

Database Meeting 2016-02-10

Date

10 Feb 2016

Attendees

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

Discussion items

docker

- docker config issues on centoOS7 at in2p3 cluster

serf/consul

- per docker experts, we'd need to bootstrap container to point to consul server, not worth it, it is an extra layer
- instead, put config in read only small docker volume, that is easiest way to manage configuration for docker
- related story: DM-2675

mariadb

- running into 2 issues with mariadb
 - "load data local infile" - default mariadb configuration disallows it. Fix pushed during meeting, see DM-5122
 - related to looking for mysql binaries, it looks for mysqld_safe in the path, we now have it in two places (mariadb and mariadbclient). Fix pushed to a ticket during meeting, under review.

Entry points for different db servers

- need one entry point. Create new epic, see DM-5127

memory locking

- need to change configuration on in2p3 cluster to increase hard limit

secondary index

- seeing good performance up to 10 billion entries
- seeing comparable load performance for several smaller tables and one larger, most likely because didn't parallelized, and loaded data pre-sorted

RFC-132/RFC-133/RFC-134

- currently can't do director-child join, have to allow it, or allow director-director. Also, child tables currently don't have overlaps
 - create a story for these analysis, see DM-5128
- do estimate how much effort it is to do multi-DR shared scans on the same cluster (multi-DR in more general terms == multi-tenant)
- possibly different DRs in different VMs?
- or scan coordinator, and time slicing (allocating cluster for each scan for a fraction of time)
- keep 2 latest at ncsa, older public release in the cloud (maybe google, maybe amazon, maybe private)
- yes we need tool for repartitioning to allow forward-porting science from older to newer releases (for small, L3 data sets)
- could potentially go with much smaller chunks which would allow keeping the same partitioning and cross-DR joins.