

Database Meeting 2016-01-13

Date

13 Jan 2016

Attendees

- [John Gates](#), [Fritz Mueller](#), [Vaikunth Thukral](#), [Andy Salnikov](#), [Brian Van Klaveren](#), [Unknown User \(npease\)](#), [Fabrice Jammes](#), [Serge Monkewitz](#), [Andy Hanushevsky](#), [Kian-Tat Lim](#), [Jacek Becla](#)

Discussion items

JTM

- Entire DB team plans to be there Mon-Wed. Even Fabrice coming
- Jeff T. from Oxford coming too

docker

- running multi node integration test inside docker - query ran interactively works, the very same query ran non-interactive fails with http connection to webserv problem. Still puzzled
- finally can get core file inside dockers
- writing tech note about debugging inside docker

in2p3 cluster

- in2p3 will be upgrading the 1/2 of cluster to deal with kernel panic

mariadb

- problems with mariadb / mysql solved by switching to newest version of mysql proxy
- we are still depending on mysql client because mysql python depends on mysql, switching would impact the rest of the stack
- maybe ask SQUARE to test the stack with mariadb client?
- alternatively, when we switch to sqlalchemy that will get solved
- note that we will need mariadb for pluggable authentication everywhere anyway

shared scans

- still looking how to avoid scheduling shared scan for the same chunk on multiple replicas of that chunk
- ideas tossed around involve "relative ordering" to enforce sending all queries to "primary" copy
- Would like to send interactive queries to least loaded copy
- still thinking. Not urgent, this is just resource optimization

db/table mgmt

- document as comment in DM-2802

Planning S16 cycle

- output captured in JIRA and a spreadsheet. The spreadsheet will be published soon