

Database Meeting 2022-11-09

Date

09 Nov 2022

Attendees

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

Notes from the previous meeting:

- [Database Meeting 2022-11-02](#)

Discussion items

Discussed	Item	Notes
✓	Project news	<p>Fritz Mueller:</p> <ul style="list-style-type: none">gave a presentation on Qserv at ADASSthere is a 3-page document to be DMTN-243 at Pub Board for review: https://dmtn-243.lsst.io/action item for Igor Gaponenko: send Fritz Mueller a few sentences on the parallel ingest system to be included in the publication <p>The upcoming event:</p> <ul style="list-style-type: none">next Mon & Tue: 4th Data Facilities Workshop - 2022-11-14/15 <p>Fritz Mueller on the office move:</p> <ul style="list-style-type: none">there is a lot of pressure to move out of B50a decision has to be made ASAPmost likely it's going to be the first-floor office (yet to be settled within the team)the equipment will be moved by SLAC (there is a special team of people who do it)

<p>✓</p> <p>(continued from the previous week)</p>	<p>A bug in Qserv czar when handling failed chunk queries</p>	<p>Context:</p> <ul style="list-style-type: none"> the topic was discussed at the previous meeting Database Meeting 2022-11-02 creating table indexes for the sub-chunk column didn't seem to help to address the issue <ul style="list-style-type: none"> a "byproduct" of the investigation was an interesting discovery on the ordering of rows in sub-chunk and chunk tables in which rows belonging to sub-chunks are not clustered together. This would dramatically reduce the performance of the subchunk materialization algorithm, even if the sub-chunk index is present in the tables. Note that the order of the rows is defined by the partitioning tools (<code>sph-partition</code> and <code>sph-partition-matches</code>). A few ideas on how to improve this were discussed between Fritz Mueller and Igor Gaponenko, including improving the partitioner, postprocessing (sorting rows) partitioned contributions by the ingest workflows before ingesting them into Qserv or doing this inside the Ingest system. <ul style="list-style-type: none"> Action items for Igor Gaponenko: <ul style="list-style-type: none"> Experiment with reordering rows of the fully ingested chunk and overlap tables to group the rows based on their sub-chunk numbers to see if such optimization would speed up the sub-chunk materialization. This experiment assumes the table indexes are already present. make a JIRA ticket if there will be any benefits in the reordering John Gates commented on this by saying that the reordering <i>may</i> not help since sub-chunks are normally materialized from tables that are already locked in the memory of the worker host. Igor Gaponenko replied: we still have cases (deployments, configurations) where the memory locking mechanism may not work. The experiment will tell. the problem was the root cause of the OOM crash at the workers JIRA ticket was opened by John Gates <div data-bbox="459 646 1313 814" style="border: 1px solid orange; padding: 10px; margin: 10px 0;"> <p>⚠ DM-36935 - Jira project doesn't exist or you don't have permission to view it.</p> </div> Igor Gaponenko has started testing the effect of this improvement using the <code>slac6</code> Qserv <p>Igor Gaponenko:</p> <ul style="list-style-type: none"> should simplify the management scripts for <code>slac6</code> to help John Gates manage Qserv may run the low-level experiments with manually materializing sub-chunk tables and running MySQL queries against those to see if there are some scalability issues within MariaDB itself for slowing down these queries when there are a few hundred in MEMORY tables in use
<p>—</p> <p>(no news, postponed till the next meeting)</p>	<p>User-generated data products</p>	<p>TBC...</p>
<p>✓</p>	<p>qserv - operator and qserv - ingest</p>	<p>One topic from the previous meeting:</p> <ul style="list-style-type: none"> the FrDF team is looking at developing a better Parquet-to-CSV translator (to be written in C++) with a possible option of integrating the one into the partitioning tools any news here? <p>Fabrice Jammes and Igor Gaponenko discussed an option for formal schema validation:</p> <ul style="list-style-type: none"> the Ingest API's services will report the expected schema in some formally recognized schema definition format the ingest workflows would have to employ the corresponding formal schema validation tools to make sure the JSON files to be used by the workflow are compatible with the expected schema the very same should be done for the <code>qserv-ingest</code>'s own file <code>metadata.json</code> since the schema of the file may also evolve over time, and different copies of the files might be laying around (GitHub, file systems, etc.). altogether, this early checking should reinforce the ingest campaigns and prevent the "midflight" failures <p>Fabrice Jammes Where to put the documentation on the versioning of <code>metadata.json</code>?</p> <ul style="list-style-type: none"> Fritz Mueller: unlike the Ingest API, <code>qserv-ingest</code> is facing operations. So, it needs to be documented in the relevant (for operations) area of the documentation. <p>Igor Gaponenko: there is a new feature of the Ingest system that's about to be merged into the <code>main</code> branch of Qserv:</p> <div data-bbox="415 1734 1269 1902" style="border: 1px solid orange; padding: 10px; margin: 10px 0;"> <p>⚠ DM-36722 - Jira project doesn't exist or you don't have permission to view it.</p> </div>

	Status of Obs Core tables	Andy Salnikov : the most recent status of the project can be found at: Live ObsTAP service deployment
---	---------------------------	---

Action items

