

2016-06-02 - 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)
- 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](#)
- [Bethany Seeger](#)
- [Longshou Situ](#)
- [Unknown User \(acoburn\)](#) ★
- [Yinlin Chen](#)
- [Esmé Cowles](#)
- [Michael Durbin](#)
- [Kevin Ford](#)
- [Bill Branan](#)
- [David Wilcox](#)
- [Katherine Lynch](#)
- [Michael Gibney](#)
- [Andy Wagner](#)
- [Andrew Woods](#)
- [Doron Shalvi](#)

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](#)
 - i. A.Woods - clustering in AWS with objects and binaries in MySQL works
 - c.  Previously identified issue with ModeShape5 and Federation not picking up new files is resolved
 - d. Ready for community testing – focus on piloting migration via /fcr:backup and /fcr:restore
 - e. Production Fedora release of ModeShape5 is pending ModeShape release 5.1.0 to include [bugfix](#) – although testing can start now!
2. An `RdfSource` that is not a `NonRdfSourceDescription` currently describes itself. Should that be reflected in the Java API?
3. Open Repositories [Tech Meeting agenda](#) planning
4. Message serialization format (see [proposal](#))
5. Filesystem Federation Use Case & deprecation from core
 - a. [UPenn Use Case and additional info](#)
6. [Multi-tenancy](#) - needs from Hydra-in-a-box
7. Migrating and Round-tripping content... tooling needed
8. Use of Maven BOMs in the fcrepo4 codebase... embrace or abort? And generally, how are dependencies managed?
9. [Commons RDF](#) in fcrepo-kernel-api
10. Logging [recommendations](#) and [style-guide](#)
11. ...
12. Status of "in-flight" tickets

13. 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

- ModeShape5 update. There is a [branch](#) that is available and ready for community testing. The branch appears to be ready to be used, though it will not be merged into master until after a particular Modeshape5 bug is fixed in their 5.1.0 release.
 - Modeshape5 clustering has been demonstrated to work, but with caveats (it uses a single central database for objects and shared filesystem/database for binaries). Federation-related bug has been resolved with Modeshape5.
 - There will need to be tooling for migrating content, possibly with fcr:backup and fcr:restore or with some other utility.
 - Release timeframe for Modeshape 5.1.0: it appears to be scheduled for mid July. There is keen interest from the Hydra community with the release of Sufia7 w/r/t minimizing content migrations. With a three-week release candidate cycle, a Fedora-on-Modeshape5 release wouldn't happen until August at the earliest.
 - Clear communication and good tooling will be important.
 - There is interest in using S3 as a shared binary store in a clustered context. [Bill Branan](#) is working on a modeshape-based binary store connector supporting S3 integration. This may also be very useful for those interested in clustering.
- An `RdfSource` that is not a `NonRdfSourceDescription` currently describes itself. In the java code, the `NonRdfSource` class has a `getDescription()` method. By moving that method into the `FedoraResource` class, it could tidy up the code considerably (e.g. for example, this [snip](#) [pet](#) would be much shorter). It may also open up possibilities for specialized features related to provenance. There were no objections.
- Open Repositories tech meeting: if you plan to attend, please add your name to the [agenda](#). Everyone is encouraged to add agenda items. The focus will likely be related to the API specification, Fedora data migrations, etc.
- Message serialization format (see [proposal](#)). At some point it would be good to align the audit module's use of PREMIS and this use of PROV. Given that audit and messaging have similar but ultimately different goals, it may not be possible. The proposal uses JSON-LD, and if we use that serialization, we need to ensure that this JSON-LD must be parseable by real JSON-LD parsers. Everyone is encouraged to review and comment on this proposal.
- Filesystem Federation Use Case & deprecation from core – what are the reasons for deprecation?
 - In the process of specifying the Fedora API, this is not something that would be part of that specification; that is, alternate Fedora implementations would not necessarily implement this feature.
 - In the context of digital preservation, when a Binary is "in" Fedora, it doesn't necessarily mean that Fedora is managing that Binary (e.g. via external content or the filesystem-connector).
 - Filesystem federation has been seen to be unstable at scale
 - Penn [Use Case](#) is to use the filesystem metadata dynamically (e.g. preserving access controls), which isn't addressed by the external message body construct.
 - API-X may be an answer to this.
 - More to be discussed next week
- [Multi-tenancy](#) - needs from Hydra-in-a-box: please read and comment on the linked document. This will be discussed next week, too.