

**Software Architecture Working Group
Minutes of November 12, 2009 Meeting**

December 7, 2009

Agenda

1. Updates and Announcements
2. R5.1 status and performance
3. Statistics API for R5.2
4. Sound disseminator for R5.2
5. Private/public portal specification

Announcements and Quick Updates

The annual NJEDge conference will be held November 18 – 20. Isaiah and Ron will be participating and talking about the installation and configuration of RUcore for NJcore. Local file ingest will be a part of Fedora 3.3. We will need to decide when to upgrade from Fedora 3.0. Sho is investigating a number of PHP extensions (e.g. for character translations). We need to schedule the migration to the SOLR/Lucene search engine but will first need to review the work involved and performance characteristics before proceeding. We have some money to purchase mass storage this year and will need to submit a PO before June, 2010. Given the storage management issues we are encountering on lefty64, Ron asked Kalaivani and Isaiah to develop a plan and process for managing storage on lefty64. The process should include as a first step to give us some “breathing room” on lefty64. For ongoing management, we should assign owners to the major directories who will delete files that are no longer needed.

Status of R5.1 and Performance

Kalaivani reviewed test status, indicating that there is a bit more testing to be done with faculty deposit and marc export. It is noted that java bridge will no longer be need – Jeffery will update the configuration document. Regarding performance, Jeffery rebuilt indexes using only the bibliographic data (i.e. no full text indexing). Performance on lefty64 (R5.1) looked quite good when searching long unquoted strings. Initial searches ranged from 4 to 6 seconds and subsequent searches were significantly less than 1 second. Dave has built indexes on mss2 using amberfish 1.4.1 and he will compare performance results with amberfish 1.6.x.

Statistics API for R5.2

Jie reviewed the statistics API. We'll need to address or reinstate the top node of the collection architecture in order to get stats for all collections. Jie will investigate adding in performance stats for the search engine. The specification is ready for posting and Jie will review the user interface at the next sw_arch meeting.

Sound Disseminator for R5.1

Jeffery reported on the sound disseminator. For R5.1, we will resort to a simple play back with no segmentation or structure map. Jeffery will post the specification. We will rewrite the disseminator for R5.2 to support segmentation.

Private/Public Portal Specification

Chad reviewed the specification for privatizing portal access and creation in R5.1. Details are in his specification which will be posted.

Other Items

We discussed problems with ETD embargoing and symlink deletion. For R5.1, we will check if shib is disabled and if an xacml policy exists, in which case no link to the pdf will be created. For existing ETDs, we will replace PDF1 with EMBARGOPDF-1 and add the xac ml policy that restricts this pdf. The PDF1 is then deleted. (Note, someone will have to restore these files to PDF1 after the embargo period.) Dave may want to delete symlinks but this, strictly speaking, is not necessary. (What about the export of embargoed ETDs to UMI?)

Agenda Items for the Next Meeting

- Status of R5.1 and performance
- Statistics user interface
- VMC video ingesting
- R5.2 review

rcj – 12/07/2009