

2019-10-17- 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. [Peter Winckles](#)
3. [Jared Whiklo](#)
4. [Bethany Seeger](#)
5. [Thomas Bernhart](#)
6. [Aaron Birkland](#)
7. [Andrew Woods](#)
8. [Ben Pennell](#) ★
9. [Ben Cail](#)
10. [David Wilcox](#)
11. [Peter Eichman](#)

Agenda

1. Announcements
 - a. [Sprint 2](#)
 - b. ?
2. See if we can get a unified vision for how we want Fedora to interact with OCFL, so that we can propose one of the following to the leadership group:
 - a. Recommend pushing for changes to the OCFL spec (i.e. a mutable head), so that both versioned and un-versioned content can be "in" OCFL
 - b. Accept the immutability of object versions in OCFL, and recommend one of the following
 - i. Fedora shall store un-versioned content outside of OCFL in a fedora-specific layout
 - ii. Fedora shall disallow un-versioned content entirely
 - c. Recommend Fedora use something other than OCFL for Fedora's persistence layer, and create external tools for exporting/publishing to OCFL when desired.
3. [This document](#) describes the core issues and outlines the pros and cons of the implementation possibilities described in the preceding agenda item. Recommending a particular implementation is beyond the scope of the document; it simply seeks to provide an objective (to the degree possible) analysis to serve as a baseline for an informed discussion. Comments and suggestions are welcome, but please respect the scope of the document.
4. Other questions to be answered
 - a. > "When making a request during a transaction, should the ACL of the HEAD version of the resource be evaluated, or the ACL of the version of the resource that exists in the transaction?"
 - i. Related to <https://github.com/fcrepo4/fcrepo4/pull/1559>
 - ii. from [Decisions and Open Questions](#)
 - b. What is the approach we want to take for adding digests to external binaries?
 - i. Reference: [2019-09-26 - Fedora Tech Meeting](#) (Notes: 3.a.)
 - c. What do we want our approach to be for loading Fedora6 resources? lazy or eager?
 - i. Related to <https://github.com/fcrepo4/fcrepo4/pull/1557>
5. <add topics here>

Tickets

1. In Review

| T | Key | Summary | Assignee | Reporter | P | Status | Resolution | Created | Updated | Due |
|---|-----|---------|----------|----------|---|--------|------------|---------|---------|-----|
|---|-----|---------|----------|----------|---|--------|------------|---------|---------|-----|

| | | | | | | | | | |
|--|--------------|--|-----------------|-----------------|--|--|----------------|--------------|--------------|
| | FCREP O-3698 | Java client improperly encodes filename field | Ben Pennell | Ben Pennell | | | Fixed | Apr 15, 2021 | Apr 15, 2021 |
| | FCREP O-3695 | Single-request transactions should not use db transaction tables | Unassigned | Peter Winckles | | | Ready for Test | Apr 12, 2021 | Apr 29, 2021 |
| | FCREP O-3692 | Add metrics to migration-utils | Michael Ritter | Danny Bernstein | | | Fixed | Apr 08, 2021 | Apr 29, 2021 |
| | FCREP O-3691 | Improve performance of search queries that return rdf types | Danny Bernstein | Danny Bernstein | | | Fixed | Apr 08, 2021 | Apr 29, 2021 |
| | FCREP O-3690 | Add a configuration option for setting the S3 endpoint | Peter Winckles | Peter Winckles | | | Ready for Test | Apr 05, 2021 | Apr 29, 2021 |
| | FCREP O-3686 | fcrepo-migration-validator: validate state in datastream | Michael Ritter | Danny Bernstein | | | Fixed | Mar 29, 2021 | Apr 26, 2021 |
| | FCREP O-3664 | fcrepo-migration-validator: Add Datastream version validation for size | Michael Ritter | Danny Bernstein | | | Fixed | Mar 07, 2021 | Mar 22, 2021 |
| | FCREP O-3612 | Test performance of DB rebuild optimization | Danny Bernstein | Danny Bernstein | | | Fixed | Jan 27, 2021 | Feb 23, 2021 |
| | FCREP O-3593 | API Test Suite: 4.1.2-B | Danny Bernstein | Andrew Woods | | | Fixed | Jan 11, 2021 | Mar 11, 2021 |
| | FCREP O-3561 | Add commit() and rollback() support to DbSearchIndex Impl | Danny Bernstein | Danny Bernstein | | | Fixed | Dec 04, 2020 | Apr 29, 2021 |
| | FCREP O-3260 | Document OCFI limitations related to paths and OS limitations | Peter Winckles | Andrew Woods | | | Ready for Test | Mar 26, 2020 | Mar 03, 2021 |
| | FCREP O-2740 | Float value not updated when current value rounds to the new one | Yinlin Chen | Kevin Ford | | | Fixed | Mar 27, 2018 | May 30, 2018 |

