

2016-04-28 - 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

- [Jared Whiklo](#)
- [Bethany Seeger](#)
- [Andrew Woods](#)
- [Allen Flynn](#) ★
- [Namita Bahulekar](#)
- [Longshou Situ](#)
- [Unknown User \(acoburn\)](#)
- [Nick Ruest](#)

Agenda

1. [4.5.1 RC-3 testing status](#) - and release plans, this Friday!
2. [Release candidate policy](#) - feedback on wording?
3. F4/Modeshape5
 - a. [PR status](#):
 - i. Good
 1. Mode5 PR works
 2. Config: File works
 3. Config: MySQL works
 4. Removed default `fcrepo.modeshape.configuration`
 - ii. Bad - not so bad
 1. Config: PostgreSQL not working - but likely easy
 2. Config: Clustering - untested, but should be able to work
 - a. Binaries are copied across instances
 - b. Use central object db
 - b. Migration testing volunteers needed
4. Review [Versioning Spec - meeting](#) @2pm ET, Tues May 3rd
5. Thoughts on multi-tenancy
6. ...
7. 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

(1)

4.5.1. RC-3 testing status

Release is this Friday!

email will be used for communications tomorrow

Question about BaseBox - update to Basebox discussed

(2)

Release candidate policy

Feedback received from Fedora committers list

Comments requested

This policy makes it clear to the community how long we will test releases and which email lists will be used for notification, with templates to make it easier to notify going forward.

The JIRA queries can be troublesome to generate. To make sure we get the "right things" in the JIRA queries it would help to have some additional guidance on running the JIRA queries.

Suggestion accepted to work out goals for 4.5.1 release for improvements to description/instructions for JIRA query process. To be further developed once goals set.

When you put a release in JIRA, JIRA will collect up the issues for you. A future conversation is planned to discuss the JIRA queries associated with release.

(3)

Update on Modeshape 5

Pull Request (PR) Status

In general Modeshape 5 works but there are a small number of issues. There are a handful of configurations (example configs) that people use in the `fcrepo.modeshape.configuration`

There is a default for LevelDB, but we are not using LevelDB anymore, so that goes away. Similar to Modeshape 4, Modeshape 5 allows you to have two different storage locations, one for the "objects" and one for the binaries. The binaries by default go to the filesystem as before. There is also a configuration for storing the objects just on the file system and that is some compressed "proprietary" Modeshape way. That seems to work fine. Likewise the objects like what we have had, there is a configuration where the objects are stored in MySQL which seems to work just fine.

We are advocating, it seems, that people use one of the database configurations for their Fedora installs going forward.

Goals: (1) Easy for people to spin-up and test Fedora. (2) Want people to use a configuration that is intended when people spin-up for PRODUCTION.

When using the MAVEN-JETTY plug-in, it is set up to use the file config. However, when using 1-click or Tomcat, those do not have a system property provided and therefore those DO NOT spin up and the you get errors saying Modeshape configuration not set.

So 1-click is no longer 1-click as a result.

For 1-click, we are trying to keep the bar low so people can spin up Fedora and play with it, so it is a goal to update the 1-click and/or its MAVEN plug-in to afford 1-click. A way to do this would be to have a default value hard-coded in to the `repository.json` file OR `repo.xml` file that tells which `repository.json` file to use. Right now it pulls in the configuration that is there, however a change here will need to be figured out now for Modeshape 5.

More Issues with Modeshape5 [BAD or NOT SO BAD](#)

PostgreSQL is not working. Seems like it should. If someone familiar with PostgreSQL (version 9.5, 9.4 were tested) and its drivers could take a look, that would be helpful. The error noted is a driver error that we did not used to see.

Modeshape 5 does have clustering. This is untested. Infinispan was set-aside in part because of clustering issues. Modeshape 5 says its ideal for "non-heavy" binary installations. The implication is that binaries get copies across the instances in a blocking fashion, and so uploading may be (too) slow. Modeshape also suggests that one use a centralized database configuration for the objects. There is a clustering capability. There are some constraints. The constraints seem to be reasonable for performance reasons. A configuration has been added but not tested.

