

2016-05-26 - 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

- [Jared Whiklo](#)
- [Nick Ruest](#)
- [Bethany Seeger](#) ★
- [Longshou Situ](#)
- [Unknown User \(acoburn\)](#)
- [Yinlin Chen](#)
- [Esmé Cowles](#)
- [Michael Durbin](#)
- [Kevin Ford](#)
- [Peter Eichman](#)
- [Bill Branan](#)
- [David Wilcox](#)
- [Katherine Lynch](#)
- [Michael Gibney](#)
- [Emily Morten-Owens](#)
- [Andy Wagner](#)

Agenda

1. ModeShape5 update
 - a. Need assistance [enabling PostgreSQL](#) in Modeshape5 branch ([MySQL works](#))
 - b. Need assistance [testing/enabling clustering](#), which is [available with Modeshape5](#)
 - c. Pending ModeShape release 5.1.0 to include [bugfix](#)
2. Repository configuration (e.g. repository.json): should this be standardized in any way or do different Fedora implementations use whatever configuration style/content that makes sense for the implementation?
3. An `RdfSource` that is not a `NonRdfSourceDescription` currently describes itself. Should that be reflected in the Java API?
4. Logging [recommendations](#) and [style-guide](#)
5. Deprecation of Filesystem Federation – benefits, drawbacks, alternatives?
6. 4.6 release schedule
7. Removing deprecated (JCR-based) Import/Export endpoints – move code to -exts repository
8. Message serialization format (see [proposal](#))
9. Is there an LDP client available on which to base a new Fedora client?
10. Multitenancy
11. Status of "in-flight" tickets

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

1. Modeshape 5 update – revisit this agenda item one next week, when [Andrew Woods](#)
 - a. Aaron Coburn tested with the recommended currently released version of postgres driver and it worked fine. Esme will also try it out.
 - b. Enable clustering in modeshape 5 – no comments - talk to [Andrew Woods](#) for more input
 - c. Esme reports that there is a bug in modeshape 5 that prevents it from starting at certain points of the day. There is a bug where the time gets changed and the timestamps are then incorrect.
 - i. <https://issues.jboss.org/browse/MODE-2603>
2. repository.json - standardize it? [Unknown User \(acoburn\)](#) – repository.json file is currently passed right into modeshape. Do we want to put a fedora stamp on it and standardize it or just leave it up to the implementations. After some more thinking [Unknown User \(acoburn\)](#) thinks we should just leave it up to whatever the impl wants to do. [Esmé Cowles](#) agrees. If there is a use-case where there are files on a system, we may want to do something like this, but in general we should leave it up to the implementations.
3. [Unknown User \(acoburn\)](#) has a recent pull request - should the fedora kernel api have a new method called something like getDescription() and the implementing class would return the description - ie, a binary would do something different than a container class. This would make some of the code simpler in some case. PR that demonstrates need for this change: <https://github.com/fcrepo4/fcrepo4/pull/1040/files#diff-d372035f89c8d0846920e20bd13e0145R514>
 - a. talk more next week when [A. Soroka](#) is around
4. Style guide - involves a PR that [Unknown User \(acoburn\)](#) made. [A. Soroka](#) is suggesting we have some kind of policy or recommendation on how logging works.
 - a. came from – https://github.com/fcrepo4/fcrepo4/pull/1033#discussion_r63703110
 - b. general idea of not printing stack traces. [Esmé Cowles](#) thinks developing a style for this would be a nice thing to do.
 - c. general idea is yes - it's a good idea, but idea of what it is and how implemented would need to be thought through.
5. Deprecation of file system federation
 - a. [Esmé Cowles](#) and Mike Durbin - two who initially wrote code and had use cases for it. Because you can't edit it with REST API it's read only. You can use it to store files, but you can use it to store metadata usefully. Linking is the same as if things were on an external webserver. If you have files you want to serve from disk, then you should use a webserver. You can link to them using the same external content approach. Why should we keep federation, given that?
 - b. Kate Lynch - they use it in the infrastructure of repo currently being developed. Interested in exploring the possibility of them maintaining it if it were to become an external feature.
 - i. tried turning off federation and then used external system. Alt is to put files into Fedora directly. W/o connector you can't do the file system pattern w/in fedora. If that's what's needed then you need to figure out another way.
 - ii. Explore using a regular webserver in this scenario to see if it meets needs. If not, explain use case and we can discuss again.
 - iii. There is a bug with federation in modeshape 5 that prevents federation from seeing new files when Fedora's running.
 - iv. Reason for federation - the file system is a remote cloud base location. File system source of metadata, then redirect
 - v. One path is move to external modules and have it be a plugin from there. It could be maintained separately.
 - c. Is there anyone else outside who might be using federation feature? Might be worth asking around to see who might be using it. Wouldn't want to be a roadblock for people to upgrade.
 - d. Who is using it (maybe):
 - i. U. of Virginia
 - ii. Cleveland Museum of Arts – not sure if they are actually using it.
 - e. Suggestion is to broadcast this more loudly to ensure that the broader community knows about it and has a chance to see how it might effect their infrastructure. The group is planning on doing this.
6. 4.6 release schedule [Bill Branan](#) - Hydra In a Box is curious about the schedule.
 - a. Trying to balance the next Hydra release based on when Fedora next releases.

