

Database Meeting 2015-10-21

Date

21 Oct 2015

Attendees

- [Andy Salnikov](#), [Andy Hanushevsky](#), [Serge Monkewitz](#), [Brian Selvy](#), [Vaikunth Thukral John Gates](#), [Unknown User \(npease\)](#), [Fritz Mueller](#), [Jacek Becla](#), [Mike Kelsey](#)
- [Kian-Tat Lim](#) for parts of the meeting

Discussion items

1 min round table status

- AndyS waiting for review (DM-3574), Serge will do soon
- Brian writing requirements about authorization/authentication. Create story to capture the work
- Vaikunth playing with elk + our log files. Debugging problem with services failing after they run for long time
- John working on thread poll / scheduling
- Nate: getting metadata out of fits file. Need to talk to AndyH about css story
- Fritz: dht prototyping, demo likely end of this month. Helping AndyH
- AndyH getting up to speed with running qserv
- Mike Kelsey (working with us @50% for 3 months, possibly longer): started looking at objectId/chunk index optimizations.
- Jacek: provenance prototype

ObjectId / Chunk index discussion

- two use cases: user queries, and data loader
- for data loader, need to load 40B in 24h, so 1/2 million/sec
- build index locally on each worker (one index for all chunks), preferably serve it from RAM, or at least SSD
- serve it through xrootd, or memcached
- Redis might be an option, it has some extra smarts eg for compressing ranges
- for memcached
 - put $o(1000)$ elements per bucket to fill the pages we transfer through network
 - we could further rely on replication to speedup recovery
- for loader use case, preload cache for area that loader works on
- conclusions: try xrootd and memcached at in2p3 cluster and see what throughput we could get

- related epic:



[DM-2119](#) - Jira project doesn't exist or you don't have permission to view

it.