

2015-07-02 - 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)
- U.S.A/Canada toll free: 866-740-1260, participant code: 2257295
- International toll free: <http://www.readytalk.com/intl>
 - Use the above link and input 2257295 and the country you are calling from to get your country's toll-free dial-in number
 - Once on the call, enter participant code 2257295
- 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

- [A. Soroka](#)
- [David Wilcox](#)
- [James R. Griffin III](#)
- [Aaron Birkland](#) ★
- [Stefano Cossu](#)
- [Unknown User \(daniel-dgi\)](#)
- [Andrew Woods](#)
- [Yinlin Chen](#)
- [Jared Whiklo](#)
- [Doron Shalvi](#)

Agenda

1. Officially deprecate fcrepo-message-consumer?
2. Performance testing needed!
3. Current Priorities

i. Performance	Unable to locate Jira server for this macro. It may be due to Application Link configuration.
ii. Single su	
	Unable to locate Jira server for this macro. It may be due to Application Link configuration.
iii. fcr:metac	Unable to locate Jira server for this macro. It may be due to Application Link configuration.
iv. Point-like	Unable to locate Jira server for this macro. It may be due to Application Link configuration.
v. Camel R	Unable to locate Jira server for this macro. It may be due to Application Link configuration.
vi. Migration	

• wait for mike

vii. Bugs

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

4. Service provider idea? <https://groups.google.com/d/msg/fedora-tech/UovyLGZeToE/INRtWZZ4JM4J>
5. 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.

6. 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. Seems to be general consensus for supporting fcrepo-camel-toolbox over fcrepo-message consumer. Shall we deprecate the consumer? If so what would that mean? Does git need to be re-organized at all?
 - Does fcrepo-consumer do anything that fcrepo-camel-toolbox doesn't?
 - Yes - it serializes rdf to disk as a preservation strategy. See also [FCREPO-1519](#)
 - What is the deprecation strategy? How long will it stick around?
 - This is the first to be deprecated
 - There have been several "will not fix" issues
 - There are a couple of policies on the wiki, but no deprecated modules policy
 - How is this module currently treated in the code and release process?
 - There are several projects in github that get simultaneous releases, and some with special documentation handling. Published to maven central, release published to github. This is one that gets published to maven and github as part of the fcrepo release process
 - Let's look at the overall release strategy. Currently, code is separated into fcrepo4 and fcrepo4-labs..
 - There should be a process for graduating from labs to fcrepo main project
 - It would be nice if all the release artifacts were coming out of fcrepo main project. Some come from fcrepo labs
 - fcrepo labs projects are pulled unintentionally in the release process anyway to bump up their version numbers, so dependencies stay in sync
 - ... so we need to define a graduation process and deprecation process
 - There is some disagreement that there is a natural progression from labs to core.
 - What does it mean to be in the main fcrepo organization? What are the core modules that define fedora, architecturally? During a release, what does the fedora organization commit to releasing?
 - The list of "core modules" and "commit to releasing" is likely different
 - What happens if a popular labs module breaks? Is it proper to expect committers to fix it?
 - There are projects built in jenkins that get built: <http://jenkins.fcrepo.org/>. We are notified when they break, and feel responsible to fix them
 - Is bumping the version for a satellite project mean advertising it as part of the supported release project? Bumping is an easy thing, but could be (mis) interpreted as an advertisement of support. We may want to say "do not use same versioning scheme as fcrepo project". Ran into this issue with ontologies as well
 - We have consciously broken Fedora projects into modules, to make main codebase as small as possible, have modules easy to use with the codebase. If we have them use their own versioning scheme, that may introduce a fair amount of overhead in the release process
 - Some of these modules may depend on Fedora APIs and core objects - but these are fairly unchanged with point releases. Should these be bumped? We don't really have a good definition of our public APIs. Some external modules use impl stuff.
 - If the public APIs aren't sufficient, we need to update them. Why does this module require functionality that is not in the public API?
 - Proposal: Start with effort to publish public API, discover which modules implement that, and consider that the core. Then define modules that should remain in repo because they are part of the central effort. Need a separate procedure for "taking responsibility" for something
 - Vote: Nobody opposed to deprecation
 - Action item: Announce the deprecation on lists. Maybe someone needs it, and wants to own it
 - Action item: Define a process for deprecation
2. Performance
 - Many aspects not yet tested
 - What does the community think is important to test?
 - Sharding?
 - How big can Fedora grow?
 - How many properties on objects?
 - How many total binaries, in different size ranges?
 - SCAPE has done clustering for ingest, did not really notice an improvement. Is it worth calling them?
 - We need to have real data
 - Definitely hit up the migration folks, as they are or will be shuffling in significant data.
 - One institution plans to migrate several terabytes of images, ingested as first-class resources. May federate them at first, then migrate. Not ready for this now, but it should happen at some point.
 - Action item: Create JIRA tickets for desired performance tests.

3. Is there any update on Modeshape ticket?
 - No. it continues to be kicked down the road
 - Please vote or comment on the issue!
4. Add service provider topic to top of next week's agenda, as we ran out of time!