

2018-09-27 - 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://duraspace.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

1. [Danny Bernstein](#)
2. [Peter Eichman](#)
3. [Bethany Seeger](#) ★
4. [Jared Whiklo](#)
5. [Andrew Woods](#)
6. [Kevin Ford](#)
7. [Ben Pennell](#)
8. [Aaron Birkland](#)

Agenda

1. Announcements
2. Sprint 2 Strategy and Goals
 - a. Release candidate
 - i. Vagrant
 - ii. Camel toolbox
 - iii. java client
3. Open questions

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

ii. The cause of the 500 is

```
javax.jcr.RepositoryException: The session with an ID of '060742fc6' has been closed and can no longer be used.
```

- iii. It appears to be an OSX 10.13.x issue : [Jared Whiklo](#), [Peter Eichman](#), and [Danny Bernstein](#) have reproduced.
 1. [Danny Bernstein](#) confirmed It was not a problem for OSX 10.10.5
 2. Any ideas of how to solve it?
 - iv. If we can't resolve it, what are our options?
 1. Leave the tests in place - and document the problem in the README.
 2. @Ignore the tests until we can sort out the issue or it sorts itself out (OS fix)
 3. ?
- b. Data URIs in the constrainedBy link header
 - i. <https://github.com/fcrepo4/fcrepo4/commit/750426101de9330cafaa88838a3d694d1ef69580#diff-14e2f9192b9abd40557a43f9342f4481R44>
 - ii. Proposal: Remove and use standard URIs instead like we do with all other constraints.
 - c. Prefer headers on Mementos: <https://github.com/fcrepo4/fcrepo4/pull/1420>
 - i. Question1: Should the original resource return <> a ldp:Container and <> a ldp:RDFSResource when omitting server managed triples (currently these are being omitted in OriginalResource but not the Memento)
 - ii. Question2: How should fedora respond to requests to Prefer omit
 1. What would a use case for omitting server managed triples in memento?
 2. Proposal 1: Mementos should behave like the original resource
 - i. Rationale?
 - ii. Action: based on the answer to Question1 either omit the ldp types from memento or add them to the original resource response
 3. Proposal 2: Mementos should not respond to Prefer omit
 - i. Rationale?
 4. Proposal 3: ?

d. [All AuthZ tests assume literal agent values](#)

- i. What is the best way to fix this problem?
- 4. <your agenda item here>
- 5. 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.

6. 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.

7. 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. Next sprint starts next week!
2. Lots of work happened during the sprint - many CTS tests filled in. Nice job!
3. Sprint 2 Strategy and Goals
 - a. Goal: get to the Release Candidate
 - i. There are a number of tickets there that we know we want to address. Priority on those first.
 - ii. Setup Release Testing work - pages, testing that the process still works and address any issues seen.
 - iii. Documentation
 - iv. Jared suggests focusing on the maybe 5 tickets that are bugs, leaving new features for later.
 1. Suggestion of looking at all tickets to prioritize them.
 - v. Someone can look at fcrepo-camel-toolbox and java client to get an idea of what needs to be done there.
 - b. Secondary Goal: tool set. Priority would be to update camel-toolbox and java client first, vagrant last (as it depends on the other things and may not be much more work to update).
 - i. Ben P. suggests documenting the changes that are needed to the fcrepo tool set: fcrepo-camel-tool box, java client.

4. Open questions:

- a.

