Winckles		IN REVIEW	Ready for Test	Apr 05, 2021	Apr 29, 2021
	FCREP O-3686	fcrepo-migration-validator: validate state in datastream	Michael Ritter	Danny Bernstein		IN REVIEW	Fixed	Mar 29, 2021	Apr 26, 2021
	FCREP O-3664	fcrepo-migration-validator: Add Datastream version validation for size	Michael Ritter	Danny Bernstein		IN REVIEW	Fixed	Mar 07, 2021	Mar 22, 2021
	FCREP O-3638	Create indices on search tables	Danny Bernstein	Danny Bernstein		IN PROGRESS	Unresolved	Feb 17, 2021	Apr 29, 2021
	FCREP O-3632	fcrepo-migration-validator: Use Immutable OCFLRepository	Unassigned	Danny Bernstein		IN PROGRESS	Unresolved	Feb 15, 2021	Feb 15, 2021
	FCREP O-3612	Test performance of DB rebuild optimization	Danny Bernstein	Danny Bernstein		IN REVIEW	Fixed	Jan 27, 2021	Feb 23, 2021
	FCREP O-3595	API Test Suite: 4.3.3.1-A	Danny Bernstein	Andrew Woods		REOPENED	Unresolved	Jan 11, 2021	Feb 10, 2021
	FCREP O-3593	API Test Suite: 4.1.2-B	Danny Bernstein	Andrew Woods		IN REVIEW	Fixed	Jan 11, 2021	Mar 11, 2021
	FCREP O-3561	Add commit() and rollback() support to DbSearchIndexImpl	Danny Bernstein	Danny Bernstein		IN REVIEW	Fixed	Dec 04, 2020	Apr 29, 2021
	FCREP O-3260	Document OCFL limitations related to paths and OS limitations	Peter Winckles	Andrew Woods		IN REVIEW	Ready for Test	Mar 26, 2020	Mar 03, 2021
	FCREP O-3194	Setup docker hub	Thomas Bernhart	Thomas Bernhart		IN PROGRESS	Unresolved	Jan 30, 2020	Apr 23, 2021

	FCREP O-3186	WebAC authorization gets wrong list of RDF types when checking access to binaries	Peter Eichman	Peter Eichman		IN PROGRESS	Unresolved	Jan 27, 2020	Feb 03, 2020
	FCREP O-2865	Camel Toolbox should not index mementos	Mohamed Mohideen Abdul Rasheed	Danny Bernstein		REOPENED	Unresolved	Sep 13, 2018	Apr 26, 2019
	FCREP 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
	FCREP O-2413	Adopters Guide: Versioning	Danny Bernstein	Andrew Woods		IN PROGRESS	Unresolved	Mar 08, 2017	Aug 08, 2018
	FCREP O-2299	Import - Export Videos	Youn Noh	Andrew Woods		REOPENED	Unresolved	Nov 14, 2016	Aug 14, 2017

Showing 20 out of 24 issues

Ticket Summaries

1. Please squash a bug!

T	Key	Summary	Assignee	Reporter	P	Status	Resolution	Created	Updated	Due
	FCREP O-3697	Multi-request transactions can be put in a bad state	Unassigned	Peter Winckles		OPEN	Unresolved	Apr 14, 2021	Apr 29, 2021	
	FCREP O-3678	Fedora 5 and SSL configuration not working	Unassigned	Jaime Penagos		OPEN	Unresolved	Mar 19, 2021	Mar 19, 2021	
	FCREP O-3655	application/link-format returns multiple entries for the same second	Unassigned	Jared Whiklo		OPEN	Unresolved	Feb 25, 2021	Mar 03, 2021	
	FCREP O-3635	Fedora 6 does not start on Jetty 11 or Tomcat 10	Unassigned	Peter Winckles		OPEN	Unresolved	Feb 16, 2021	Feb 19, 2021	
	FCREP O-3620	Fedora 5 -> 6 upgrade fails if a Fedora 5 export has retrieved the external binaries	Unassigned	Jared Whiklo		OPEN	Unresolved	Jan 29, 2021	Apr 29, 2021	
	FCREP O-3589	Membership /containment not returned as inbound references	Unassigned	Ben Pennell		OPEN	Unresolved	Jan 06, 2021	Jan 27, 2021	

