

# Database Meeting 2016-07-06

## Date

06 Jul 2016

## Attendees

Fritz Mueller, [Andy Hanushevsky](#), [Vaikunth Thukral](#), [Unknown User \(npease\)](#), [Andy Salnikov](#), [Fabrice Jammes](#), [John Gates](#), [Brian Van Klaveren](#), [Igor Gaponenko](#)

## Discussion items

Welcome Igor Gaponenko!

### L1/AP Database

- Code to instantiate an L1 database based on the proposed schema has been completed.
- Working with this now to create a simulated data stream and populate a test database.
- Need to talk to [Serge Monkewitz](#) to find out status of HTM indexing extension to sphgeom.

### Docker

- Swarm/OpenStack bug found: need to shutdown VMs before snapshotting, not just wait for cloud-init completion
- Fabrice waiting for NCSA OpenStack to come back up to finish testing of his Swarm rebase now that bug addressed
- Fabrice working on Docker usage tech note
- Fabrice curious about benefits of microservices. Fritz and Brian to investigate

### PDAC

- Brian is working on containerizing DAX services. Needs full LSST stack base container, because DAX services need Butler.
  - Found some issues with scripts along the way (non-existent jre package on debian, build fails on VM <2GB RAM, etc).
  - CI is building stack images, but they are really large (2-3GB). What is going on? Needs to be investigated.
  - Probably we will end up with separate base containers for Qserv and DAX because Qserv has a lot fewer stack package dependencies
  - Working out how configuration for DAX service containers should best be done
- Igor has talked with K-T and Greg to track down the SDSS Stripe82 catalog load data. Some is at NCSA, but some also at IN2P3?
- Igor planning to move a part of the data to IN2P3 to start experimenting with the loader script.

### Shared-Scans

- "Snail-scan" added, which will end up executing queries expected to exceed the scan period of all other scans. Code still needed to feed this scan.
- John is testing/validating memman changes from this week, concurrent with ongoing review.
- John will be adding some plumbing which tracks statistics on tasks.

### Butler

- Nate has two tickets wrapping up review. Needs to disentangle some commits, and move some work to follow-up stories.

### X-SWAP

- Vaikunth showed some baseline graphs of memory, IO, CPU usage on an idle cluster. Memory stats look anomalous. Investigating
- Tried to gather 24hr graphs under load from ruqueries.py, but services crashed after 14hrs. Also needs to be investigated.
- John has added the requested query begin/end messages to qserv logging

### Documentation

- Brian working on DAX API docs, has run into many tooling issues re. Sphinx and CI. Next stage is to get it working with Travis CI.
- Nate will be updating LDM-463 draft branch as part of his upcoming Butler commits.

### Misc

- John investigating czar-config crasher. Looks to be related to log4cxx; need to build out-of-stack.
- Andy H reports 4.4 release of XRootD looking likely in the next few weeks. After this, xrdssi branch can be merged to master.
- Andy H working on XRootD logging enhancements as well.