

2009-02-17 - Committer Meeting

Present: Andrew, Bill, Eddie, Chris, Kai, Gert

Welcome Gert, longtime committer. (Introductions)

Andrew:

- Has been working with Clover, a coverage tool from Atlassian. Dan is looking at licensing. Being open source/non-profit, we can qualify for free licenses. Will have clover running during the build process. Outstanding question: how we get a bunch of per-machine licenses at once? And how do we distribute them to committers? Eddie asked: does it work with ant? Has seen with maven. Andrew: yes.
- Experimenting with Derby. Two ways: one is standalone, the other is embedded. Main interest is in doing it in embedded mode, standalone is nice if it's easy. Working on DDLConverter, looking at installer next.

Bill:

- Continuing with Admin UI work. Lately been stalled by issue with file uploads with Flex. Unfortunately a little different based on how browsers work. With AuthN required, problems. Off, things are ok. Firefox and other non-ie browsers will send their ingest requests in a separate session. Wondering if can pre-emptively authenticate. Most other pieces of the Admin UI are now working.
- REST API: Noted and fixed rest api issues. Object properties now includes state of the object. URL for the datastream info... can request datastream.xml. API-M getDataStream has been added to REST API. CreateObject also now returns PID in the response, in either case. Thus, there will be several API changes with next release. Will ping Eddie and Matt on the API changes when finalized.
- Has been sitting in on discussions with Dan and Sun to set up test harness. Current plan is to have 3 Sun machines. One running Fedora, one running MySQL, one running the test suite. Goal is to get higher performance. Try to determine where there are benefits to having certain types of hardware. Also seeing what can change hardware-wise to optimize things. Sun wants to figure out what the recommendations are for people using Fedora on Sun hardware. 32GB ram, 10-100TB disk. Will have the lab accessible to Fedora Commons. Kai: interesting. So far, our tests have been on commodity hardware. Glad to help with input. Wonder if it would be useful to create benchmarks. If there's one set of data we can use across environments.
- Also fixed fedora-find bug.

Chris:

- Was at dev8d last week, great event.
- On Friday, gave talk on Fedora basics and upcoming features/interesting stuff. Pointed to wiki (now publicly visible) to orient new developers /committers: <http://fedora-commons.org/go/fcrepo> With the new website, this is where the link for "Fedora Developers" or similar will go.
- Asger Blekinge-Rasmussen (State and University Library of Denmark) gave a presentation on his Enhanced Content Model work: <http://fedora-commons.org/confluence/x/SICn> and it looks very promising. They have also developed an interesting web-based object builder tool (they contracted out to do the work, but will be making it available via Apache2). Want to help as much as possible on getting the community's eyes on this and getting their feedback. We discussed the possibility of bundling this somehow, or making a link to it prominent on a future Fedora release as a means to get more eyes on it. It probably won't be ready for any kind of bundling in time for Fedora 3.2 (early spring), but likely sometime after OR'09. Still more discussions to come on this. To aid in his development, Asger has also written an improved set of documents on API-A and API-M and is now helping to port it over to our site so others can benefit.
- Matt Zumwalt gave a demo of MediaShelf's ActiveFedora. This is similar to Ruby's ActiveRecord concept, and the basic idea is that it makes working with Fedora objects very easy for Ruby developers. The model of what you're working on is specified in ruby code, and the idea is that it could be (but isn't currently) persisted in a Fedora content model object as a description. Matt has also had some discussion with Ben O'Steen who is doing a similar thing, but with Python (possibly using JSON as the format in which to describe the model?). They had been thinking that it may be possible to converge on a single format (rather than making it Ruby-only or Python-only) and will be investigating this more. There was also some good discussion on the differences between the ActiveFedora approach and Asger's work. It was noted that they came at it from very different angles; Asger thinking more about adaptability and use in many contexts, and MediaShelf thinking more about ease-of-use when putting together UIs quickly. Matt is also hoping to have a "Fedora-backed Application in 15 minutes" demo for OR'09.
- Opened FCREPO tracker for public submission of bugs/features. **Action: Still need a way to mark "verified" status on items.** For the time being, we can look at dates and submitter username. **Action: Move remainder of FCSUPPORT items pertaining to repo software over.**
- Applied Renauld Waldura's patch to fix FCREPO-407

Eddie:

- Just talked to Glen Robson. Did orig SWORD impl for Fedora 2.2. Looking to finish the work to a) upgrade to SWORD 1.3 and work with Fedora 3.1 by end of month. Currently nothing in their subversion on that project. Sword folks want to move toward using AtomPub.
- Mulgara 2.0.8 update, looking fine before conf. On hold, what to do about native store. Use new XA1.1 as default? Not as well tested, will force a rebuild. For now, maybe makes sense not to worry about it. Can go either way at last minute if we need. Will probably be another point update of Mulgara. Choice btwn old and new storage is a config in the jar.

Gert:

- Mostly listening today. Currently has two projects using Fedora and GSearch.
- Likes having public committer meeting notes

Kai:

- Main interest is in performance/bottlenecks.
- Found spot in Fedora code where a sql query is triggered and runs for each authentication 8-10 times. More detail here: <http://fedora.fiz-karlsruhe.de/docs/Wiki.jsp?page=Sql> **Action: Chris will enter in tracker.**
- Also found Fedora triggering GC after each ingest. Has done large ingests and monitored heap. Gets much better without explicit triggering. The 1.5 JVM really does a much better job at GC. No problems. Mote detail: <http://fedora.fiz-karlsruhe.de/docs/Wiki.jsp?page=Java%20Heap%20%26%20GC%20Tuning> **Action: Chris will enter in tracker.**