

Database Meeting 2016-03-16

Date

16 Mar 2016

Attendees

John Gates, Fritz Mueller, Brian Van Klaveren, Vaikunth Thukral, Serge Monkewitz, Fabrice Jammes, Unknown User (kelsey), Unknown User (npease), Jacek Becla, Kian-Tat Lim

Discussion items

X16 is here!

- Planning: W16_03 sprint is closed. Fritz working with JIRA and with Jacek's planning spreadsheet to load the X16 sprints. Expect X16_03 sprint to be loaded and start 17 Mar 2016.
- X16 will have a lot of focus on documentation refresh. Most team members will have documentation tasks assigned early in the cycle. Please update LDM docs on the provided "/draft" integration branches.
- If you are documenting something and are in doubt where it goes, just go ahead and capture the doc; we can find the right place for it.

W16 leftovers

- Fritz, data distribution prototype: re-planned to finish in X16.
- Serge, spatial query utils for Butler: split off sphgeom Python wrapper part; expected to close before end of week. Utility script layer above the wrappers re-planned for X16.
- Nate, Butler multiple-repo infrastructure: reviewed, but not passing CI yet; expected to close before end of week.

Large query results

- Test remove response queuing on czar: with the original queuing, it is possible to blow up the czar on a large result from many workers. Without the queuing, the czar code apparently survives, but something untoward happens in the proxy layer after the results table is populated. John is continuing to investigate.
- runqueries.py, which we want to use to test large results and shared scans on the IN2P3 cluster, is having some difficulty. It seems the python script loses mysql connections to the proxy part way through. John is continuing to investigate.
- John noticed during large-results testing that when a worker fails to respond, the error condition does not make it all the way back to the user. Partial results are returned to the user, but no indication that a failure occurred and that the results are incomplete is provided.
- John has also noticed that result and message tables on the czar don't seem to be cleaned up in all circumstances. Possibly related to killing the czar docker container?

VO

- Brian has done some prototype experiments with thrift, protobuf, and capnproto to assess whether any of these might be useful as an alternative to writing a VOTable Binary2 parser from scratch. Results were not terribly promising. Discussions with Walter Landry at IPAC also lean towards Binary2 parser from scratch.
- Brian has developed a simplified VOTable implementation with JSON guts instead of XML. Will probably be useful as an internal format. To discuss with Trey at Monday's Data Access meeting.
- Tatiana would like to get UCDs along with data when issuing a query. PQL with canned queries might be a way to accomplish this.
- We need to assess viability of leveraging astropy VOTable code. Will introducing an astropy dependency sooner than later be a problem?

Docker

- **MULTI-NODE INTEGRATION TESTS NOW RUNNING IN TRAVIS CI, INTEGRATED WITH GITHUB!!** Thanks, Fabrice – this is super cool and useful! See build badge in README on github. See also Travis jumping in on all pull-requests and running the multi-node integration test.
- Docker bug can cause cached container graph from previous releases to become corrupt. One-time fix is to delete the container graph db and pull images fresh. Fabrice will see that this is done on the IN2P3 cluster.
- We could use some attention to clean container exit. Right now qmeta contains stale info from previous runs, etc. Might also relate to result and message table leftovers on the czar.

X-SWAP

- Vaikunth also having some trouble with the runqueries.py script
- Some of the requested monitoring infrastructure at IN2P3 does not seem to be working yet – Vaikunth to follow up with Yvan.
- More info needed by Vaikunth about the new query id feature (John to provide)

Misc

- Fabrice is currently migrating data to the second set of 25 nodes at IN2P3, so we can begin to use them as an additional test cluster.
- Vaikunth has a provisional fix deployed for the connection timeout bug, and test is running (fingers crossed!)
- Serge to extend sphgeom with HTM indexing code, per request from Simon
- Exposure metadata ingest: may be possible to use a fixed schema, e.g. CAOM