

# Database Meeting 2022-01-26

## Date

26 Jan 2022

## Attendees



- Igor Gaponenko John Gates Andy Salnikov Unknown User (npease) Andy Hanushevsky Fabrice Jammes Colin Slater Fabrice Jammes

## Goals

- See a list of topics below

## Discussion items


Time	Item	Who	Notes
	Project news	Fritz Mueller	<p>Fritz Mueller Remote work location/permission forms are required to be filled out before the end of this week (Friday).</p> <p>Fritz Mueller A project management meeting is going on in Tuscon right now. Concerns on the multi-site processing.</p> <p>Presentations made at the last week's <a href="#">2nd Data Facilities Planning Workshop - 2022-01-19/20</a> are linked to that page. <a href="#">Fritz Mueller's</a> notes on the Workshop:</p> <ul style="list-style-type: none"><li>NCSA decommissioning LSST-allocated hardware end of October</li><li>ComCAm is the main focus. This Summer we expect enough k8s nodes at USDF and IDF</li><li>Qserv wasn't specifically mentioned.</li></ul> <p>Fritz Mueller Replacement for NCSA resources (for Qserv): mentioned the problem to Will who agreed that IDF could be used for testing Qserv at scale if needed after we lose NCSA. Qserv doesn't worry LSST management because is primarily required for the final stages of DRP. Though, the team still needs sufficient resources to test scalability.</p> <p>Igor Gaponenko we need early access to Kubernetes infrastructure at USDF as early as possible to ensure the infrastructure quality, stability and performance are adequate for Qserv.</p> <ul style="list-style-type: none"><li>Fritz Mueller SLAC experience with this technology is rather limited realistically speaking just one person - Yee). NCSA was managing Kubernetes manually, and SLAC is going to do so. It's expensive to pay for external support (Google?) for Kubernetes infrastructure.</li></ul> <p>Andy Hanushevsky discussed using polling <a href="#">Rucio</a>. Another interesting idea was to use Kafka in Rucio.</p> <p>Fritz Mueller mentioned that <a href="#">Andy Salnikov's</a> work on APDB would be affected by the OGA's "rack" at USDF, that's supposed to be commissioned before this Summer (for ComCam). No hardware has been ordered yet. This creates uncertainty in planning.</p>
	Getting ready for DP0.2		<p>Fritz Mueller commented on plans for loading a slice of DP0.2 into IDF. We're past the milestone by 1 month. Reasons: delayed migration of the Kubernetes-based deployments to the new ("lite") containers. Still see some bumps in that direction. Options:</p> <ul style="list-style-type: none"><li>wait for another week before the new version of the operator and ingest will be ready,</li><li>or use existing Qservs at IDF and existing ingest tools for that</li></ul> <p>Fritz Mueller would like to retain existing data during Qserv upgrade and avoid reloading all (including existing) data from scratch as proposed by <a href="#">Fabrice Jammes</a>.</p> <ul style="list-style-type: none"><li>MariaDB and schema upgrade problems need to be addressed here.</li><li>We need to make a progress in this (upgrade) path.</li></ul> <p>Igor Gaponenko DP0.2 comes with an extended schema (compared with DP0.1) which may need to be tested. Do we have the schema? Do we have sample data? Shall we test load a slice of DP0.2 into NCSA first?</p> <ul style="list-style-type: none"><li>Fritz Mueller Both are available. Data exist as Parquet files that need to be translated into the CSV/TSV format. Data may be at a shared filesystem at NCSA. We need to contact <a href="#">Hsin-Fang Chiang</a> to figure out where to find it.</li><li>Fritz Mueller loading into the <b>small</b> cluster at NCSA is a good idea. We would have to do it anyway.</li></ul>

