

2017-09-28 - 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

- [Danny Bernstein](#)
- [Aaron Birkland](#)
- [Jared Whiklo](#) ★
- [Peter Eichman](#)
- [Joe Harrington](#)
- [Joshua Westgard](#)
- [Yinlin Chen](#)
- [Ben Pennell](#)
- [Bram Hauer](#)
- [Andrew Woods](#)
- [Bethany Seeger](#)

Agenda

1. Update from "[Alignment to Spec](#)" sprint
 - a. Documentation updates go to: [Fedora 5.x Documentation](#)
2. [Oxford Common Filesystem Layout](#)
3. Wrapping up fcrepo-import-export-verify: <https://github.com/fcrepo4-labs/fcrepo-import-export-verify/pull/48>
4. Tickets In-Review

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

5. ..

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

Actions

1. API Alignment Sprint

- The first sprint moved alot of CRUD and fixity alignment issues, did lots of planning on versioning and authz.
- The second sprint carrying on and trying to move the versioning into a final design that will work and handling loose ends around Authz.
- Meeting at 12:00 ET today to talk about the actual planning for versioning.
- Important information from Sprint 1 - Ensure the versioning design implementation as no overlap between versioning sprint team members from 1 to 2
- Meeting Tuesday Sept. 26 with some Sprinters and Memento folks to:
 - Checked the design document.
 - Checked the delta document.
 - Checked the JIRA tickets created around versioning.
 - Asked questions of the Memento folks around where there were some confusion.
- [Bethany Seeger](#) - Memento folks thought the best plan for action was that the LDPRv LDPCv and all LDPRm(s) could have seperate ACLs.
 - LDPRm(s) would all share a single ACL, not per Memento.
- Another call with the Memento team on next Tuesday so having questions fleshed out before that call would be beneficial.
- WebAC codebase has been merged into the core codebase, now working on hooking it up and looking at ways to enable/disable AuthZ at deploy time.
- Consider moving the old WebAC repository out of the fcrepo4 organization, but annotationing the README that it is still the major AuthZ for 4.7.
- SOLID Spec for aclAgentGroup being a list of vcard members, where agentClass is for public. Note at the end of this section (<https://github.com/solid/web-access-control-spec#public-access-all-agents>) was confusing.
 - Should we remove agentClass or are we just adding agentGroup for lists (which we previously used agentClass).
 - [Andrew Woods](#) believes to limit use of agentClass to foaf:Agent for now. acl:agentClass should only be used with foaf:Agent to ascribe public access. [Esmé Cowles](#) and [Daniel Lamb](#) agrees.
- 5.x documentation space has been created to capture the changes.

2. Oxford Common Filesystem Layout

- There have been discussion around preservation sensibilities with respect to Fedora.
- With the API specification and defining user level interactions and allows an abstract layer to separate client from Fedora.
- The sentiment of the OCFL is the same in how information is persisted on the filesystem. What is persisted and what information and what format.
- In the Modeshape we would not force Modeshape to use this, but could bundle functionality to also synchronously or asynchronously persist to the filesystem using the OCFL initiative.
- Could be useful for interoperability scenarios (ie. import/export, BagIt creation, Camel serialization service) where the point is to have something persisted and using a filesystem abstraction layer.
 - A little uncomfortable if a repository would be expected to make accessible it persisted storage on the filesystem in this format.
 - When talking about a backup you need a static system and most times the information persisted to the filesystem is not sufficient to perform a restore. You would need to include additional information.
- From a preservation perspective (and Modeshape) the binaries are relatively safe, we lose the filename (is replaced by the SHA-1). The metadata is in the database. The nature of the database in which the data is stored could be considered fragile. From a longer term preservation stand-point and recognizing the application (Fedora) as the steward of the content now, but it leaves open the possibility of making transition to future stewards.
 - perhaps the journaling used in Fedora 3 could be a consideration.
 - Fedora becomes the files and everything else (the software) is fungable. The OCFL becomes the Fedora API.
 - Fedora had beauty with the persistence on disk, but performance issues in reading the data off files at every turn.
 - Users can see the data and that data can be moved around.
 - Camel serialization toolkit and import/export so you are not messing with live data.
 - It would be possibly to store everything as files on disk, binaries as themselves and RDF sources as also files. Possible but not necessarily part of the specification. Each institution had a different use case.
- Is there enough interest and value seen to people in the community for the OCFL as a specification or best practice guidelines?
 - Broadly how do you persist links between resources once they are on the filesystem.
 - Import/Export made their own layout, which is different from BagIt (for example). So this could be useful for standardization around this type of output.

