

# Database Meeting 2016-09-21

## Date

21 Sep 2016

## Attendees

Local: [Unknown User \(npease\)](#), [Igor Gaponenko](#), [Vaikunth Thukral](#), [John Gates](#)

Remote: [Fabrice Jammes](#), [Brian Van Klaveren](#), [Andy Salnikov](#), [Fritz Mueller](#) (for first half)

## Discussion items

Fabrice:

- Waiting on NCSA machines, sent requirements to them, preparing to use puppet there, implementing prototype for qserv
- Some Swarm issues are fixed, xrootd starts now, but docker network doesn't work correctly. Some other features not available
- Ok to keep working on it till PDAC is up
- Moving data to OpenStack cluster for large scale CI, using ceph

AndyS:

- No progress on supertask
- Both AndyS and Gregory busy on other stuff
- L1DB - writing some tests on the beefy in2p3 machine, preliminary tests show better performance than regular disk
- Running on SSD and NVMe, looks similar so far

Brian:

- Code of Conduct 95% established for LSSTC-Slack
- Working on docker files for webserv etc
- Working with Steve for recipe for docker @ NCSA
- AndyH - mlock available only on new linux kernels, but container kernel is based on host OS, discussion needed?

Igor:

- Jason Alt says PDAC is online, can log in to nodes
- Ready to try to load stripe82 dataset
- Chunk size needs to be decided, asked Serge and waiting on reply
- Running some tests for settings
- Data migration ~30% complete, ~2 weeks more for 100% (IN2P3 fraction 50%, NCSA 50%), these are calexp
- Loading - considering partitioning first then using loader
- EFD discussion with Fritz/Mario on Tuesday

John:

- DM-6814 reviewed and being tested, needed updated xrootd with SID changes
- Missing rows problem - hit czar with lots of queries, it returns fewer rows than expected. Will test and push to review after DM-6814 testing
- System able to handle large load (cancel, wait etc) but only runs 4-5 FTS at a time
- Malformed queries create executive and register it, causes segfault, made ticket for a quick fix
- Some more issues that have been shelved for now
- Retry mechanism will need to be reworked

Nate:

- Composite datasets - implemented basic GET, now working on PUT and making unit tests
- Once these are working, getting it reviewed and closing ticket
- Look at further use/edge cases
- All reqs dumped in confluence page
- Jim Bosch to work the API

Vaikunth:

- Found and worked around bug in elasticsearch query results for hosts
- Have basic query throughput (QTP) quantified
- Running tests again after cluster restart - profile QTP against system params?

Other:

- Setting up interview 2 with candidate
- xrootd bug for gcc version 6 or above fixed (DM-7679)