

# Database Meeting 2016-06-22

## Date

22 Jun 2016

## Attendees

Fritz Mueller, Andy Hanushevsky, Vaikunth Thukral, Unknown User (npease), Kian-Tat Lim, Andy Salnikov, Fabrice Jammes, Brian Van Klaveren

## Discussion items

### L1/AP Database

- AndyS has been on vacation or otherwise occupied with Atlas, so no work logged on this last week. Has been thinking about requirements for L1 /AP database input simulator, and will write these up upon returning to SLAC next week.

### Butler

- Repo refactor ticket (DM-6459) has implementation now in review.
- Design document for config persistence (DM-5548) is up for review/input.

### Docker

- Oualid has finished up cluster spin-up script and it has been reviewed. Fabrice to merge tomorrow.
- Fabrice's Swarm scripts will be rebased on Oualid's changes, then go up for review.
- Fabrice is working in Qserv container infra. next steps tech note.
- Swarm scripts are currently only work on OpenStack. To work with bare metal at IN2P3 or PDAC would require some prerequisites on the cluster. Fabrice to create tickets for setting these up at IN2P3. We'd like to test there and use the new scripts for PDAC.
- Brian has begin work on containerized DAX services. Experimenting with Jenkins at SLAC and Travis CI for container builds.

### XSWAP

- Vaikunth continues work on cleanup of log data transformation scripts. Looking to produce a few plots by the end of the sprint.
- Vaikunth would like to run some tests to see if randomized input query sets (as currently generated by runqueries.py) produce statistically stable measurements. If this is the case, we can leverage our existing tools for upcoming profiling experiments. If not, we'll need some new tools that provide repeatable query sets.
- Vaikunth will be updating and using IN2P3 nodes 124-149 for these experiments.

### Documentation

- Brian working with Jonathon to set up infrastructure for DAX API documentation.