

2017-11-30 - 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](#) ★
- [Peter Eichman](#)
- [Andrew Woods](#)
- [David Wilcox](#)
- [Carrick Rogers](#)
- [Michael Durbin](#)
- [Esmé Cowles](#)
- [Seth Shaw](#)
- [Bethany Seeger](#)
- [Aaron Birkland](#)

Agenda

1. Announcements
 - a. [Candidate Recommendation released!](#)
 - b. AWS Graph Store ([Neptune](#))
 - c. Oxford Common Filesystem Layout [meeting... tomorrow, Friday, Dec 1st](#)
2. Needing volunteers
 - a. [New Fedora Clustering configuration](#)
 - i. Testing volunteers?
 - b. Fedora workshop leader for [Code4Lib 2018](#) in Washington, DC, February 13-16
 - i. [Esmé Cowles](#) and [Yinlin Chen](#)
 - c. Facilitator for Fedora user group meeting at 2018 Digital Initiatives Symposium in San Diego, April 24
 - i. [Danny Bernstein](#)?
 - d. Fedora exhibitor at [Southern Miss IR Conference](#) in Hattiesburg, MS, April 26-27
 - i. ??
 - e. Creating an auto-invite Fedora slack page?
 - i. <https://fedora-project.slack.com/>
 - ii. Like: <http://slack.samvera.org/>
 - iii. Using: <https://github.com/rauchg/slackin>
3. 4.7.5 release - Planning for January 2018
 - a. Are there any missing bug fixes?
 - b. Release manager?
4. Discussion of Real World Objects and Fedora
 - a. Any objection to relaxing Single Subject Restriction?
5. 5.0.0 release
 - a. [API Alignment](#)

Unable to locate Jira server for this macro. It may be due to Application Link configuration.

Unable to locate Jira server for this macro. It may be due to Application Link configuration.

6. Beyond 5

- a. Persistence?
- b. Journaling?
- c. Simple, synchronous query?

7. Tickets re

Unable to locate Jira server for this macro. It may be due to Application Link configuration.

Unable to locate Jira server for this macro. It may be due to Application Link configuration.

Unable to locate Jira server for this macro. It may be due to Application Link configuration.

[repo/kernel/modeshape](#)

- waiting on feedback from [C](#)

- ready to go? [Danny](#)

- d. Unable to locate Jira server for this macro. It may be due to Application Link configuration. - ready to go? [Danny](#)
8. ...
9. Tickets In-Review

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. Announcements
- a. Fedora API specification was just released.
 - i. Hoping in the February timeframe to have 2 verified implementations with the [test suite](#).
 - ii. What are the 2 implementations might be, the Modeshape one, the other is possibly Cavendish there is a Python implementation that [Simeon Warner](#) is working on.
 - iii. [Esmero](#) is a possible PHP implementation, not much movement lately.
 - iv. Northwestern was doing implementation in Python with an AWS lambda backend but had to set it aside.
 - v. [Derby](#) is another possible contender to implement the API.
 - vi. [Lambdora](#) - an AWS fedora
 - vii. [Trellis](#) - supports all of the underlying specs, but may not conform to all of the fedora-specific requirements
 - b. Amazon now supports a graphstore that support SPARQL query.
 - c. Call tomorrow at 11:00 ET for the Oxford Common Filesystem System.
2. Activities
- a. Postgres clustering configuration, Andrew spoke with Rob Kaufman and he has implemented this at UCSD. He says it supports multiple masters for writes, which would be very interesting. Can we get someone to stand-up a cluster and test this. [Carrick Rogers](#) might have some resources to stand up some instances to test this.
 - b. Anyone going to Code4Lib on helping [Esmé Cowles](#) with a workshop. [Michael Durbin](#) will be there and can help out.

- c. Digital Initiatives Symposium needs someone to help facilitate a Fedora user group meeting, [Danny Bernstein](#) will be there he thinks, [Set h Shaw](#) is going for Islandora and might be able to help.
 - d. Fedora exhibitor at the Southern Miss IR conference, do we have anyone to do this and/or possibly a proposal for the topics? Anyone in the area? Contact [David Wilcox](#) if you might be interested.
 - e. There is a Fedora slack channel, but there is a pattern for creating an auto-invite. Samvera has setup this type of work. Declan or Mike Giarlo have set it up or are maintaining it.
3. 4.7.5 Release
- a. Shooting for a release in January. Any reason not to do it now? Just the time of year?
 - b. Any other bug fixes?
 - c. Any volunteer to do the release?
4. Single Subject Restriction
- a. Would this allow you to persist subjects that are not within their repository.
 - b. Relaxing the single subject could allow several bad patterns, all descriptions on one resource for example.
 - c. They want to have various stuff in their resource, some is of a real world object that is published elsewhere.
 - d. Is it Fedora's job to restrict people to keep them from making bad decisions.
 - e. There are easy ways to do this by persisting these statements in a serialized XML.
 - f. If we are going to maintain the SSR then we need to ensure that we well document the principle that leads to that decision.
 - g. In regards of named graph provenance you might want separate audit statements held in different places.
 - i. a resource in the repository is a named graph, and it would be plausible that a resource could contain statements that describe the provenance of other resources in the repository.
 - ii. If you want to manage the above as a separate resource, that could not be done well.
 - iii. Storing ontologies as a RDF source is not able to be stored in the current repo based on SSRs and blank nodes.
 - h. No ideological attachment, but there might be real world reasons for having it.
5. 5.0.0 release
- a. Have a look and please add any notes you might have for future discussion
6. Beyond 5.0.0 - Areas for improvement
- a. Have a look and please add any notes you might have for future discussion