

Database Meeting 2016-04-13

Date

13 Apr 2016

Attendees

Fritz Mueller, John Gates, Brian Van Klaveren, Andy Hanushevsky, Vaikunth Thukral, Andy Salnikov, Serge Monkewitz, Fabrice Jammes, Unknown User (kelsey), Unknown User (npease)

Discussion items

Documentation

- John has updated shared-scan documentation in LDM-135 and pull-request to draft is up
- Brian starting work on DAX service updates to LDM-463
- AndyH starting work on XRootD-related updates to LDM-135
- Fabrice has a student (Qualid Achbal) who has fixed some errors in the qserv quick-start guides

L1 database design refresh

- AndyS continuing this effort, and is preparing a list of questions for other teams.
- Will need to make arrangements for testing the new design at scale. To be factored into F16 planning.

Shared-scan testing

- **Test results so far indicate memory-managed shared-scans working per design expectation.**
- Interesting unexpected result: solo full-scan query takes **longer** than five full-scan queries simultaneously (~10% longer on a six-hour query). John has added some additional log messages and will investigate further. Will also try experimenting with Linux kernel IO scheduler selection to see if that impacts this.

Secondary index

- Mike has been working on secondary index construction code. Considering some refactoring so more work could happen concurrent with data load.

Butler

- Nate and KT continue to work through Butler architecture issues, downstream of the detailed use-case meeting with Gregory
- Nate writing up RFC regarding repo-of-repos design
- Serge is finishing up spatial-indexing utility for use with Butler. Paul Price has provided review input which Serge is incorporating. Paul will take over the rest of this task, re. integration with pipeline tasks, since it directly involves pipeline code which Paul is currently scheduled to refactor.

XSWAP

- Vaikunth now has a proof-of-concept **Kibana dashboard up on the cluster** (see <http://ccqservmon.in2p3.fr:80/goto/2f1b3b15c5985d875a72de044bfc051> from within the SLAC network). Work continues refining this.

Misc

- Fabrice working on worker configuration enhancements. Refactoring and simplifying, and has found lots of unused worker config code that may be removed.
- Vaikunth and John have been investigating a czar -> meta db timeout issue.
- Fabrice's student Qualid will be investigating cluster-on-demand with large scale datasets, based on either Grid'5000 or possibly OpenStack. Potentially quite useful to us to for gathering our KPMs on a more continuous basis.
- New job posting is generating some interesting resumes! Fritz is setting up phone screens, and will post resumes to google docs folder for other team members who are interested.
- Serge now working on HTM additions to sphgeom library per request from Simon. Will probably take him through next week, and will be the last work of his tenure with LSST. Thanks, Serge!
- Serge to visit SLAC May 9-10 for knowledge transfer.
- Fabio will install mariadb monitoring extensions on the IN2P3 test cluster, which should be quite useful for devs when integrated with Vaikunth's dashboard work.