

2019-10-03 - Fedora Tech Meeting

Time/Place

This meeting is a hybrid teleconference and slack chat. Anyone is welcome to join...here's the info:

- Time: 11:00am Eastern Daylight Time US (UTC-4)
- Audio/Video Conference Link: <https://lyrasis.zoom.us/my/fedora>
 - Dial-in:
 - +1 408 638 0968
 - +1 646 876 9923
 - +1 669 900 6833
 - Meeting ID:
 - 812 835 3771
- Join fedora-project.slack.com on the "tech" channel

Attendees

Part 1:

1. [Danny Bernstein](#)
2. [Andrew Woods](#)
3. [Peter Winkles](#)
4. [Jared Whiklo](#)
5. [Bethany Seeger](#)
6. [Youn Noh](#)
7. [Thomas Bernhart](#)
8. [Daniel Lamb](#)
9. [Aaron Birkland](#) ★
10. [Ben Pennell](#)
11. [Rosie Le Faive](#)
12. [Dan Field](#)
13. [Ben Cail](#)
14. [Mohamed Mohideen Abdul Rasheed](#)
15. [Joseph Rhoads](#)
16. [David Wilcox](#)

Agenda

1. Announcements
 - a. [OCFL Editors meeting](#)
2. Fedora 6 and OCFL: managing active updates
 - a. auto-versioning
 - b. versioning on-demand
 - c. no versioning
3. Discuss [Ben Pennell's proposed additions to the persistence API](#)

Tickets

1. 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.										

2. 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.

3. 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.

4. 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.

Notes

1. Announcements

- OCFL Editors Meeting
 - Primary objective of the editor meeting this week is to tie up any loose ends for the 1.0 specification
 - Wrap-up documentation, press release, push the validator forward
- Sprint coming up on Nov 4
 - Signup page: [2019 Fall Sprints - Fedora 6](#)
- Polished sprint summary video COMING SOON
- Thomas will be testing the migration utils next week
 - Andrew was going to do a short video about the tool as well
 - Its still fairly "experimental"
 - Dan and Andrew available to field questions
- 3 Dans on the call. Three.

2. Fedora 6 and OCFL: managing active updates

- Possible approaches:
 - Mutable head, where live updates go into a head directory
 - Or putting changes off in a separate directory
- Editor sense so far is to use a deposit directory outside of the root (which is no longer in the OCFL storage root)
- OCFL's storage root is intended to be preservation storage, not live update space. Things in the root are immutable. Might even be appropriate for "Warm storage" where can only be written once (how would that work with inventory updates?)
- How can we encourage fedora users to minimize the amount of time that stuff is in an unversioned state.
- Andrew proposing toggle-able on-demand versus auto-version option
- Bethany suggesting that serialization of take place asynchronously
- Jared - if we're going to be maintaining two storage approaches, might as well maintain one that gets converted OCFL after a while
 - bethany hating the idea of double storage
- dbernstein suggesting if object in motion, only things that have change are in the directory external to OCFL
- andrew - deposit directory is the mutable head. Transaction isolation within this directory.
- jared - maybe could represent changes in database rather than on disk.
- Peter is wondering if we are expecting the OCFL library to implement this mutable head, or if that is a fedora concern
 - Andrew was assuming that the clients had a notion of deposit directory
 - Peter was only considering this for versions that were being immediately committed, so only lasts for duration of method call
 - It would be possible to make it longer lived, but it would be a considerable rework of the library to support. Its possible, but Peter doesn't like it.
 - Peter feels like OCFL is not a good storage format for in motion data, so fedora would have ownership of everything prior to committing to OCFL
- How does OCFL feel about external binaries?
 - It hasn't been talked about a great deal, its just a link to something that's not there.

- OCFL would just be a set of metadata about files
 - Having OCFL directly track external binaries in its inventories is not currently supported
 - But there may be a distributed storage root concept in later versions of OCFL.
- Object stores
 - there wouldn't be notions of move operations, mutable head sort of goes out the window
- What are the characteristics that fedora wants out of an external deposit directory?
 - Principle of least surprise - so users don't wake up in 2 years and find nothing in OCFL
 - Autoversioning by default, if performant for normal workflows?
 - Make clear what the versioning approach is in the UI/API
 - Ben Cail and Joseph would want to use push to OCFL at end of every action/transaction
 - danny lamb - islandora would probably be the staging ground rather than the fedora mutable head
 - Ben Cail has a similar situation with
 - Peter is stating that there is data duplication, there are copies before ingest, and then moved into OCFL for preservation purposes.
 - Ben Cail - the OCFL client would become a fedora specific client if it handled all the object build up stuff.
 - Ben P - mentioning that in his local repo, as objects added to a collection it would produce new versions of the collection if any state was tracked about the children (in his case, PREMIS of the collection is updated). This was to try to illustrate that in situations where resources are added and never modified, there may still be fedora objects that receive many modifications in order to track new objects.

Actions

- [Aaron Birkland](#) to look explore notion of OCFL client with database as authoritative metadata source + asynchronous writing of the inventory.json file
- [David Wilcox](#) will review the NDSA matrix and pull out the concrete technical requirements that could be considered during the Fedora 6 development.