

# Database Meeting 2016-03-09

## Date

09 Mar 2016

## Attendees

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

## Discussion items

### Alert production DB design

- Andy S. has started research task with LSE-163, and is starting to have a few questions/concerns, mostly about temporal data. Serge points out that there is some more design info in LDM-135 section 3.1. Andy will continue to read/digest and formulate specific questions.

### PanStarrs

- IN2P3 reiterates desire to prepare infrastructure for reception of PanStarrs data well ahead of data arrival. What format will the data be, and what is the anticipated loader architecture?
- Serge mentions a rumor overheard that PanStarrs data was delivered to MAST as 5000 lbs. of hard drives...
- In any case, PanStarrs data delivery to us still not anticipated before summer.
- Fritz to find out who to ask (and ask) about delivery format. Data format will drive loader architecture, and then information on this can be forwarded to IN2P3.

### Secondary index

- Mike has started digging in to qserv code to find interfacing points for the new secondary index design. By end of meeting had a convincing list of integration points.

### Unique query ID

- Feature almost done; the qmeta part is now on master, and the qserv part is in review.

### BLOB issues

- Additional VARBINARY unit test added and is working.
- BLOB with proper escaping now passes through the partitioner successfully.
- There is still some BLOB quoting issue with qserv itself (accessing a blob through qserv and mysql can produce different results). Fabrice will open a new issue to investigate.

### Service timeout issue

- John and Vaikunth found some bad pointer-nulling code that is a smoking gun. Vaikunth has a fix and is preparing to test.

### Large results

- John has been experimenting with a flow-control change on the cluster (change to threading on czar so that it pulls more data from workers only after previously pulled data have been deserialized and committed to the results table). This does seem to work around the particular czar-explosion-on-really-large result that has been seen previously, but has exposed another problem further upstream, seemingly in the proxy layer. John and Andy S. are investigating.
- This flow control change will require some additional threading rework on the workers as well to be completely useful.

### Xswap

- All the desired monitoring tools and infrastructure are now in place on the IN2P3 cluster.
- Vaikunth says elastic search functionality is nearly ready for use by qserv devs as well.
- Will need a little more work to adapt to the new query ID feature once it lands on master.

### Butler

- Serge is nearly done with spherical geometry Python wrappers and scripts
- Nate is currently blocked on some Butler issues awaiting review/feedback from KT
- Nate planning to use RFD/RFC mechanisms moving forward to ease visibility issues with other teams
- Previously planned concrete work queue is at low water mark, and another Butler "customer survey" will be needed soon for planning

### VO

- Brian is now prototyping a TAP-enabled dbserve to get a handle on some of the issues
- Brian working with a modified TAP which has JSON guts instead of XML. This eases exploration/prototyping, and might eventually be pushed back upstream as a proposal to IVOA.

- Brain mentions that XML-heavy VO standards might be a lot more easily/readily implemented in Java rather than Python. That's a pretty big change from the current plan, however. For future consideration after we learn more from the current prototyping.

#### IN2P3 cluster

- Fabrice will add a task to get second set of nodes at IN2P3 upgraded and running
- Account has been created for Mike and he is able to access

#### Hiring

- Nate and Brian have provided input to Fritz on the job listing for the new req. Fritz to wordsmith and forward to Jacek by the end of this week.