<p>Problems with Python 3.10, SQLAlchemy, and MariaDB 10.5 in the Kubernetes-based ingest workflow.</p> <p>Discussed <a href="#">Fritz Mueller</a>'s idea of putting Qserv and the operator into a single Git package.</p>	<p><a href="#">Fabrice Jammes</a></p>	<p><a href="#">Fabrice Jammes</a> Qserv ingest is broken after migrating to the new version of MariaDB. A problem exists with the validation stage. The loading itself works.</p> <ul style="list-style-type: none"> <li>• there is a hack that might be employed to work around the problem, such as removing SQLAlchemy and replacing it with a simple tool.</li> </ul> <p><a href="#">Igor Gaponenko</a> can it be all written in GoLang instead of Python?</p> <ul style="list-style-type: none"> <li>• <a href="#">Fritz Mueller</a> thinks that GoLang should not be used for any projects beyond the operator.</li> </ul> <p><a href="#">Igor Gaponenko</a> can we have a single <b>base</b> container for everything in Qserv, including Kubernetes-based ingest?</p> <p><a href="#">Igor Gaponenko</a> can we turn <a href="#">Fabrice Jammes</a>'s ingest framework into a ready-to-use tool that would be configurable by users to ingest their data? It would be easier for the users to do declarative ingests rather than writing customized code on top of the framework.</p> <ul style="list-style-type: none"> <li>• <a href="#">Fritz Mueller</a> good idea worth considering!</li> </ul> <p><a href="#">Fritz Mueller</a>: we're still having troubles with coordinating releases between Qserv and the operator. There are rough edges in an area of configuring Qserv from the operator. We need a solution here:</p> <ul style="list-style-type: none"> <li>• submodules as proposed by <a href="#">Fabrice Jammes</a> is not the most optimal option here (though, it's possible to use them)</li> <li>• we might set up GHA to do cross-repo CI. Though, it's a complicated problem as it requires configuration efforts</li> </ul> <p><a href="#">Fritz Mueller</a> another possibility would be to put everything (Qserv and the Operator code) into one repo. The team should think about possible complications of that.</p> <p><a href="#">Igor Gaponenko</a> what about dependencies? The operator may have dependencies that might interfere with the ones of the Qserv container.</p> <ul style="list-style-type: none"> <li>• <a href="#">Fritz Mueller</a>: that has solutions.</li> </ul> <p><a href="#">Unknown User (npease)</a> is concerned regarding the built time of the operator.</p> <ul style="list-style-type: none"> <li>• <a href="#">Fabrice Jammes</a>: that is very fast.</li> </ul> <p><a href="#">Igor Gaponenko</a> what does the "great unification" would mean from the source tree and build tools (CMake) perspective?</p> <p>* <a href="#">Unknown User (npease)</a> combining the source trees is not a problem. Different tools would be still used for each.</p> <p><a href="#">Igor Gaponenko</a> integrating the ingest framework into the single Git package will create problems for potential users who might need to customize it for specific ingests. A solution would be to turn the ingest framework into a customizable ready-to-use tool instead of being the framework.</p>
<p>Progress on topics discussed at the previous meeting <a href="#">Database Meeting 2022-01-19</a></p>		<p><a href="#">Igor Gaponenko</a> published an extended version of the proposal for <a href="#">Possible deployment and management scenarios for Qserv</a>. Started working on improving the inner communication protocol of the Replication/Ingest system to make it "cloud-friendly"</p> <div data-bbox="592 1119 1382 1289" style="border: 1px solid orange; padding: 10px; margin: 10px 0;">  <a href="#">DM-33376</a> - Jira project doesn't exist or you don't have permission to view it.     </div> <p><a href="#">John Gates</a> merged</p> <div data-bbox="686 1350 1476 1520" style="border: 1px solid orange; padding: 10px; margin: 10px 0;">  <a href="#">DM-31537</a> - Jira project doesn't exist or you don't have permission to view it.     </div> <p>into the <b>main</b> branch.</p> <p><a href="#">Igor Gaponenko</a> Any news on the new version of XROOTD that was mentioned at the last meeting?</p> <ul style="list-style-type: none"> <li>• <a href="#">Andy Hanushevsky</a> improved the version of the redirector has increased parallelism by an order of magnitude (depending on the number of cores available on a machine). It will be available in release 5.4.1. (to address <a href="#">Igor Gaponenko</a>'s concerns) No backward-incompatible changes should be expected in the configuration system. Still need to check if the affinity-specific improvements (developed for Qserv and used by <a href="#">John Gates</a>) would still work in the new version.</li> </ul>

	Documenting Qserv		<p><a href="#">Fabrice Jammes</a> reported troubles with the LSST documentation portal &amp; tools.</p> <p><a href="#">Fritz Mueller</a> proposed to move the documentation tools into the Qserv build container rather than a separate documentation container as it's implemented now.</p> <p><a href="#">Fritz Mueller</a> Move version-specific documentation into Qserv Git package and keep using Confluence for external (catalog) references only. LSST Doc has many benefits - it automatically links to the Git release.</p> <p><a href="#">Igor Gaponenko</a> expressed concerns regarding the organizational structure and visual styles of the Sphinx-generated documentation adopted by LSST.</p> <ul style="list-style-type: none"><li>• <a href="#">Fabrice Jammes</a> suggested looking at other tools that are used for generating Kubernetes documentation. The latter looks nicer and more convenient (useful and convenient navigational structure) compared to what LSST is relying upon.</li></ul>
--	-------------------	--	---

## Action items

- ☒ [Igor Gaponenko](#) should fill out the remote work forms before the end of the week.
- ☐ [John Gates](#) should fill out the remote work forms before the end of the week.
- ☐ [Fritz Mueller](#) expects the team to provide feedback on an idea to unify all Qserv-related code (Qserv itself, the operator, and the ingest framework) into a single Git package.
- ☐ [Fritz Mueller](#) will make the PR on adding logging to **qhttp**

 [DM-33088](#) - Jira project doesn't exist or you don't have permission to view it.

- ☐ [Fritz Mueller](#) will find a location of the DP0.2 schema and data for [Igor Gaponenko](#)