12 issues

2. Please squash a bug!

| Key | Summary | T | Created | Updated | Due | Assignee | Reporter | P | Status | Resolution |
|-----|---------|---|---------|---------|-----|----------|----------|---|--------|------------|
|-----|---------|---|---------|---------|-----|----------|----------|---|--------|------------|

| | | | | | | | | | |
|--------------|---|--|--------------|--------------|------------|-----------------|--|------|------------|
| FCREP O-3697 | Multi-request transactions can be put in a bad state | | Apr 14, 2021 | Apr 29, 2021 | Unassigned | Peter Winckles | | OPEN | Unresolved |
| FCREP O-3678 | Fedora 5 and SSL configuration not working | | Mar 19, 2021 | Mar 19, 2021 | Unassigned | Jaime Penagos | | OPEN | Unresolved |
| FCREP O-3655 | application/link-format returns multiple entries for the same second | | Feb 25, 2021 | Mar 03, 2021 | Unassigned | Jared Whiklo | | OPEN | Unresolved |
| FCREP O-3635 | Fedora 6 does not start on Jetty 11 or Tomcat 10 | | Feb 16, 2021 | Feb 19, 2021 | Unassigned | Peter Winckles | | OPEN | Unresolved |
| FCREP O-3620 | Fedora 5 -> 6 upgrade fails if a Fedora 5 export has retrieved the external binaries | | Jan 29, 2021 | Apr 29, 2021 | Unassigned | Jared Whiklo | | OPEN | Unresolved |
| FCREP O-3589 | Membership /containment not returned as inbound references | | Jan 06, 2021 | Jan 27, 2021 | Unassigned | Ben Pennell | | OPEN | Unresolved |
| FCREP O-3585 | Rollback with auto-versioning disabled can leave Fedora in an inconsistent state | | Dec 09, 2020 | Feb 03, 2021 | Unassigned | Peter Winckles | | OPEN | Unresolved |
| FCREP O-3584 | Resource locking does not take into account ghost nodes | | Dec 09, 2020 | Feb 03, 2021 | Unassigned | Peter Winckles | | OPEN | Unresolved |
| FCREP O-3538 | Wildcard searches and ids containing * | | Nov 19, 2020 | Nov 19, 2020 | Unassigned | Peter Winckles | | OPEN | Unresolved |
| FCREP O-3521 | Ensure all RDF formats supported in fcrepo-import-export also available in fcrepo-upgrade-utils | | Nov 06, 2020 | Nov 30, 2020 | Unassigned | Danny Bernstein | | OPEN | Unresolved |
| FCREP O-3506 | OOM on GET of root with many child containers | | Oct 29, 2020 | Nov 16, 2020 | Unassigned | Andrew Woods | | OPEN | Unresolved |
| FCREP O-3423 | ACLs not displayed properly | | Sep 12, 2020 | Nov 18, 2020 | Unassigned | Peter Winckles | | OPEN | Unresolved |
| FCREP O-3286 | Exception on malformed Link header | | Apr 25, 2020 | Nov 18, 2020 | Unassigned | Andrew Woods | | OPEN | Unresolved |
| FCREP O-3252 | Disallow creating mementos within the same second | | Mar 17, 2020 | Nov 18, 2020 | Unassigned | Peter Winckles | | OPEN | Unresolved |

| | | | | | | | | | |
|-----------------|--|---|-----------------|-----------------|------------|------------------------|---|------|------------|
| FCREP O-2263 | Reserved Characters in fragments are encoded by UriBuilder |  | Oct 09, 2016 | Nov 27, 2020 | Unassigned | Christopher Johnson |  | OPEN | Unresolved |
|-----------------|--|---|-----------------|-----------------|------------|------------------------|---|------|------------|

