

Database meeting 2022-07-13

Date

13 Jul 2022

Attendees



- Igor Gaponenko Andy Salnikov Colin Slater Fritz Mueller Andy Hanushevsky Joanne Bogart

Notes from the previous meeting

- Database Meeting 2022-07-06

Discussion items

Discussed	Item	Who	Notes
✓	Project news	Fritz Mueller	PCW 2022? SLAC performance review: <ul style="list-style-type: none">Igor Gaponenko should do it before the end of the week.
✓	Status of DP02	team	<p>Fritz Mueller:</p> <ul style="list-style-type: none">DP02 is going on quite well. No user complaints so far on the missing truth tables <p>Fritz Mueller: status of the RefMatch table MatchesTruth?</p> <ul style="list-style-type: none">the special partitioning tool sph-partition-matches was improvedthe partitioned version of the table MatchweseTruth was loaded into qserv-intCSS was manually fixedthe RelationGraph looks fine, the rest of the code looks fine toohowever, the query evaluation for the table is still not workingthis requires further debugging using gdbIgor Gaponenko will ingest these tables into the small cluster at NCSA to help Fritz Mueller's effortAndy Salnikov has proposed adding the integration test for the RelationGraphIgor Gaponenko should add support for the RefMatch tables to the Replication/Ingest systemFritz Mueller will extend the integration test to support RefMatch tables <p>Colin Slater: status of the input data of the ForcedSourceOnDiskObject table?</p> <ul style="list-style-type: none">is about to run the final fixup stagethe data will be ready within a timeframe of 1 daythe number of the Parquet files will be greatly reduced (no splitting tracts to patches) which would help a lot with speeding up the Parquet-to-CSV translation <p>Igor Gaponenko: further improvements to the Replication/Ingest system based on the DP02 experience</p> <p>Colin Slater: need to improve the monitoring of the Qserv usage in IDF</p> <ul style="list-style-type: none">Fritz Mueller mentioned Fabrice Jammes's idea to use Prometheus monitoring by pulling metrics from QservColin Slater proposed to pull metrics from Qserv and ship them to Google for monitoring, aggregation, and visualizationFritz Mueller right now we are not logging anything from Qserv since it's too expensive. So, we don't even have the searchable log
✓	Development infrastructure for Qserv		<p>Igor Gaponenko: need a replacement for Qserv development platform that we're losing at NCSA in 1 month (the August 15th deadline)</p> <ul style="list-style-type: none">still, no progress on the temporary instance based on the loaner hardware at SLACthe IT is quite busy working on the USDF infrastructure <p>Fritz Mueller will work with SLAC IT and Richard to accelerate this. Also, in September we are supposed to have 4 former master[1-4] machines as the temporary development platform. As the meantime, the IDF development cluster (qserv-dev?) is the only option.</p>

<div></div> <p>NOTE: the topic will continue to be discussed next week after Fabrice Jammes will get back from vacation.</p>	qserv-ingest	Igor Gaponenko Fabrice Jammes	<p>Context:</p> <ul style="list-style-type: none">Igor Gaponenko: extending the versioning mechanism in REST API of the Replication/Ingest and the implementation. This requires making changes to the JSON config files (to store the version numbers) and the ingest workflows (including <code>qserv-ingest</code>). There is the PR on the Git package <code>qserv-ingest</code> in the JHIRA ticket mentioned below. <div><div></div><div>DM-35456 - Jira project doesn't exist or you don't have permission to view it.</div></div> <ul style="list-style-type: none">The CI is failing on the PR due to Repl service version incompatibility. Two PRs exist in the scope of the same JIRA ticket. Both migrated to version 9 simultaneously.Do we need the common Python API for the Replication/Ingest system? We have presently 2 separate implementations: one in the Qserv container's CLI (source path <code>src/admin/python</code>), and the other one in the Git package <code>qserv-ingest</code> (source path <code>rootfs/python</code>). This seems like a duplicate effort. <p>Fritz Mueller proposed to:</p> <ul style="list-style-type: none">give the high priority for moving <code>qserv-ingest</code> into <code>qserv</code>(after that) work on the common Python API to the R-I system to be shared by the integration test and the workflow <p>On the Repl API versioning:</p> <ul style="list-style-type: none">Fritz Mueller will prepare a write-up in Confluence to propose a solutionthe topic will be discussed next week

Action items

