

# Database meeting 2022-07-27

## Date

27 Jul 2022

## Attendees




- Igor Gaponenko Andy Hanushevsky Andy Salnikov Fritz Mueller Fabrice Jammes John Gates Joanne Bogart

## Notes from the previous meeting

- Database meeting 2022-07-20

## Discussion items

Discussed	Item	Who	Notes
✓	Project news	Fritz Mueller	
✓	DP02	team	<p>Igor Gaponenko: a subset (one tract) of the table <code>ForcedSourceOnDiaObject</code> has been loaded into <code>qserv-int</code>. There is one column <code>skymap</code> that's present in the <code>Parquet</code> files, but not in the TAP schema.</p> <p>Colin Slater: what's the final status of the <code>Parquet</code> files for <code>ForcedSourceOnDiaObject</code>? Are we ready to ingest the rest based on the preliminary tests of one tract ingested into <code>qserv-int</code>? What should we do about the extra column?</p> <p>Action items for Igor Gaponenko:</p> <ul style="list-style-type: none"><li>proceed with translating and ingesting the full version of the table</li><li>eliminate the column <code>skymap</code> during the table partitioning phase.</li></ul> <p>A general discussion on how to configure (annotate) Ingest to eliminate unwanted columns and manage type mapping</p> <p>Fritz Mueller on <code>RefMatch</code> columns:</p> <ul style="list-style-type: none"><li>not much progress in debugging</li></ul> <p>Action items:</p> <ul style="list-style-type: none"><li>Igor Gaponenko: ingest truth tables into the <code>small</code> Qserv cluster at NCSA</li><li>Fritz Mueller: instrument <code>RelationGraph</code> with more debugging and work with Igor Gaponenko to deploy the test version of Qserv at NCSA and used it for testing.</li></ul>
✓	Qserv in IDF	Fritz Mueller	<p>Context (reported by Fritz Mueller):</p> <ul style="list-style-type: none"><li>stress testing Qserv (Fritz Mueller was asked to run this test by Frossie)</li><li>ended up with 20 parallel queries run in Qserv or no more</li><li>a few mysteries in the TAP query queuing and pooling have been discovered with these tests</li><li>the load testing is creating some turbulence affecting the users</li><li>good news: Qserv has been good in digging out itself from the "avalanche" (in its current version)</li><li>bad news: scaling in digging out doesn't look good. Qserv slows down faster than expected (it <i>should</i>?) Perhaps the slowness is caused by some config parameter (number of threads, etc.). John Gates has been tasked to look at that.</li><li>another observation: some queries were left in a strange state (still <code>EXECUTING</code>) after they got explicitly canceled by Fritz Mueller. Fritz Mueller's theory is that cancelation is a heavy procedure. Asked John Gates to investigate.</li><li>The dashboard seems to have performance issues when displaying long lists of past queries.</li><li>Sometimes only the headers are shown on the Dashboard and no content<ul style="list-style-type: none"><li>Igor Gaponenko reported his observations on the cross-site complaints made by Chrome browser</li><li>This needs to be investigated</li></ul></li><li>more pressure to set up better monitoring of Qserv</li></ul> <p>Action items for Fritz Mueller: discuss the next steps with Frossie</p>
✓	Status ob the <code>ObsCore</code> table extractor	Andy Salnikov	<p>Has prepared the draft version of the technote on extracting the data:</p> <ul style="list-style-type: none"><li><a href="https://dmtn-236.lsst.io/v/DM-35532/index.html">https://dmtn-236.lsst.io/v/DM-35532/index.html</a></li></ul>

	Temporarily solutions to replacing a loss of the development/testing platform at NCSA (August 15th)	<a href="#">Fritz Mueller</a>	<p>The development cluster for code development and interactive <code>itest</code>:</p> <ul style="list-style-type: none"> <li>The short-term solution: <a href="#">Fritz Mueller</a> will commission a big VM in the Google Cloud</li> </ul> <p>The mid-to-large-scale tests:</p> <ul style="list-style-type: none"> <li>setup a separate Qserv cluster at IDF (or use <code>qserv-dev</code>)</li> </ul>
 (postponed till the next week because <a href="#">Fritz Mueller</a> is presently oversubscribed)	Merging three Git packages <code>qserv</code> , <code>qserv-operator</code> and <code>qserv-ingest</code> into <code>qserv</code>	<a href="#">Fritz Mueller</a> <a href="#">Fabrice Jammes</a>	<p>Context:</p> <ul style="list-style-type: none"> <li>it was discussed last week</li> <li><a href="#">Fritz Mueller</a> was going to work on integrating/cloning <code>qserv-operator</code> into <code>qserv</code> this week.</li> <li>any progress?</li> </ul>
	Status of <code>qserv-operator</code> and <code>qserv-ingest</code>	<a href="#">Fabrice Jammes</a>	<p>For the new version of <code>qserv-ingest</code>, issues have been observed when ingesting DP02 into <code>qserv-dev</code> (namespace <code>dp02</code>):</p> <ul style="list-style-type: none"> <li>There are the InnoDB locking problems <ul style="list-style-type: none"> <li>Tried to solve this by explicitly locking the table</li> <li>This problem is not seen in IN2P3</li> </ul> </li> <li>MySQL client to the <code>ingest-db</code> timed out <ul style="list-style-type: none"> <li>his problem is not seen in IN2P3 either</li> <li>the client was keeping the connection for 7 hours and then disconnected</li> <li>a fix to add automatic retries has resulted in another class of problems reported by the Replication Controller during transaction commit time. The Controller was reporting the error <code>"Connection reset by peer"</code> on queries sent to the Replication database</li> <li><a href="#">Igor Gaponenko</a> has a theory that MySQL server may be misconfigured (the connection limit is too low, or timeouts are too low). This needs to be investigated.</li> </ul> </li> </ul> <p><a href="#">Igor Gaponenko</a> has evaluated the connection limit parameters of the Replication database and found that the parameter's value is too low <b><code>connect_timeout=10</code></b> as compared with what's set for other MySQL databases (czar and workers) of the cluster <b><code>connect_timeout=28800</code></b>.</p> <ul style="list-style-type: none"> <li>action item: locate at what code this setting is done (Qserv container's "entry-point"?), fix it, and let <a href="#">Fritz Mueller</a> include it into the new Qserv release to be deployed in IDF tomorrow during the Thirdday maintenance window.</li> </ul> <p><a href="#">Fritz Mueller</a> proposed to look at an existing solution for using an on-the-shelf queue manager to support the ingest workflows. The scaling issues are only going to get worse over time.</p>

## Action items