Servlet-auth: This has not worked for years, perhaps since we pulled out the "RBACL" code. Suggestion made to remove servlet-auth from main code base. Are there issues removing servlet-auth from the main code base? (None raised.)

*Bug in Modeshape - Modeshape would start giving integration test failures (one unit test failure) and may be attributable to a recent Modeshape update ... at Modeshape you can specify when Modeshape does garbage collection, maybe for binary resources that are no longer referenced where binaries go when they get deleted. By default the garbage collection takes place at Midnight. So, during the testing, the time calculation gets messed up and so an assertion in line 1995 throws an error. Seems like there should be a simple fix for this bug. It is suspected that a unit test of the 24 hour clock for garbage collection would provoke this bug and cause the error.

*Federation not refreshing when new files added. Federation works but when you add files to the file system it does not seem like caching logic is the same.

Summary - there is a pull request targeted at master so people can look at it and play with it at this point.

Question - is the Modeshape 4 series still supported? (Unknown.)

Question - is the Fedora Tech Community going to support Modeshape 4 backends going forward? This seems to depend on how long Modeshape will support Modeshape 4.

> Wondering whether the Fedora Tech Community supports a F 4.x or new F 5 versioning and related support for Modeshape 4 and Modeshape 5. Before a Modeshape 5 branch of Fedora gets started, the two questions above need to be answered.

Question (reiterated): What thoughts do we have about how Modeshape 5 gets rolled in and supported?

> In the simple case where this is rolled over on top of the Modeshape 5 implementation, could this be a 4.6.0 release? Or is that too soon?

Significant testing needs to be done with Modeshape 5.

Modeshape 5 will require a migration for everyone running Fedora 4. It would be good if people became available to do some testing of their own Fedora instances. For that reason, 4.6.0 may be coming too soon to expect the migration to Modeshape 5. The issues reviewed today need to be worked out before migration can take place. There is some time involved to do the work needed to prepare for a Modeshape 5 migration...

(4)

Versioning Spec Review Meeting has been set for Tuesday, May 3rd, 2pm ET. All are reminded to read the spec and invited to attend the meeting next Tuesday.

(5)

Thoughts on **Multi-tenancy**

Question: Are there specific use cases?

Consider Hydra-in-a-Box ... this initiative is pushing toward being able to implement Hydra-in-a-box in a multi-tenant situation. Recognizing that there is some interest and others are working on aspects of this multi-tenancy and that Morris' team at UMichigan is interested too, this topic is being surfaced.

Brainstorm: If there are ideas or models for multi-tenancy support, let's discuss and consider them.

So, if different regions are set up, can resources be looked across tenants (across regions). The assumption is operating as if each tenant had its own repository.

So, would we have multiple resources without parents? A variation would be pushing the resources down a level so that a root exists that only the "global" admin has access to.

At present, when there are binary resources and object resources, there is a single directory. So, with multi-tenancy, would completely separate binary file directories be needed. This seems like a potentially desirable requirement. Consider costing, security, etc.

In the ModeShape context, the notion of workspaces has not been discussed and it seems like workspaces would support this file segregation. Separate storage locations could also be configured. This requires further thought and exploration.

Consider soliciting use cases with a lot of detail at this point to help guide the discussion on multi-tenancy, recognizing that generating use cases may not be enough to understand a full requirements set for a variety of reasons.

NEXT STEPS and ACTIONS

- a. Move forward with ModeShape5 - who is interested in collaborating on the PR and issues listed above re: ModeShape 5?
 - PostgreSQL - Jared W.
 - Clustering -
 - ModeShape bug - who can/will identify this bug for ModeShape, perhaps with an associated unit test? - Longshou S.
 - 1-click run issue and MAVEN - address problem of having 1-click run with 1-click - Aaron

Possible Solution Way 1 - hard code in a default configuration

Possible Solution Way 2 - use MAVEN profile to adjust the configuration

For today, we will defer going through "in flight" tickets. Are there other comments or issues to discuss?

For next 2 weeks, a placeholder agenda page will be created and the dialog can continue. Jared W. or Aaron will run the agendas x 2 weeks.

<call ended 11:54 ET>