15 issues

3. Tickets resolved this week:

| Key | Summary | T | Created | Updated | Due | Assignee | Reporter | P | Status | Resolution |
|-----|---------|---|---------|---------|-----|----------|----------|---|--------|------------|
|-----|---------|---|---------|---------|-----|----------|----------|---|--------|------------|

No issues found

4. Tickets created this week:

| Key | Summary | T | Created | Updated | Due | Assignee | Reporter | P | Status | Resolution |
|-----|---------|---|---------|---------|-----|----------|----------|---|--------|------------|
|-----|---------|---|---------|---------|-----|----------|----------|---|--------|------------|

No issues found

Notes

1. Announcements

- Islandoracon was great, there was a fedora special interest group that met during lunch and did fedora presentations
 - Bento boxes and preservation. Drupal 8
 - Can pop in and out components like Fedora. For some, its considered "in" the repository once it is in fedora.
 - A lot of people looking forward to Fedora 6
 - Workshop to get people to migrate to Islandora 8, which required moving from F3 to F5 went well.
- Sprint 2 - It would be helpful to try to decide some of the Open Questions before the next sprint [Decisions and Open Questions](#)

2. Uniform vision for how Fedora should interact with OCFL

- Jared - it does not make sense to store fedora objs in two different places. It puts a burden on Fedora for the sake of OCFL
 - he likes OCFL, but wants a transparent file layout from the application perspective.
 - Fedora deals with so many in motion objects, it may not be worthwhile
- Andrew - historically users have not differentiated between in-motion and preserved
- Bethany - could there be multiple clients operating on the OCFL?
 - Andrew - we should be very clear that only multiple read clients
- Is option A a change to the spec?
 - That is what Andrew as thinking, that there would be a HEAD version in the spec
 - As compared to rewriting the most recent version
- There are four options, as outlined in the agenda subpoints and the document
 - Jared: A => C (falling back to C if A fails)
 - BenP: A => C or B1
 - Bethany: A => B2 => C
 - AaronB: A => B2 => C
 - BenC: A => B1 with encouragement to users to version or B2 (doesn't like C)
 - PeterW: A (cleanest but seems unlikely) => C
 - Thomas: A => C
 - DannyL: C (unless vociferously opposed by leaders) => A
 - AndrewH: Doesn't feel like he can voice an opinion about what fedora should do
 - AndrewW: A (wants fedora to be as close to OCFL as possible)
- PeterW - Another option is to serialize to non-OCFL and periodically push to OCFL
 - Jared - similar to B1 in that there are two places where content is stored
- Andrew - Not clear that having a mutable head directory would take away from the value of the version directories
- Aaron - A is somewhat out of our control, since it depends on the OCFL editors. It is not so much a technical issue.
- David was suggesting that we might just do A even if it wasn't in the spec
 - PeterW - Fine if you don't care about validation, but the bigger issue might be that the inventory would be referencing content that isn't allowed, so the object might not be readable.
 - Assuming that the root inventory file would point to HEAD as the current version
 - Aaron - since the version inventories are the canonical inventories, then might not even need to add anything to the root inventory.
- RECOMMENDATION:** Option 'A', if supported by the OCFL specification

4. Other questions to be answered

- a. "When making a request during a transaction, should the ACL of the HEAD version of the resource be evaluated, or the ACL of the version of the resource that exists in the transaction?"
 - Jared - it would require two actions, first updating ACL and then starting transaction and doing work
 - BenP - it seems like a special behavior to only work against the HEAD of the ACL versus everything else in the repo
 - **DECISION:** Treat ACL's as normal, so use the version in the transaction unless this causes significant technical issues
- b. What is the approach we want to take for adding digests to external binaries?
 - We talked about this some on Sept 26, Andrew remembers that we wanted fedora to retain fixity for internal binaries. For external, users can add their own. For internal, users can't override the internal digest.
 - Fedora has its own internally managed checksum for internal content
 - Users can add digests to external content
 - Should fedora try to fill in a checksum if none is provided?
 - Jared - spec says that ext's must respond to want-digest. So do we need to store it internally? How do we fulfill
 - **DECISION:**
 - Fedora to retain fixity for internal binaries
 - For internal, users can't override the internal digest, but can add supplemental digests
 - For external, users can add their own

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.