

# 2016-08-11 - 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](http://irc.freenode.net)

## Attendees

- [Andrew Woods](#)
- [Longshou Situ](#)
- [A. Soroka](#)
- [Unknown User \(acoburn\)](#)
- [Esmé Cowles](#)
- [Yinlin Chen](#)
- [David Wilcox](#)
- [Katherine Lynch](#) ★
- [Nick Ruest](#)
- [Benjamin Armintor](#)
- [Bethany Seeger](#)
- [James R. Griffin III](#)
- [Aaron Birkland](#)
- [tamsin johnson](#)
- [Allen Flynn](#)

## Agenda

1. New Fedora committer: [Aaron Birkland](#)
2. [Fedora 4.6.0 \(Last MadeShape4\)](#) - release testing status
3. Weak/Strong E
4. Recent v
5. Review I
6. Review ar
7. New bug
8. Work on
9. ...
10. Status of in-flight tickets

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.

## Ticket Summaries

1. Please squash a bug!

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. Tickets resolved this week:

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.										

## 3. Tickets created this week:

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.										

# Minutes

- Congratulations and welcome to New Fedora committer [Aaron Birkland](#).
- [Fedora 4.6.0](#) (Last ModShape4) - release testing status
  - Testing is ongoing for Fedora 4.6.0 (last Modshape 4 release). [A. Soroka](#): thanks to everyone who has been helping with testing, and thanks to those doing the tedious work of manual testing in web browsers.
- Weak/Strong
  - See 

Unable to locate Jira server for this macro. It may be due to Application Link configuration.
  - See
  - The LDP spec says we should use if-match when performing an update to make sure it does not overwrite another update. The HTTP spec currently does not allow weak e-tags on update. Currently, if a client uses strong e-tags, If-Match works. If the client uses weak e-tags, it returns "condition unsatisfied." [tamsin johnson](#) commented on this issue, Fedora attempting to use strong validators: <https://github.com/fcrepo4/fcrepo4/pull/1089#issuecomment-23918773>. There is a stringent list of requirements for strong e-tags. Not sure if there is a way to meet them for RDF unless you implement some method of internal version control.
  - Proposed solution: use the timestamp to indicate updates in a method similar to e-tags functionality. This provides a workable alternative to the e-tags If-Match method of avoiding conflicts. This recommendation could go to the LDP community and be offered for consideration as a proof-of-concept workflow for RDF resources to protect edits.
  - Any LDP clients that are using the recommendation and are written to spec, when updating resources via PUTS, would be incompatible with Fedora in this implementation. However, Fedora is already incompatible. Any LDP client that is trying to faithfully follow the spec won't be able to make updates unless they attempt If-Match, recognize the failure, and try another method, i.e. the header method. If a to-spec LDP client gets a weak e-tag in the response, it shouldn't accept it for validation. There is not currently a way for things to Just Work as they exist.
  - The software in its current state allows you to accept a weak e-tag and scrub the W/ to pass it along as a strong e-tag. This is a loophole that should be closed, but there is a concern about blocking off functionality to do so. Correcting client behavior on the server side should be a policy-driven decision, not a software-driven one. The PR's purpose is primarily to make this more explicit with a more useful error message, and close a loophole.
  - Additionally, the way e-tags are used is not a back-and-forth, but two stateless requests that take place at two separate times. Scrubbing the W/ from the weak e-tag doesn't make it match the spec.
  - In sacrificing this, we sacrifice caching and weak concurrency. If we no longer accept weak e-tags (do we? have we ever?), the weak concurrency is taken up by the If-Not-Modified header, which is practical. Caching is all right as well because the tags are still out there for utilization.
  - The LDP community discussed in the RFC-7232 process that weak e-tags should be allowed. The spec processes for these two things overlapped in a way that caused LDP to make an impracticable recommendation that is under review. LDP says server should require If-Match, but if they're not sending strong e-tags (which no one has figured out how to do), it makes it impossible to require If-Match. This may hopefully be resolved in discussions in the LDP group.
  - Immediate-term resolution: do nothing. Raise this with the LDP group and see if we can come up with a better solution in the next few weeks. We may be overlooking a nasty edge case by leaving this in for now, but let's try to come up with a bigger solution that's generally accepted, and implement that in the near term.
- Recent v
  1. 

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

2. Unable to locate Jira server for this macro. It may be due to Application Link configuration.
- - Review Import/Export [design priorities](#)
    - [Nick Ruest](#) has been spearheading. Looking for consensus on proposed priorities. If you have not yet done so, please review the design priorities.
    - No pushback from those who responded, just consensus so far. Scheduling the sprint for Phase 1 priorities soon.
  - Review and contribute to "[Fedora value proposition](#)" statements
    - [David Wilcox](#) has put together a voting system with a variety of phrases that describe what Fedora is/does.
    - People are encouraged to vote on the ones that resonate with them, add comments if any.
    - Timeframe for voting and submissions: within the next week or so. Folks on a Steering Group call on Tuesday will be going through the data that's there.
    - The goal of this is to figure out which phrases Fedora stakeholders find the most important. Statements are about "what is Fedora?" and not "what could Fedora be?"
  - New bug: Unable to locate Jira server for this macro. It may be due to Application Link configuration.
    - T if this can be caught at the
  - Work on Unable to locate Jira server for this macro. It may be due to Application Link configuration.
    - through most of the nodes to return the ones needed. To address the issue, [Widesshape](#) queries need to be optimized to not do this, targeting the actual resolution to use query resolution facility where possible, to cut down on node iteration. [Ben](#) has done some work around this and had some positive performance findings. [Esmé](#) trying to replicate [Ben](#)'s performance findings without success. This could mean the performance is impacted by something deeper in the stack.
    - Second subticket: preventing additional jcr:node lookups with data we already have (was this correct?)
      - [Esmé](#) & [Aaron](#) working on this. Looking for volunteers to help with this work on in its current state.
      - Two complexities: 1) the path and URI are not a perfect match, 2) the node lookup for the subject is not at issue, the lookups of N child nodes are. In order to put an identifier in the triple, you have to get the node to get that identifier. This work aims to make it so that the application does not need to get the node to get the identifier for child node references.