

Database Meeting 2016-04-06

Date

06 Apr 2016

Attendees

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

Discussion items

Data sets

- Pan-STARRS availability pushed back to September earliest.
- SUI/T is recommending getting started with the Summer2013 SDSS Stripe 82 data set.
- WISE data set from IPAC would be interesting and useful. Serge will inquire about availability to assist with this.
- Fritz, Jacek, and Brian met briefly with Tony Johnson and Jim Chiang to discuss their potential use of qserv to host the DESC Twinkles data set. The data would be available sometime around October 2016, and is estimated at ~1TB of total storage. DESC is looking around for some appropriate hardware to use to host this; if they can find this we will help them get started with qserv for evaluation.

Documentation

- John has started work on updating the shared-scan sections of LDM-135.
- Brian to start on DAX services documentation in LDM-463 this week.

L1 database design refresh

- A meeting ([notes here](#)) this past week about the proposed contents of VOEvents generated as alerts has exposed some additional design concerns. One example: if alerts contain identifiers for objects only, then the event consumers will need L1 db access to resolve those identifiers. This raises scaling concerns (e.g. hundreds of simultaneous L1 queries following the publication of an "interesting" event.)
- AndyS is experimenting with some dummy databases at this point to try and work out some of the indexing details.

Shared scans

- Proxy/czar threading issue has been resolved, and high volume query script is now working again on the cluster. This unblocks shared scan testing.
- Some preliminary shared scan tests on the cluster look promising; John is on vacation the rest of this week but testing will resume when he returns next week.

VO

- A prototype ADQL parser in Python (ported from Presto SQL parser) has been completed and looks promising for use in dbserve TAP/ADQL support.

Butler

- Nate is working on persistence improvements for butler configuration. Proposed implementation breaks something in pipe tasks currently; investigations continue.
- Design discussions were held this past week with Nate, Gregory and KT about repo search and put behaviors. This raised some issues about how some of Gregory's use cases should best be accommodated as Butler configurations; design work with Nate and KT continues in this regard.
- Nate is starting to look at object storage architecture task.
- Serge continues work on Python utils for butler spatial queries. Found and fixed a bug in sphgeom convex hull algo that was blocking.

Secondary Index

- Mike is looking at the existing secondary index generation code, which is in Python, to get a grip on what might need to be changed to accommodate the new secondary index design.
- Current implementation generates secondary index in a second pass, after the loader has distributed data to the workers. Some questions here:
 - What will need to happen with respect to chunk replication for this scheme?
 - Could/should we build the secondary index "on the fly" instead, while the loader is partitioning the data?
- Another outstanding question is whether horizontally partitioning the secondary index on a single machine will provide worthwhile performance gains (laptop tests were inconclusive in this regard).

Docker

- Yvan will be setting up a registry proxy cache at IN2P3 which should help with docker pull performance.

XSWAP

- Vaikunth expects to have an elastic search database and dashboard up soon, showing perf stats (cpu, memory, io-bandwidth) on a running cluster.

- Will work with John next week to separate query IDs and timestamps to their own log to optimize log scan times.

Hiring

- HR working to finalize Igor Gaponenko's start date – May 15 likely.
- New job posting is up in Taleo now (requisition [#2175](#)). 5 replies so far (as yet unchecked).

Misc

- Fabrice is working to improve qserv worker configuration, specifically with respect to shared scans. A question has arisen about whether we should extend the existing INI file based mechanisms, or introduce a new YAML based mechanism. Team decided after some discussion to stick with the existing INI mechanism for now.
- Vaikunth talked with Jacek about vertical partitioning tests to date, and it has been decided to conclude these investigations and write them up.
- Meta-issue from Jacek: watch for discrepancies between baseline docs and other docs or expressed/assumed capabilities! Jacek has seen several of these come to light recently. It is all of our collective responsibility to actively report/track these (file a bug in JIRA with component "Design Documents" or "Requirements Documents".)