

2021-04-08- 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/j/396037251>
 - Meeting ID: 396 037 251
 - Find your local number: <https://lyrasis.zoom.us/u/ad6Xb7q3ia>
 - Join fedora-project.slack.com on the "tech" channel
 - Self-register at: <http://slack.fcrepo.org/>

Attendees

1. [Danny Bernstein](#)
2. [David Wilcox](#)
3. [Jared Whiklo](#)
4. [Peter Winckles](#)
5. [Daniel Lamb](#)
6. [Ben Pennell](#) ★
7. [Thomas Bernhart](#)
8. [Calvin Xu](#)
9. Mike Ritter

Agenda

1. Announcements
 - a. [2021-10 Camel Toolbox Sprint](#)
 - b. North American User Group
2. Pilots / Testing
3. Pre release
 - a.
 - b.
 - c.
 - d.
 - e. Unable to locate Jira server for this macro. It may be due to Application Link configuration.
4. Any new
 - a.
 - b. Unable to locate Jira server for this macro. It may be due to Application Link configuration.
 - c.
5. Your topic
 - a. Modeshape CI <https://github.com/fcrepo-exts/modeshape/pull/3>
 - b. fcrepo-java-client question
 - c. migration util still have performance issue after tuning java heap size and turning off checksum validation

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
 - North American User Group
 - [2021-04 Online Fedora Users Group Meeting - North American Timezones](#)
2. Pilots / Testing
 - Performance testing by Andrey, Peter and Jared
 - Andrey and Jared were seeing fairly bad performance
 - Andrey has performed the same tests on Fedora 4 and 5, and is seeing significantly worse performance on 6
 - Will check to see what the specs of the machine running the tests were
 - Peter and Danny have been unable to reproduce the poor performance
 - Including search index commits in same DB commit as the other db commits seems to maintain normal performance
 - Search queries are fast except when getting back RDF types. Danny is working on resolving this in a separate PR.
 - Comparing POSTs and PUTs - Jared reported that PUTs were a little bit slower
 - Single object vs Multi-object transactions
 - At the moment all transactions assume they are multi-object
 - It may reduce the database latency significantly if single object transaction commits were implemented, but at the cost of considerable complexity
 - what percentage of the latency is due to database? To determine how much of a concern this is.
 - Calvin testing migration-util with larger heap, checksum turned off
 - Initially faster, slowed back down to 25k per day after about 1.5 million objects
 - Previous runs used about 30% cpu, but not it is only using about 7%
 - It seems to still be processing image files, which are not particularly large
 - UVA originally ran their migration before alpha2, had significantly slower performance in their second migration after alpha2
 - Ben C created a ticket to look into slow resumption of migration-util, which iterates back through all the Fedora 3 objects.
3. Pre release short list of bug fixes and improvements
4. Any new tickets/issues to be considered for the release
5. Other topics
 - fcrepo-java-client
 - Need to follow through on java 11 update and update of integration tests.
 - Thomas will take a look at the existing two tickets for updating it
 - Modeshape
 - Danny will look at the gh actions PR
 - Will update Seth's PR to run gh actions to make sure tests pass
 - We will review the unreleased updates to modeshape since the last release was cut, to determine if we want to sync up our fork
 - ocfl-java database
 - Peter will be checking on its usage of postgres for inventory caching and distributed object locking, as there are additional costs incurred by its use of separate processes per connection that he had not realized

Actions

- Look through migration util commits since December to see if there are any potential causes of performance drop off.