

Database Meeting 2016-01-27

Date

28 Jan 2016

Attendees

- [Fabrice Jammes](#), [Unknown User \(kelsey\)](#), [Andy Salnikov](#), [Unknown User \(npease\)](#), [Fritz Mueller](#), [John Gates](#), [Vaikunth Thukral](#), [Brian Van Klaveren](#), [Andy Hanushevsky](#), [Serge Monkewitz](#), [Jacek Becla](#)
- [Frossie Economou](#) (for part of the meeting)

Discussion items

DM-3985

- hostname, dns issues. Solved by using NAT, and docker spy

ELK

- latest version (which has useful modules) incompatible with syslogng (which is used on IN2P3 cluster)
- ask IN2P3 to consider switching to new ELK and not using syslogng

Releases

- need a designated person to bless qserv / dax releases - Fritz
- planning to migrate [Qserv Release Procedure](#) to a tech note
 - issues with readthedocs because of eups (readthedocs needs to build docs, needs setup.py)
- eups tags for qserv release - have two options:
 - a. proper tags displayed, at the expense of extra recompilation if tag changes but code did not
 - b. proper tags not displayed, but we skip recompilation if code did not change
- we opt for a.
- would be better to tag everything, yes, including 3rd party modules
 - K-T/Mario would rather not tag 3rd party with monthly or weekly tags
- would be good to not hide the original tag, e.g. if we tag mariadb as 10.6.1.lsst1, we want to see that version, and m.2016.01 would hide it
- and eups manifest can hold only one tag per module
 - Frossie will think how to satisfy both needs 😊

storage/IO updates

- hardware for czars
 - on average we will have merge from 3 machines, 10 Gbits/machine
 - disk speed (across all chars) should at least match network speed (so 40Gbits/sec)
 - might consider SSD for the results, streaming will happen in small pieces, lots of random IO
- recovery from failures
 - reserve % of IO for contingency behavior
 - industry standard ~50%. Can't afford
 - resource mgmt / throttling query throughput might help
 - consider 3 replicas instead of 2. Disk are getting denser
- intermediate results
 - small, even at 6 GB/ high vol query x 100 = <1 TB across 100+ machines
- Do a tour of LDM-141 storage/IO spreadsheet for the DB group

anaconda dependencies

- update *requests*, capture dependencies in .table file before removing anaconda
- flask now packaged together with 3 other dependencies. It is a bit hacky, there is a silent dependency...

prep for JTM

- all: should think what to discuss with the other 5 teams, will reuse the google sheet we used before.