

# Database Meeting 2016-03-30

## Date

30 Mar 2016

## Attendees

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

## Discussion items

### Documentation updates

- AndyS, Mike, and Nate so far have committed documentation updates to the draft branches of LDM-135 and LDM-463. Others will be contributing in the upcoming April sprint.

### Large results

- memmanreal updates from AndyH are reviewed and on master
- John is still investigating cluster problems
  - Lots of cluster script debuggin this past week with Vaikunth and Fabrice; new scripts working now
  - Added an instance-counting hack to various classes in search of memory leaks, but no smoking guns so far
  - Work with AndyS to add some debug logging to the proxy; this revealed that proxy is blocked by czar from handling subsequent queries until all subchunk queries of the current query are dispatched. John looking at queuing dispatch to a separate thread to work around this.

### L1 database design

- AndyS. has finished initial investigations and written them up. Serge has reviewed/commented.
- Some cross-team meetings scheduled for next week re. alert production and L1.

### VO/ADQL

- Exploratory TAP/JSON work on dbserve is in to master
- Brian needs to check in with Tatiana and Gregory about their deployment needs for this. Will discuss at DAX meeting this Monday.
- DESC has shown some interest in the TAP/JSON dbserve
- Brian has started work on a python SQL parser, based on porting/modifying the ANTLR4 grammar used by Presto so it can be used with the PLY parser framework. The idea is to use this to do ADQL validation and pre-processing in dbserve. The resulting grammar could feasibly be used as input to Bison for a C++ parser as well.

### Butler

- Serge's work on spatial query utilities for Butler has been delayed by some non-LSST work, but should finish up by the end of the week nonetheless.
- Nate now working on persistence improvements for Butler configuration.

### Vertical partitioning

- Vaikunth found that adjusting optimizer search steps over a wide range of values made little difference to either query compilation or query run times for the types of joins being tested. If these joins are representative of the types of joins we would use for object and source tables, then consensus is the testing to date is sufficient. Vaikunth to check this with [Jacek Becla](#), and write up results if so.

### XSWAP

- IN2P3 informs Vaikunth that elastic search databases have been set up. Vaikunth to begin trying to work with these soon.

### Docker

- Generalized deployment scripts were a little more complicated than Fabrice originally thought. Working now.
- We are reaching the limitations of what shmux can do for us as a primary tool on the clusters. For example, status reporting is problematic (John and Vaikunth as cluster customers both emphatically agree on this point). We should try using something like SWARM as an alternative soon.
- Yvan is investigating setting up a local docker repo at IN2P3, which may address some performance problems we have experienced using the public Docker Hub. We should consider whether SQuaRE or NCSA should host one of these for us in the mid/long-term.

### Misc

- Connection timeout fix is post-review, heading for master soon
- Serge has been bogged in some EI Cap problems on his mac. Will investigate lsstsw workflow to see if this can help.
- Qserv access in the long-term: will this be TAP-only via dbserve, or will we support direct mysql protocol to czar for external clients as well? This decision has implications for how we architect qserv. For example, if TAP/dbserve access only, then we could potentially dispense with mysqlproxy, move much more functionality directly into czar, have only a single parser, etc. But the downside would be worse tooling for debug, integration test rework, etc. Should decide on this soon...