Database Meeting Notes 2014-05-21

Date

May 21, 2014

Attendees

- Jacek Becla
- Serge Monkewitz
- Unknown User (danielw)
- Unknown User (bchick)
- Andy Hanushevsky
- Fabrice Jammes
- and Robert Lupton, Unknown User (robyn) for the git discussion

Goals

- Discuss blockers, short term plan and design
- Design on git merge procedure and release cycle

Discussion Items

DM-198 (RefMatch)

- need two solid days, requires some changes to TablePlugin
- TablePlugin does a lot, might consider spliting it or organize code better
- problem with paranthesis, not added where they should, need to open ticket, Serge will create the ticket

Logging

- Boost and printf style, have version of code that does both.
- Outstanding problem: performance when threashold is not met (turning off the logging code), time varies: 30nanosec - 2 milisec on slow lsst-dev03
- menu of various implementations, describe delays, features. Then discuss. Will have that tomorrow.
- requirements are in tension with one another
- in qserv we prefer printf style
- proposed name for the module: "log"

Issue with **comments showing up as Robyn's** if someone replies to automated email from Jira

- Iain is looking into it

DM-70

- Moving to June, won't happen in May (work with start in May, soon)

DM-630 review

- can wait until andyS is back

Git

- preferred process:
    - do all dirty work, squashing and rebasing in personal branch called /u/<userName>/DM-<issueNumber>
    - prepare the final, clean personal branch, send for review
    - reviewer needs to check and OK history cleanness
    - after review done and comments applied and personal branch is clean again, push to official branch /tickets/DM-<issueNumber>
    - merge with master with no-fast-forward
- rules
    - force push for personal branches - no problem at all
    - force push for official branches OK (once enabled) as long as the branch was not pushed to master. If it was pushed to master, that branch is read only forever
    - no force push to master, even if history is messed up a little
    - new people should not be allowed to push to master until they are well trained and prove they understand git well enough
- let's see if the rest of DM buys into such model, we will consider adjusting if they decide to do something else.

Fabrice will send info to qserv-l about **git bisect** scripts he has access to

Fabrice has access to **20 virtual core** (through OpenStack) at IN2P3 for testing

- useful for testing different platforms
- useful to test system dependencies that are needed
- that will be used for specialized tests, we will still rely on buildbot for regular automated builds
- need to find if/how developers can automatically trigger build of a specific branches in buildbot

memory leak in new xrootd client

- Daniel should send valgrind log

Loader

- work started during dev week
- #1 priority for Serge after RefMatch done

Optimizations to **empty chunk list**

- this is related to data loading
- partitioner can generate that, easy when it has access to all data that is loaded
- **Generating objectId index** is related to loader too

empty chunks and objId index work

- look into that after loader is ready (or during work on loader)
- ~1 month of worth
- Daniel will open ticket

Qserv Release cycle

- options on the table: 1 per DR (every 6 months) vs monthly
- will do monthly
- freeze date: 5 working days before end of each month
- these 5 days reserved to make the release solid
- only commits with blockers/critical fixes allowed during that period
- Fabrice will help out with testing, Jacek will coordinate
- monthly releases will correspond to monthly sprints in Jira Agile

Action Items

- ☑ Serge Monkewitz create ticket about issues with parenthesis
- ☐ Unknown User (bchick) document menu of different logging implementations
- ☐ Fabrice Jammes send info about git bisect to qserv-l
- ☐ Jacek Beclafind if/how developers can automatically trigger build of a specific branches in buildbot
- ☑ Unknown User (danielw) send valgrind log
- ☑ Unknown User (danielw) open ticket about empty chunk and objId index work