

Database Meeting 2022-11-02

Date

02 Nov 2022

Attendees

- Igor Gaponenko Andy Salnikov Fritz Mueller Andy Hanushevsky

Notes from the previous meeting:

- Database Meeting 2022-10-26

Discussion items

Discussed	Item	Who	Notes
✓	Project news	Fritz Mueller	<p>Project news:</p> <ul style="list-style-type: none">no DMLT this weekgave a Qserv presentation at the ADASS conference <p>Office spaces at B48 (ROB)</p> <ul style="list-style-type: none">Fritz Mueller and John Gates visited both officeswe will continue discussing the topic next week when John Gates will be able to join the conversation
✓	Supporting user-generated data products in Qserv	team	<p>Context:</p> <ul style="list-style-type: none">started discussing the topic at the previous meeting <p>Igor Gaponenko will begin looking at how to improve <code>github</code> to support multi-part form data in the request's body using Boost Beast</p> <p>Fritz Mueller will talk to Flossie about where to put the new services</p>
✓	Code reviews	team	<p>Context:</p> <ul style="list-style-type: none">the team has shrunkAndy Salnikov is now 100% on MW + APDBongoing development efforts in the R-I system (where we still have a lot of work to get ready before the LSST data releases) and "new" Qserv (soon) requires reviewers <p>Options:</p> <ul style="list-style-type: none">external solutions: Igor Gaponenko automated (AI-based) code reviews in GHA to do the basic code analysis. Examples: https://codeball.ai/internal solutions: Fritz Mueller proposes a series of "code tours" to help the remaining reviewers to stay up-to-date with the ongoing effort to improve the R-I system and Qserv. <p>Fritz Mueller had a negative experience with the older pre-AI tools</p> <p>Igor Gaponenko may begin experimenting with the new generation of automated tools to see if there will be any benefits</p> <p>Discussed the following options to help reviewers understand the reviewed tickets:</p> <ul style="list-style-type: none">developers should write good descriptions of the proposed changes in JIRA ticketsspecifically for the R-I system, Igor Gaponenko will make a presentation on the design of the systemwe may also begin making presentations at the group meeting to cover various topicshaving design documents available to reviewers would be helpful as well

	(Possible) bug in Qserv <code>czar</code> when handling failed chunk queries.		<p>Context:</p> <ul style="list-style-type: none"> it was discussed at the previous meeting John Gates began looking at the problem using the Qserv cluster at USDF (<code>slac6</code>). Issues with <code>sudo</code> the privileges of John's account were seen. We also got an official bug report for Qserv that <i>may</i> be relevant in this context: <ul style="list-style-type: none"> https://community.lsst.org/t/truth-match-and-forcedsourceondiaobject-tables-are-available/7088 <p>Fritz Mueller on what we know so far:</p> <ul style="list-style-type: none"> mysterious OOM worker restarts need to be investigated. We don't know what causes that. it's been observed that it takes too much time for Qserv to process sub-chunks. <ul style="list-style-type: none"> Igor Gaponenko Do we have a problem with materializing those? Fritz Mueller Missing indexes for <code>subChunkId</code> on the truth table? the data model of the truth tables (RefMatch) in DP02 is complex. There is a synthetic unique key that users are not aware of how to use it's difficult for scientists to write meaningful and efficient queries BTW, we won't have these tables in the LSST DRs. Though, we will still have other RefMatch tables. <p>Igor Gaponenko suggests that improving <code>czar</code>'s retry logic for failed chunk queries could help. One of the improvements would be to introduce delays before retries in order to give the restarted workers extra time to recover. Otherwise, we would be just wasting the attempts. Perhaps, we could switch from the retry count-based logic to the time-based? Basically, instead of counting the number of failed attempts count the amount of time permitted for recovery.</p> <p>Fritz Mueller provided extra info on the shared scan optimization not working at IDF due to memory locking not working</p> <ul style="list-style-type: none"> had a conversation with John Gates John Gates thinks that if the memory locking is not working then we should turn it off since there may be other bad effects we need to identify the config option to disable the locking and try the one at IDF <p>Igor Gaponenko the query that is giving us troubles is known. This is the same query that was tested at SLAC (Qserv at USDF)</p> <p>Igor Gaponenko in order to improve the stability of Qserv, we may also experiment with increasing the replication level to 2. This requires increasing the storage capacity of Qserv workers from 10 TB up to at least 16 TB or higher:</p> <ul style="list-style-type: none"> Fritz Mueller will look at options here <p>Fritz Mueller has proposed the following plan on how to proceed with the investigation to address both the short-term goal (of making IDF users happy) and the long-term one (improving Qserv):</p> <ul style="list-style-type: none"> Step 1: Deploy the table index for <code>subChunkId</code> on the truth table (the director table) Step 2: Experiment with "Igor's crazy" queries at IDF with the memory lock switched on and off to see the effect of the switch on the performance, memory, and stability of Qserv workers. Step 3: If turning the memory locking off still won't help then bump the replication level to 2 to see if that will help. Step 4: And, finally, begin looking at Qserv code to solve the problem in the long run.
 (postponed till the next meeting)	qserv-operator and qserv-ingest	Fabrice Jammes	<p>One topic from the previous meeting:</p> <ul style="list-style-type: none"> the FrDF team is looking at developing a better <code>Parquet-to-CSV</code> translator (to be written in C++) with a possible option of integrating the one into the partitioning tools any news here?

Action items

