

Database and Data Access Meeting 2016-05-18

Date

18 May 2016

Attendees

- @AndyS @gapon @vaikunth @natepease @jgates @FabriceJammes @MichaelKelsey @AndyHanushevsky

Status Reports

Fabrice: Monitoring work (



DM-5885 - Jira project doesn't exist or you don't have permission to view it.

). Will sit with Fabien to provide

monitoring reports from large scale tests. Fixed docker git script (



DM-6130 - Jira project doesn't exist or you don't have permission to view it.

). Command line complex, will provide docs. Will work on

Docker swarm (



DM-5967 - Jira project doesn't exist or you don't have permission to view it.

). With Walid (student at IN2P3) working

to set up scripts – can he get NCSA account?

Nate: Butler stuff. RFC-184 closed as adopted; working on implementation. Storage factorization.

Vaikunth: Progress cleaning up Kibana dashboard, now gets only nodes we care about. Can adopt and do analysis on data; some binning effects due to pre-aggregation?

Andy S: Finishing tech note about L1 database schema, revision, etc. Need reviewer before publishing (soon!). Next step to start prototyping implementation. Biggest unknown: new L1 for each release, or accumulate forever? Will look at MySQL vs. PostgreSQL. Must be partitioned, based on time -- newest is active; close when full and open new one. Avoid searches across multiple partitioned.

Igor: First week; getting up to speed. Still need HipChat account, NCSA, IN2P3 (request access to QServ cluster). Want to build/test software before leaving. Get VMware installed on Mac.

Michael: Adding objectID to internal containers in lsst/partition, in order to dump secondary index.

John: Threads throttled much more than expected on cluster (



DM-5910 - Jira project doesn't exist or you don't have permission to view it.

). Possible solution: use single pool of threads to allocate

among requests; enforces maximum total usage. Bigger problem: proxy gets too much data back, then dies -- eats system resources rapidly then stops, taking czar with it. Very narrow window of "safe" queries which both exercise new code and doesn't kill proxy with results.

Andy H: Will discuss face-to-face w/John some changes to simplify his proxy stuff. Logging in XRootd essentially done except for interface to QServ

logger (



DM-4473 - Jira project doesn't exist or you don't have permission to view
it.

). Sequencing releases to coordinate with QServ

in process. Working on memory management stuff as well.