

Database Meeting 2022-02-23

Date

23 Feb 2022

Attendees


- [Igor Gaponenko](#) [Fabrice Jammes](#) [Unknown User \(npease\)](#) [Andy Salnikov](#) [Joanne Bogart](#) [Colin Slater](#) [John Gates](#) [Fritz Mueller](#)
-




Goals

- Please, add topics for discussion as needed.

Discussion items

Time	Item	Who	Notes
	Project news	Fritz Mueller	<p>Anything interesting from DM Leadership Team Virtual Face-to-Face Meeting, 2022-02-15 to 17</p> <ul style="list-style-type: none">• mostly concerned about "missing functionality" in light of the upcoming telescope commissioning• scaling issue of the processing pipeline after moving from HSC to DP0.2• transitional plan for SLAC (USDF)• many small files in Butler are a big concern. Rucio doesn't scale well for millions of small files.• image serv support to be added for DP0.2• meeting "embargos" <p>News from USDF:</p> <ul style="list-style-type: none">• hardware is coming and being installed• HSC reprocessing will be tested at SLAC using USDF resources (Fritz Mueller is involved in this) <p>No additional requirements/expectations for Qserv have been aired so far, except ingesting user data products:</p> <ul style="list-style-type: none">• A lot of uncertainties in many areas, including API, authorization, registering tables in the TAP service. Another problem is retaining large result sets obtained from Qserv (5 GB scale, etc.). How to handle this?• How do ingest from the perspective of "naive" users?• Igor Gaponenko service-like ingest workflow for users might be a solution.• concerns regarding user data quality (depending on the data format, data may need to be further processed/sanitized)

.	Status of DP0.2	<p>Igor Gaponenko a slice (140 GB, 40 chunks) of the catalog ingested into the <code>small</code> Qserv cluster at NCSA</p> <div data-bbox="654 180 1443 348">  DM-33733 - Jira project doesn't exist or you don't have permission to view it. </div> <p>Fabrice Jammes ingest workflow readiness for ingesting into IDF?</p> <p>Colin Slater what's the scale of the catalog compared to DP0.1? What's the ETA for the catalog inputs?</p> <ul style="list-style-type: none"> • roughly x5 increase in the amount of data (wider tables) • data will be ready in late March • ingesting will start in April • overall DP0.2 readiness is June 2022 <p>Joanne Bogart can we talk again about the extra tables for that catalog?</p> <ul style="list-style-type: none"> • response received via the team's Slack channel: <div data-bbox="654 636 1443 940"> <p><i>Access to DC2 truth info such as SN parameters, stellar types, variability light curves, cluster locations, etc., has been requested by DP0 delegates, and the Rubin project has inquired if those data could be made available.</i></p> <p><i>JB: I don't know how they will be used or whether the content will ultimately end up in Qserv.</i></p> </div> <p>DESC truth match tables may be added.</p> <p>Image metadata (expo) tables may be added to support images.</p> <p>A possibility of having match tables between DP0.1 and DP0.2. Maybe, depending on a request.</p> <p>Possible issues:</p> <ul style="list-style-type: none"> • 5x more data to lock in RAM for each chunk when processing JOINS • DP0.1 has 2 TB in the total spread between 1000 chunks which requires 2 GB of RAM per chunk on average • it means 10 GB/chunk for DP0.2 • the expected dynamic range of the chunk size may be a potential issue here • using SSDs may offset the need in locking chunks in memory
	Supporting modalities of Qserv deployments	<p>This is a big topic that was discussed at the previous meetings. A number of ideas have been expressed, including database initialization into <code>init</code> containers. Fritz Mueller may have a proposal.</p> <ul style="list-style-type: none"> • the best practice in Kubernetes is to put databases and the relevant <code>init</code> containers in separate from the application pods • the app pods would have <code>init</code> container waiting for the database <code>init</code> containers to be ready • A problem is that database sockets can't be shared between pods in this architecture • also root access will be required to the databases (why?) • to address this, Fritz proposes to colocate pods on the same node. This still needs to be verified. • IG: are <code>init</code> containers sequenced across different pods of the same deployment? • IG: what about schema migration? Would it be done via the database pod? • IG: the Replication system needs root access to the CZAR database <p>Fritz Mueller will put these ideas into a proposal to be discussed further. The document will be shared with the team before the next group meeting.</p> <p>Relevant changes in the <code>entrypoints</code> are expected.</p>

	<p>Adding Replication system's "registry" (redirector) to the Replication system's deployment</p>	<p>Igor Gaponenk</p>	<p>The context:</p> <div data-bbox="613 159 1399 327">  DM-33376 - Jira project doesn't exist or you don't have permission to view it. </div> <p>Besides adding a new service, schema change in the Replication database will be required.</p> <p>Fabrice Jammes <code>qserv-operator</code> needs to be modified and extended.</p> <p>Fritz Mueller will build a new Qserv release (or tag Qserv) based on the current state of the <code>main</code> branch before</p>
	<p>Progress on other topics discussed at the previous meeting Database Meeting 2022-02-09</p>		<p>The "Grand unification" of the Git packages</p> <div data-bbox="654 533 1443 701"> <ul style="list-style-type: none">  DM-33618 - Jira project doesn't exist or you don't have permission to view it. </div> <p>Moving Qserv documentation into the container</p> <ul style="list-style-type: none"> Fritz Mueller has it almost ready to go (after the PR will be created and reviewed): <div data-bbox="654 825 1443 993">  DM-32709 - Jira project doesn't exist or you don't have permission to view it. </div> <p>Fabrice Jammes do we still need to discuss migrating configurations of the tests catalogs at https://github.com/in2p3-dp0 to the current API of the Ingest system?</p> <ul style="list-style-type: none"> ingest workflow config files would need to store the API version as well. This will be added later in Git package <code>{{qserv-ingest}}</code>

Action items