- b. [Unknown User \(acoburn\)](#) - three items he hopes would be in 4.6 release
 - i. How fedora module dependency chain works. The Fedora Http module depends on modeshape module - modeshape module could be entirely pluggable. There is a pull request that takes this part way. A user of Fedora wouldn't notice a difference, but the code would be differently organized.
 - ii. Import/export endpoints taken out. There is currently a notice about how it will be removed in 4.6. It's very JCR specific.
 - 1. Are the endpoints necessary for upgrade to Modeshape5? The endpoints could be part of an external module to access it for the upgrade.
 - 2. This hasn't been discussed a whole lot yet.
 - iii. How messages are serialized - the current format is a JMS message. At some point it would be nice if Fedora supported more than just JMS, so a standard format would need to exist. Currently is a header only format. The header of the property of what's changed in Fedora is wrong and hard to make work correctly. It would be good to get consensus on what that serialization should be.
 - 1. We should change this as soon as possible so that people who are writing code against this don't get too far in the process and have to change a bunch of things. All of the currently camel tools will need to be updated.
 - 2. [Unknown User \(acoburn\)](#) created a proposal (agenda item 8 from above) of what messaging could look like.
 - iv. [Unknown User \(acoburn\)](#) would not want to hold up a release for these.
 - c. Will 4.6 include Modeshape 5.X? Probably not 5.0, but if 5.1 is actually released on July13 – we could have release in August. That seems pretty far off, though.
7. Talked about removing JCR endpoints - fcr import/export - trying to move JCR specific features from core. Move them to extension module. Still be available, probably web-app plus, but not in the core anymore. fcrepo-serialization module would be part of this change and moved to extension. Moving JCR ones out is great. [Esmé Cowles](#) raised considerations about migration - there maybe be concerns, but if it's in webapp plus and most use that, then this may not be an issue.
- a. Maybe leave them there, but have them be Spring injected - could wire in JCR serialization or some future (maybe RDF) serialization. [Esmé Cowles](#) thinks this is the right way to move forward.
 - b. We should talk about this again when [Andrew Woods](#) is around.
8. Message serialization format - skipped, as already discussed above.
9. LDP Client available? More about the java client - there have been a few written for Fedora. The current one we have is pretty basic. Now some people are using it and are interested in having it be a bit smarter. Do you rewrite it or do you extend it to make it more sophisticated - or- do you find an LDP client that exists on which we can base this Fedora client. [Esmé Cowles](#) - There are a number of clients that are paired with other LDP implementations and not transferable. A generic java LDP client would be great or to talk to marmotta or others with existing client and maybe see if one of those could be generalized to work with more than one server. [Unknown User \(acoburn\)](#) looked at a few projects and it seems that many of them are tied to a specific framework. OSGi dependencies can make this challenging. Maybe just revamp java client
10. multi-tenancy - [Bill Branran](#) - hydra in a box wants to use fedora in a way that allows multi-tenant apps sitting on top of fedora.
- a. would like single fedora to scale out with cloud storage underneath. Currently thinking (writing?) Modeshape binary store for S3
 - b. make calls through fedora to potentially multiple tenants. Had talked to [Andrew Woods](#) regarding Modeshape workspaces to divide tenants
 - c. Question – should you be able to make links from one tentant to another? Not expecting to at ths point. One root node for each tentant
 - d. Question - binary store hierarchical tree currently based on sha1 . Would content from tenant a be completely separated from tenant b? Content from tenant a goes in one S3 bucket and tenant b in another S3 bucket. Would like the opportunity to put it in a totally separate AWS account. Keep the content separate and not intermixed at all.
 - e. Question - any links to S3 modeshape documentation? Has github repo that they just started -
 - i. There is a ticket in modeshape to create this, but was closed as nofix. [Bill Branran](#) will talk to them more about it.
 - f. If anyone has thoughts on S3 or multi-tenancy please get in touch with [Bill Branran](#).