LSST Database Hangout 2014-09-17

Date

17 Sep 2014

Attendees

- Jacek Becla
- Serge Monkewitz
- Fabrice Jammes
- Andy Salnikov
- Unknown User (danielw)

Goals

- Get together, resolve bottlenecks, make progress faster, plan, design, exchange ideas

Discussion items

Qserv on distribution server

- did successful build through buildbot
    - initially had problems with software unrelated to qserv
- if we release qserv, all other packages will get tagged (in eups),want to avoid polluting tag namespace for all packages. Ideally, want qserv separated from the rest. Hopefully this will happen in Sept. But we want the Aug release, so we tried publishing
- "-b" takes eups tag, not buildId
- published, but broke "current" (overwrote existing). Thanks to Mario for bringing it back. Suggesting improvements to the publish script, see DM-1189

Logging

- finally approved
- AndyS made improvements to logging, see DM-1180
    - Jacek tried, it passes all tests, incl qserv built with new logging (in DM-207)
- all qserv code now migrated, (worker in DM-1050, all the rest in DM-207), not pushed to master yet

DM-198 (RefMatch, Join)

- finished review and resolved during hangout

DM-199 (master-worker protocol)

- code ready to be reviewed

DM-648

- Fabrice won't have time this month to finish
- AndyS taking over
- some tests are not unit tests (require special state in database). Integrate then into the existing integration tests

tagged the tip of master of xrootd git package

- and uncovered that 24-byte-long git tag breaks building xrootd!

DM-750

- more complex than expected
- Serge pre-reviewed

DM-666

- Serge reviewed, looks fine

CSS

- transient cache in in-memory mysql table in front of zookeeper? Too heavy. Although, if we need to keep schema for tables, it'd come handy
- maybe use boost property to parse json in c++?
- or do everything in python, and fill the cache that is read from c++
- metadata can be large, schema for thousands of tables (see how many dbs we have on lsst10 now)
- if we decide to clone the metadata in cache (to avoid collisions between process that updates and czar reading), try to not clone everything, only what changed.
- use protobuf instead of json? No, not self describing?
- Do we have to send uuid for each table to the worker? Hopefully not

Action items

- [ ]