

SW Architecture Working Group
December 17th, 2015

Present: Kalaivani Ananthan, Isaiah Beard, Jie Geng, Dave Hoover, Chad Mills, Aletia Morgan (recorder), Jeffery Triggs

Agenda

1) Announcements - All

- Rhonda is away today. Aletia will take minutes in her place.
- 2016 meeting schedule. Agreed to start biweekly on 1/14/16. Aletia will do Zimbra invites.

2) 7.7 Rollout Review

- Everything seemed to go well with the update yesterday (12/16/15). RUcore was only down three hours (7-10 am). All seems to be working properly, although deposits are slowing down with the holidays. Note that CDL will be closed over the holiday, so while the servers will be running and DOIs will be minted, there will be minimal technical support if needed.
- One concern with the rollout testing is that a significant bug was not identified until it reached the staging server. Kalaivani wants to improve the testing process to avoid late surprises like this. How did a bug that appeared to be fixed on test travel to staging? An example is a bug in mapping the notes field in scholarly deposits. Is there any way to test more effectively? Yang asked if we need better specs for the testing plan that includes more specific step-by-step directions. Kalaivani asked who can (or will) write them - the developers, or the stakeholders? There is no easy solution. The changes in each new version come from multiple sources, and so the specs are scattered. Perhaps we could use some sort of cross-reference chart to show all components that have been modified and need to be tested with any new feature, with the stakeholders defining a test and participating in the testing (as Jane Otto did this time). Developers should be responsible for defining the testing protocol of more arcane technical features. Essentially, we are missing the organizational level of “analyst”, but have to make things work more effectively if possible. At this point, specs are reviewed at USG, but they may not be detailed enough to allow for clear testing protocols. So this again argues for asking the stakeholders as part of the process to develop test plans with Kalaivani.
- Specific suggestions:
 - o Include a link in software.libraries to the spec document
 - o Require stakeholders to request feature changes to specifically test these features
 - o Attempt to identify any additional possible scenarios that could affect the testing outcomes (e.g., batch uploads vs. single object uploads)
- Maybe the new bug tracking platform will help prevent such miscommunication.
- Chad noted that ultimately though, we've had a very successful year. We've rolled out 4 major and 4 patch releases in 2015.

3) RUcore 8.0 & 8.1 planning – All

- The 8.0 Fedora upgrade is set for January and February. 8.1 code freeze is set for March. During January and February we can start reviewing 8.1 features. Meetings in February are the 11th and 25th. Deadline for specs for approval is the Feb. 24th (CISC) or Feb. 25th (SW Arch). We have 4 meetings to review specs.

- Recall that 8.0 is essentially the 3.8 Fedora upgrade.
- Dave noted that we also need to upgrade PHP and a new service pack. PHP 5.4 goes EOL in January. The question is whether to upgrade to PHP 5.5 or 5.6? 5.6 is the most recent, but 5.5 is supported only for 8 more months. We need to determine whether we will attempt to make that upgrade along with the Fedora 3.8 upgrade, or do them separately to be safe.
- Timing of Fedora 3.8 – looking at the second week of January for upgrade, then test, move to staging, and rollout. We need to verify the sequence of server upgrades!

4) Email Archiving Assessment Document - Beard

- This was triggered by a request from Tom Izbicki, who is retiring. The plan is to create a dark archive record of Tom's emails. It's a one off situation, but should be aware of and follow best practices. Isaiah recommends mbox (single file for all) and eml (one file per message). This plan was approved at the most recent CISC meeting. While at this time there is no RU or RUL email preservation policy, we don't know of any plans to archive any other RUL email accounts. Isaiah is working with Tom to clean out any personal content, and Tom is identifying emails to be preserved. Dave is concerned with the policy implication of preserving selected RUL email account, and that it creates an unfortunate precedent and obligation on our part. Thus, he wants to be on record as objecting to it.

5) Google Sitemap specifications - Marker, Triggs

- We need to identify a new method and create new specifications to accommodate collections like CTKe that should not be visible. Thus, we will not yet run the video sitemap on production. There is a technical issue that has to be resolved, so this will be put on the list for work as part of the 8.1 release

6) High-resolution Image Viewing Specification for 8.1 - Beard, Mills

- Some of the options for Hi-Res image viewing have turned out to be expensive (Kakadu, JPEG2000).
- Chad has identified a less expensive approach in the use of Tiled Pyramidal TIFF (pTIFF) file formats instead of JPEG2000. As with JPEG2000, under the pTIFF standard, images are broken up in to many "tiles" and then reassembled for viewing.
- Reading pTIFF files: The best open source viewer for these files (HTML5 based) is OpenSeadragon, originally developed by Microsoft.
- Another tool, IIPImage also works remarkably well and works with JPEG2000 files. IIPImage creates standard JPEG files for presentation from PTIF OR JPEG2000. By displaying JPEG as the presentation format, and using JPEG0000 for Download only, no license costs are incurred.
- We need to confirm where in the image presentation process we would adopt the OpenSeadragon viewer vs. the existing tools; from the thumbnail (.jpg) to .ptiff at the full native resolution. Browsers can open the generated JPEG files using the OpenSeadragon viewer. IF JPEG2000 were download only most users would need a tool to view these downloads, so we need to verify how much of an issue that would be.
- A key benefit of the relatively new (v 2.0) IFFF spec is the opportunity to integrate business logic into our processes. Digilib, IPimage and others now follow this spec.
- Moving forward:
 - o Plan for any changes in metadata labels and file naming conventions for these new datastreams.
 - o Verify status of mimetype for pTIFFs. IMAGE rather than pTIFF?

- Should we delete existing JPEGs? Some may be the objects of direct legacy URLs. But in the future, we should deliver the .pTIFF files when these links come in. Chad is looking at alternatives for redirects, to optimize the processing load.
 - pTIFFs can be generated from regular TIFFs using ImageMagick. Some TIFFs will need to be pre-converted to be in a consistent TIFF format for pTIFF conversion. The pTIFF files would replace any current JPEGs. This does not affect the existing thumbnails (THUMBJPEGs). The old JPEG download stats will need to be associated with the new pTIFF views. PTIF-1, etc., could be the default label with .ptif as the extension. Legacy URLs will be quietly converted from JPEG to PTIF. [The Finereader license is not completely clear. Isaiah will get in touch with them).
- Does this need to go to CISC? **Yes.**

7) Jobs & Reports – Triggs

- Moving RARCHcontent. Jeffrey reported that we have done command line ingests of files of 14, 30, and 70 GBs on rep-test. These were done with copies symlinked to a web accessible url. Dave will work on making RARCH available for a test on rep-dev and eventually on production. In order to make RARCH accessible, we need to determine the sequence for testing – Rucore-DEVEL first, and then RUcore-TEST, or another way? Jeffrey will bring a document about this to the 1/14 meeting. Kalaivani noted that once Jeffrey has verified the function, we will need to document the process.
- PDF coversheets. Jeffrey has created all the covers possible. The only ones that would not work were PDFs that had internal restrictions. He will forward a list of the handful of restricted PDFs to Rhonda and Jane.
- SQL Errors on RUcore-Staging. Dave has seen some errors. SQL is set to latin-1 vs. UTF-8, which is generating a Fedora error on a bad chapter. So there needs to be some attention to these errors. He asked developers to look at any error emails, to see if the problem can be identified. Could it be control characters, or special characters in ETD titles? This needs to be understood.