Unable to locate Jira server for this macro. It may be due to Application Link configuration.
- ii. Mac 10.10.5 doesn't show this error. (10.10.5 = Yosemite)
- iii. 4 tests failing in FedoraLDAPIT tests. Same 4 tests failing. Something to do with setting text HTML in header - to return the web page. Maybe something to do with how Grizzly is working; may be a timing issue.
- iv. Peter E. Do we think this is a network/stack/socket communication problem? Any leads on what's causing it? He's seen some weirdness in their batch loading clients in the network HTTP/TCP stack that weren't working in Mac OS High Sierra, but worked on Sierra. High Sierra may have broken some networking issues. Danny B. suggests leaving it since it may be a Mac issue - but many of us use Mac OS's to build; we'd have to ignore those tests. Can we do a conditional ignore on Mac OS's?
- v. Jared suggests we try to look at this more. In the test, in Velocity the SessionFactory and http servlet request is injected, using context...?. Maybe something he's written there is not the best way to do this. Info he's trying to get is that it needs to know if it is a memento and if it's versionable. Maybe there's an easier/better way to get that data?
- vi. One proposal - have someone look at it now
- vii. Or make a ticket, put a conditional ignore on the tests, put some manual tests for this in the release plan, and keep moving forward, addressing this later.
- viii. Ben Pennell may take a short look at it. (He has a High Sierra mac)
- b. Data URIs in the constrainedBy link header

- i. Default response on malformed RDF is a Data URI in the constrained By link. We are not using data uris anywhere else. You can't just easily read the response and see what's happening there.
- ii. No objection to changing that. Danny B. will create a PR.
- c. Prefer headers on Mementos: <https://github.com/fcrepo4/fcrepo4/pull/1420>
 - i. Question1: Should the original resource return <> a ldp:Container and <> a ldp:RDFSResource when omitting server managed triples (currently these are being omitted in OriginalResource but not the Memento)
 - 1. The thought is that this is a new bug - that the original resource should return those two triples. Danny B will create this ticket.
 - 2. Danny B will create a documentation ticket for what the SMT's actually are.
 - ii. Question2: How should fedora respond to requests to Prefer omit
 - 1. In theory they should be the same - that what you see with the original based on the representation request, should look the same in the memento.
 - 2. Proposal 1: Mementos should behave like the original resource
 - i. We'd like to go with this.
- d. All AuthZ tests assume literal agent values
 - i. Danny believes we have some way to handle this. You can specify the baseURI for agents? The actual logic is backwards from expected. It retains the comparison between the username from current user and the agent value - is still a string comparison of just the user name. The BaseURI user webac uri prefix, its not being added to the principal from the container, but it is being stripped coming from the WebAC. Logically you end up comparing things in the correct sort of way.
 - ii. Maybe the fix is to make sure that the values going into the ACL's are URIs. Any of the agents that are being tested by the CTS should probably be URI's. Then when running against fedora, we'd have to specify the agent base URI to start up the server with.
 - iii. Need to start collecting a recipe of how to configure your fedora to run the CTS on it. (external content setup, webac agent setup, etc)
 - iv. CTS is assuming HTTP basic auth; leaving out server certs, OAuth, etc.. What should we support? What does the Fedora API have to say about how it get a user name from the request? Is anything specified for that?
 - v. WebID passed in the header? WebID TLS - auth using WebID – Solid spec mentions this by name. <https://github.com/solid/solid-spec#webid-tls>
 - vi. Maybe create an issue for this and address it during the upcoming sprint.
 - vii. Solid spec mention WebID and WebID TLS, but don't seem to be musts.
 - viii. Question of what are we assuming about WebID based on what Solid requires?
 - ix. Zimeon raised a point about users and names and we should check in with him
 - x. All of CTS assumes basic auth but some systems may not be using basic auth. Would be hard to cover testing all auth systems. We could work with folks implementing the spec and see what they are using.
 - xi. Aaron Birkland mentioned CLAW and JWT's - solution was to have the system hand off a http client off to APIX, which does all the JWT stuff. We could add an authentication interface for the CTS in order to delegate authentication to the implementation. The tradeoff – Fedora implementors will have to implement the interface. However, this approach might be more practical than trying to support specific systems folks are using.

Future Agenda Topics:

- Amherst College's acrepo apache camel services: <https://gitlab.amherst.edu/acdc/repository-extension-services>
- Reviving OAI