

Database Meeting 2015-08-12

Date

12 Aug 2015

Attendees

- [Andy Salnikov](#), [Andy Hanushevsky](#), [Unknown User \(npease\)](#), [Fritz Mueller](#), [Fabrice Jammes](#), [Serge Monkewitz](#), [Vaikunth Thukral](#), [Jacek Becla](#)

Discussion items

- check if we can talk to port 4040 on lsst-dbdev* from lsst-dbdev (Jacek)
- DM-2727 (packaging request) - resolve how to deal with packaging such modules at Bremerton
- DM-2020 and DM-2022 - will try to write up and close (Fritz)
- DM-2871 butler - really want to close it this sprint. Check with [Kian-Tat Lim](#)

DM-3161

- blocked by newly discovered issue: problems with symbols related to mysql, cssLib needs it, it is already in czarLib, czar imports both czarLib and cssLib
- right solution: split czar library into several smaller. Non trivial - need to understand dependencies between modules. Create new story in w16 to

fix that correctly. Done, see



[DM-3447](#) - Jira project doesn't exist or you don't have permission to view it.

- DM-3253 move to W16 (Jacek)
- DM-3245 move to W16 (Jacek)

Forced sources at in2p3

- don't use spatial constraints, joins are fine
- long term issue: if we have more director tables per query, want a way to specify which one is driving

Add story related to LV queries stalling for too long (would it need priorities?). Done, added a note to



[DM-2077](#) - Jira project doesn't exist or you don't have permission to view it.

, we will handle it as part of implementing shared scans.

Timeout for long running query:

- short term: make the time out very long (like we do now)
- Issue with that: we can't detect if client is mis-behaving
- predict how long a query might take and set timeout? Can't always estimate well
- periodically poke the worker: ask asynchronously: what is the status of this query? Worker should respond: queued, scheduled, working / started x sec ago

- see epic



[DM-3448](#) - Jira project doesn't exist or you don't have permission to view it.

Ask in2p3 team if ganglia monitoring at in2p3 could be made accessible from outside of slac network

Handle better connections problems. Now "uncaught exception" if # connections low and we start many queries and can't connect to mysql. Add story

about it (done, see



[DM-3449](#) - Jira project doesn't exist or you don't have permission to view

it.

). Also add comment about about #

connections in etc.cnf

Scaling tests

- ran 400K+ queries during 24h, all worked (50 simultaneous low vol and 5 high vol)
- now trying 2x more
- try queries with larger results (now on average 16KB/query)
- See email from Fabrice "one more test query" (query with large result) and see if that is still failing