2017-01-26 - 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
- Esmé Cowles
- Jared Whiklo
- Kevin Ford
- Longshou Situ
- Jim Coble
- Andrew Woods
- David Wilcox ⭐
- Yinlin Chen
- A. Soroka
- Benjamin Armintor
- Daniel Lamb
- Aaron Birkland
- Karen Estlund
- Dan Coughlin

Agenda

1. Fedora 4.7.1 release is out
   a. Release Planning
2. Fedora/API-Spec delta-document: Aaron Birkland available to spearhead?
   a. Initial planning about scope and purpose: google doc
3. Fedora/API-Spec - is it ready for initial community input (proposed deadline for input May 1)
4. Focus on intra-repository-resource-relationship performance
   a. Storing these relationships as URIs (Prefer:InboundReferences and referential integrity go away)
   b. Optimizing the backend database index
5. Fedora DevOps interest group in collaboration with Islandora and Hydra
6. Trajectory of development
   a. Common API and import/export facilitates transfer between impls
      i. Transfer impact mitigated by "external-content"?
7. Interest in Fedora Docker? ...looking for another maintainer for fcrepo4-docker
8. WebApp configuration possibilities - fcrepo-webapp-plus limitations
   a. AuthZ
   b. Audit
   c. Minter
   d. Backup/Restore
   e. API-X?
9. ...
10. Status of "in-flight" tickets

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

Fedora 4.7.1

- Released

Release Planning

- Ideally we get to a point where we can project out a date and high level features for future releases
- There are several small breaking changes after 4.7.1
    - Removal of test namespace
    - RDF 1.1 upgrade
- Given the status of formalizing the Fedora API and changing the codebase to align with it, do we want to have a major release beforehand that includes some of these small breaking changes?
- There is no technical reason why we haven't moved to semantic versioning - just waiting for an appropriate milestone such as the formalized Fedora API
- The community wants only 1 major release per year, but as many critical patch releases as needed. These would be back ported to recent releases
- Fewer releases would mean some people running snapshot releases to get the bug-fixes as they become available
- As master diverges from 4.7, back ports will become more complicated
- Some projects have frequent releases but include a stable branch for running in production
    - e.g. 4.7 branch, 4.8 branch with small breaking changes listed above, 5.0 branch with API spec
    - Supporting 3 different branches would be difficult given the size of the Fedora developer community
- Unknown User (acoburn) recommends separately versioning all modules and releasing them with some frequency
    - Dan Coughlin relates Hydra experience of trying to manage many different gems with many different versions. Difficult for new developers, hard to talk about what version of the software we are all working on
    - Unknown User (acoburn) says we could structure the implementation separately from the interfaces, and implementations write directly to the interfaces
    - We currently have a lot of modules that don't see changes from one release to the next but still see incremental version changes because Fedora is released as a monolith
    - Daniel Lamb says they deal with this issue with Islandora. Lots of modules, everything released in lockstep twice annually. Changes that touch multiple git repos can cause issues where boundaries in the code are not well-defined
        - For CLAW, each module is in its own git repo, and they will slice much more frequent patch releases
        - Everything will be tested together as a large community effort at particular points in the year
        - Properly defining the boundaries is the hard part
        - From a management perspective, it is better for modules to vary independently and avoid monolithic releases
    - Andrew Woods says this would be fantastic
        - Worth keeping management issues in view
        - Supports this refactoring effort. Need to figure out how this would get done

            - Might be some lessons to learn from Unknown User (acoburn)'s Trellis

Fedora/API-Spec delta-document

- Aaron Birkland working on this
    - Has a document defining scope and audience of effort
    - Who is this for and what do we want?
    - Line by line analysis vs. high level overview
    - Google doc, anyone can contribute
    - Next step: move to wiki page and solicit input from the community in terms of what we're looking for
    - Assemble team to divide and conquer
- Near-term benefit: helping with development, clarify tasks to be performed on Fedora implementation to align it with the spec
- Line by line comparison will help generate actionable tickets
- Aaron Birkland will move the doc to the wiki and send a message out to the community for feedback, ideally by the end of the week
- Andrew Woods looking for stakeholders to participate in this effort
    - Delta doc, completion of specs, aligning Fedora implementation with spec
    - Dan Coughlin (doc review), Jared Whiklo, Daniel Lamb

Fedora API Spec

- Are there any changes/updates that need to go in before we push this out to the community for feedback?
- Targeting May 1 for completion of feedback and initial release
    - Daniel Lamb concerned the deadline is too far in the future, may not get any feedback until close to deadline
    - Maybe move deadline up to April 1, or April 2 to avoid April Fool's Day
- Do we need multiple implementations of each component of the spec in order to make the spec official?
    - Seems like there are implementations already in progress in the community that can address this concern
- Community message: we want to reach an endpoint by the beginning of April. Short deadline for editorial changes, then start working on bringing implementations in line with spec. Need 2 or more implementations in order to make the spec official.

Focus on intra-repository-resource-relationship performance

- Recent email traffic on Hydra list on this problem
- Deeper dive phone call after this one to move this effort forward
- Has been effort in the past but nothing that has been completed
- Two approaches:
    - Storing these relationships as URIs (Prefer:InboundReferences and referential integrity go away)
        - If we drop referential integrity (persist URIs instead of references) would this negatively impact implementers?
        - Daniel Lamb: Break everything down to per-resource requests
            - Allows users to create bad/dead links, but this risk is worth it
        - Benjamin Armintor: Hard to characterize in a one-size-fits-all fashion - more important in some circumstances than others
            - If synthetic triples of different containment interaction models are not stored than this is not such a big deal
    - Optimizing the backend database index