	FCREP O-3585	Rollback with auto-versioning disabled can leave Fedora in an inconsistent state	Unassigned	Peter Winckles		OPEN	Unresolved	Dec 09, 2020	Feb 03, 2021
	FCREP O-3584	Resource locking does not take into account ghost nodes	Unassigned	Peter Winckles		OPEN	Unresolved	Dec 09, 2020	Feb 03, 2021
	FCREP O-3538	Wildcard searches and ids containing *	Unassigned	Peter Winckles		OPEN	Unresolved	Nov 19, 2020	Nov 19, 2020
	FCREP O-3521	Ensure all RDF formats supported in fcrepo-import-export also available in fcrep-upgrade-utils	Unassigned	Danny Bernstein		OPEN	Unresolved	Nov 06, 2020	Nov 30, 2020
	FCREP O-3506	OOM on GET of root with many child containers	Unassigned	Andrew Woods		OPEN	Unresolved	Oct 29, 2020	Nov 16, 2020
	FCREP O-3423	ACLs not displayed properly	Unassigned	Peter Winckles		OPEN	Unresolved	Sep 12, 2020	Nov 18, 2020
	FCREP O-3286	Exception on malformed Link header	Unassigned	Andrew Woods		OPEN	Unresolved	Apr 25, 2020	Nov 18, 2020
	FCREP O-3252	Disallow creating mementos within the same second	Unassigned	Peter Winckles		OPEN	Unresolved	Mar 17, 2020	Nov 18, 2020
	FCREP O-2263	Reserved Characters in fragments are encoded by UriBuilder	Unassigned	Christopher Johnson		OPEN	Unresolved	Oct 09, 2016	Nov 27, 2020

15 issues

2. Tickets resolved this week:

T	Key	Summary	Assignee	Reporter	P	Status	Resolution	Created	Updated	Due
No issues found										

3. Tickets created this week:

T	Key	Summary	Assignee	Reporter	P	Status	Resolution	Created	Updated	Due
No issues found										

Minutes

Agenda

1. Updates and Announcements

a. Archivemata camp

The first question was, is there any actual work that might be done with Fedora? And the answer is integrating at storage level is not good idea. Archivemata does lots of copying, moving bits, more like file system and not like a repository -- Fedora won't work well there.

However, how do you move stuff from Archivemata to Fedora? Let's let Nick and crew get sprints done and we'll have most of that done -- Archivemata produces bags and Fedora is working on accepting them.

The last is with the format policy registry. Lots of folks are interested in breaking it out of Archivemata, and amongst the people interested are [JISC](#), in the UK. They are preparing to manage an Archivemata service. So, let them get done before discussing breaking it out. (Context: many people who are using Fedora like Archivemata's format registry but may or may not want Archivemata itself, including Islandora -- the community would have interest). [Justin Simpson](#) promised to be in touch and restart conversation when they are ready.

In sum though, the vast majority of discussion was inward facing and around community development.

b. Update from import/export:

[Esmé Cowles](#): We've made good progress on export, and got basic import utility merged in, and are working on round trip for phase one objectives, w/o doing anything with bags. We're working on sample data sets and driving initial docs, on how different this utility is from RDF serializer Camel project that was the starting place for project.

We're in a good place for future sprints to tackle bags and other features. We'll have an executable jar that addresses objectives and docs for how to use utility, enough so people will be able to at least give it a try and provide feedback that could get rolled in to later sprints.

c. Two Java devs, [Bill Branam](#) and [Danny Bernstein](#), will have time to address Fedora performance issues around having many members. There's a Hydra change of direction (from members to collection) PR in, which is a mitigating change, but this will directly address this issue. (Status of the PR is the commits are not in but supporting commits are, Curation Concerns PRs are priority right now. There were some concerns, around whether solr indexing would change, but (AFAIK) concerns were addressed in it. PR should merged in next few weeks.

2. Revisit default behavior on PUT requests containing RDF bodies

The issue is a request for changing the interaction contract on PUT requests with Fedora. What we're asking for is changing the model from how it relates to server managed triples. In the current approach, if you want to do a PUT, replacing all triples, you can make a PUT with a complete representation of the resource, with both user and server managed triples in it. If you don't have server managed triples you get a 409 in response. The change being requested is for clients to only include user managed triples and if they include the server managed ones they get the 409.

[A. Soroka](#): For context -- what [Stefano Cossu](#)'s trying to do calls into question whether or not he should be using PUT at all, and he shouldn't be -- for these reasons. Here's his use case. Stefano doesn't care what the current state of resource is, he wants to mutate it (he may or may not want to remove properties) with no questions asked. He's only interested in the triples he actually wants to change. Here's why it matters: he's arriving with a request in hand, fully formed (adds or removals), -- and it's not a complete representation. It's a delta, or a future state of resource, however we want to describe it, and this is a direct conflict with what's a PUT. He's doing a PATCH. He's not doing it through a SPARQL update because it's unpleasant -- what he wants is a more convenient way to PATCH, like with RDF Patch or LDP Patch. Whether or not we act on this is up for discussion, but I feel we have to consider whether to do this with PUT.

