

2021-04-15- 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. [Michael Ritter](#) ★

Agenda

1. Announcements
 - a. [2021-10 Camel Toolbox Sprint](#)
 - b. North American User Group Debrief
2. Performance issues needing attention
 - a. Short-lived
 - i.
 - ii.
 - iii. Unable to locate Jira server for this macro. It may be due to Application Link configuration.
 - b. Migration
 - i. Do we know the problem is neither memory nor IO bound?
 - ii. Are there JVM tunings that we haven't tried?
 - iii. Is a heap dump likely to help us?
 - iv.

micrometer: Unable to locate Jira server for this macro. It may be due to Application Link configuration.
3. Pilots / Testing
4. Pre release
 - a.
 - b.
 - c.
 - d.
 - e.
 - f. Unable to locate Jira server for this macro. It may be due to Application Link configuration.
5. [Marmotta](#)
6. Volunteers to be the maintainer on <https://github.com/fcrepo-exts/migration-utils>
7. Any other new tickets/issues to be considered for the release
8. Your topic

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
2. Performance Issues
 - a. Short live transactions
 - i. More of an impact on reads than writes, e.g. getting the root resource of your repository with 1m+ items
 - ii. Ordering adds more latency
 - iii. Ongoing work for short lived transactions
 1. Long lived transactions perform full join
 2. Short transactions perform a simpler query
 3. Could be expanded to writes
 - iv. Search Index
 1. Current work is for a synchronous update
 2. Question of how much of a performance impact there is when updating the index
 3. Possibility of adding config to update async through event bus from initial discussion
 - b. Long running transactions in bad states
 - i. If changes are already committed to the ocfl layer, changes can't be rolled back
 - ii. Might be able to attempt to rollback, then mark transaction as failed
 - iii. Can also prevent the transaction from being committed at all
 - c. Migration utils
 - i. For pid list: need to stop iterating the F3 repository once we've processed all items
 - ii. metrics: Try and capture bytes/sec

Actions

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