

# Database Meeting 2015-08-05

## Date

05 Aug 2015

## Attendees

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

## Discussion items



[DM-2930](#) - Jira project doesn't exist or you don't have permission to view

it.

- use case: worker dies, leaves behind some in-memory tables in mysql server. These tables represent subchunks. There is a database with such materialized subchunks for each chunk that qserv materialized at a given time
- restarting mysql cleans up the contents of these tables, but the tables are not deleted. That is causing problems, because current logic is: if subchunk table exists, don't populate it
- conclusion: clean it up when worker starts
- not a scalability problem because we should not have a large number of chunks with materialized sub-chunks (in general it is slow, eg deleting 20K databases with 10 in memory tables in each takes 8 minutes. Note that each in memory table has .frm file on disk)

### Large scale tests

- data reloaded, started running tests again



[DM-3161](#) - Jira project doesn't exist or you don't have permission to view

it.

- once ready (soon) deploy the mysql-based kv interface, but unless integration with qserv is very smooth do not make it default right in the middle of large scale tests and the Bremerton activities.