[Andrew Woods](#): We need to verify that's correct description of Stefano's use case. And we need to reaffirm we're doing what we need to with PUT.

Aaron Birkland -- The only time I've come across Stefano's use case is in integration tests, where I don't care what the state of resource is. And that's kind of minor, so I understand the appeal of PUT for this kind of situation. That said, I think Fedora's current behavior is correct. The LDP spec is explicit on what clients should do with triples they don't care about (PATCH them through), but I can see desire for shortcuts in certain situations.

Adam: Stefano does care about the resource in the sense of he wants it to be maintained except for his changes, he just doesn't care what that state is. I respect his claim that it could be easier.

Andrew: We could just add a header -- is there another way to allow GET/PUT to work? -- It would require the client requesting minimal state -- we'd have to change the current get to support change in PUT.

Doing PUT w/o server managed triples, a user can do that with an additional header: `Prefer: handling=lenient; received="minimal"`. Both cases are handled in PUTs. You can pass in headers to get what you want. So, we'd just be losing one scenario possibly.

Does anyone know if there's a way to allow server managed triples?

There are references to it in two places in the LDP spec, our current interaction supports one scenario. And Stefano's ticket is supported in a may clause in the other.

Andrew: HTTP PUT should be idempotent and currently is not, because server managed properties change. (Later IRC discussion indicated that in fact current PUT behavior is at least arguably idempotent for a definition of that term derived by [Jared Whiklo](#) from the HTTP spec.)

Adam: How would we get around it? Stefano's request would.

Andrew: The spec says we should be able to make three identical PUTs and all three succeed and the state is the same regardless if you did it once or three times.

Adam, Andrew: The state of what, though? The server managed triples? Date modified would change .. but the state of resource is not the state of representation, so if change appears in representation but not resource is this a problem? HTTP is concerned only with representations. The representation has not changed, is idempotent but the resource has changed, date has, etc. Also concerned with does it succeed or not.

Aaron: In the example Adam gave before, with three same PUTs in minimal mode, you're modifying full representation three times. That's not what you'd expect with idempotent behavior, but comparing the two different representations of the resource. But where the current behavior is now (doesn't matter if they all succeed to meet criteria of idempotent), is clearly idempotent.

(a little mind bending)

Aaron: Idempotent is state, responses are not relevant here.

General agreement - we should come back to Stefano with we think Fedora's current implementation is correct, however, what are things you'd like to see to meet your needs? (Maybe an easier way to PATCH, etc.)

3. About Fedora 4.7 release

Andrew: I have two questions. One, is there any reason that we'd want to hold off on next release? The main thing it does is move Fedora to latest version of ModeShape, ModeShape 5 (which will help with Infinispan). And two, which is a blocker -- making sure mechanisms for migrating existing data work. Are there any other reasons we should hold off on 4.7?

[No response.] Seems to be no.

(general discussion) In fact we should release soon and let people who want to use ModeShape 5 use it more easily and will lead to better testing (results will be limited to mode 5 bugs).

Esme: I agree it's a blocker but we need someone with real production data to do it and find bugs.

[Michael Durbin](#): We won't find problems before, no one will do it. We'll have to ship and let people find problems and we'll fix them.

Andrew: Would it be useful to have a small utility for people to do migration?

[Bethany Seeger](#): Not sure if a script would be useful, but maybe just flesh out the fcr:backup/restore page a bit more?

Andrew: Considering a RC candidate would require a release manager -- who can do it? Adam?

Adam: I did 4.6 so I could avoid 4.7.

Andrew: We'll think about who can do it, and talk about how to make it happen.

4. import/export - Clarify?

Esme: No, I understand, no need to discuss.

5. Gradle in fcrepo-exts?

Bethany: Aaron Coburn and I've been talking about the possibility of [Gradle](#) in some projects in fcrepo-exts, instead of Maven. It does same thing as Maven, pckgs things. We want to see what the community thinks about some projects in fcrepo-exts using gradle.

general: Gradle uses Groovy language, Jared notes that [Unknown User \(acoburn\)](#) has seen that pax exam tests can be done but are more difficult.

Andrew: First we need to make sure Gradle does all that Maven does and that it won't impact fcrepo in some unknown way, by having some projects in fcrepo-exts using it.

Adam: Let's continue Gradle conversation next week, tell Aaron it received a thunderous 'maybe'.