

SW Arch Meeting Minutes - September 15, 2011

Agenda

- Announcements and Updates
- Progress Updates
 - R5.2.1
 - R5.2.2
 - R6.0 - New Platform
 - R6.1 - SOLR/Lucene
- Software development requirements for Fedora 3.4.2 (Postponed)
- Update of /rarch specification
- Continuation of AV objects with transcripts discussion

Announcements and Updates

In the recent CISC meeting, there was extensive discussion of an NSF grant and our ability to participate in this grant. The grant is entitled “ Metadata for Long-standing Large-Scale Social Science Surveys” and the solicitation seeks proposals that will develop tools to bridge data collection and dissemination, first concentrating on future waves and then migrating legacy waves into standard formats. In general, CISC felt we should go for this grant although it will be a “stretch” and there was some concern that undertaking this grant would overwhelm our capacity to deal with more near term issues. In the end, we felt that we will have to solve many of these problems anyway so the grant is worth pursuing.

Progress Updates

Release R5.2.1. Kalaivani reported on this bug fix release. Most of the bugs are fixed and Rhonda will contact the faculty member to test faculty collection creation. We expect to finish testing next week and will plan for a release on Sept. 26. Although downtime will be minimal, Isaiah will put out an announcement.

Release R5.2.2 (WMS Only). This release is focused on WMS only and Yang reported that code complete for bug fixes will be at the end of September. We decided to move the /rarch capability for objects LT 2 GB to a later release since it is not urgent. This capability could be put in the large file release (R6.2) or in an earlier release if it is ready. For this capability, Yang will implement another configuration option that indicates whether or not the /rarch directory should be used. Although we acknowledge that the testing interval will be tight, the public release target is October 15.

Release R6.0 (New Platform). Dave and Sho reported on the progress to achieve sanity with the new platform. The urls for testing have been distributed and initial testing shows promising results. There was a problem with Djvu which was fixed by going back to the 32-bit version. Jeffery will check with the vendor to see if we can get a 64-bit version. There are style sheet problems in the dlr/EDIT

search and Jeffery will switch to the no-parse version that is already in place and working. The new server hand-shakes with Darwin, but we will need to test the complete video ingest and play scenario. With additional testing, we are hoping to declare sanity next with for this environment.

Sho indicated that Fedora 3.5 does not really offer any major new features that are useful and there are no critical bug fixes in 3.5. In particular, we do not want to move to Akubra and should verify that it will remain an option in 3.5 and later. As a result, we will stick to the more conservative approach and upgrade to Fedora 3.4.2 which should be more stable. In order to prepare for Fedora 3.5, it would be good if we could install it somewhere so developers could begin to test on this platform. Sho will verify Shibboleth, API-A and PHP changes on Fedora 3.2.1 as a first step. This environment should be ready for developer sanity testing by September 28. Given the upgrade in PHP, we will need to check out each PHP module to make sure there is no impact from changing versions. After this test is complete, we will move to Fedora 3.4.2 (target date TBD).

Release R6.1 (SOLR/Lucene). Rhonda indicated that the vetted metadata will be available by the end of September and that we will be able to discuss the SOLR/Lucene requirements in the next sw_arch meeting (Sept. 29). She also indicated that developers are participating and making changes as they are proposed in the user services group. Chad and Jeffery will need to update the architecture specification to reflect all the changes in WMS, dlr/EDIT, and the portal software. We will also proceed with Analytic version 1.1 in this release.

Update of /rarch Specification

Isaiah discussed the addition of two new bullet items to the specification that deal with how user-uploaded content is handled and the purging/clean-up of files after ingest. Both WMS and dlr/EDIT will need to address this issue, including video files on the Darwin server. Regarding upload, the new file system may allow us to simply rename rather than do a copy/delete. Dave will check to see if this is feasible. Isaiah will do one more update to the specification and then it can be considered “final”.

Agenda Items for Next Meeting

- SOLR/Lucene requirements – Rhonda
- Architecture update for SOLR/Lucene (Chad, Jeffery)
- Report from Sho on software development changes for Fedora 3.4.2
- Continuation - Video/audio objects with transcripts
- Getting started on the large file (GT 2GB) specification
- Review of the url migration implementation specification