

Database Meeting 2022-08-17

Date

17 Aug 2022

Attendees

- Igor Gaponenko Andy Hanushevsky Fabrice Jammes Andy Salnikov Joanne Bogart

Notes from the previous meeting (2 weeks ago, before PCW2022):

- Database meeting 2022-08-03

Discussion items

Discussed	Item	Who	Notes
✓	Project news	Fritz Mueller	LSST has a new system commissioning manager: https://www.linkedin.com/in/anastasia-alexov-a4808549/details/experience/ DM JTM in Chile will happen on March 13, 2023. Join Status reviews will be happening on September 14, and 15 at SLAC.
✓	NCSA USDF, IDF	team	<p>Fritz Mueller:</p> <ul style="list-style-type: none">NCSA shutdown has been extended through the end of this week to address a problem with some files lost during file migration to SLACsome issues with storage at SLAC <p>Igor Gaponenko asked a question regarding the possibility of modernizing bbcp. Andy Hanushevsky responded that the new tool should get a better architecture to allow</p> <ul style="list-style-type: none">better memory managementcompression & encryption <p>Igor Gaponenko: Qserv setup on the temporary (loaner) hardware:</p> <ul style="list-style-type: none">6 equally configured nodes32 cores/64 threads: AMD EPYC 7543256 GB of RAM12 TB of NVMe (4 disks configured as RAID0)A snapshot of the <code>large6</code> Qserv was made at NCSA in May 2022. The snapshot will be used for setting up QservThe preliminary plan is to install Qserv using the "igor" mode with 5 workers only (one node will be running Qserv master services (<code>`czar`</code>, <code>czar`</code>'s database, XROOTD redirector, Replication Controller, Replication Worker Registry, and the Replication database). We may afford to lose one worker snapshot (out of 6 available) since the source Qserv was being run at the <code>replication_level=2</code>. <p>Fritz Mueller on the extended use of this instance:</p> <ul style="list-style-type: none">this Qserv may be used for serving DP02 and other catalogs before we'll get the permanent hardware (there is the ongoing procurement process for a batch of the first 15 nodes). <p>Igor Gaponenko: developing Qserv on the VM in IDF</p> <ul style="list-style-type: none">Spun up the VM of 8 vCPUs, 64 GB of RAM, and 80 GB of the "balanced" disk (HDD plus memory-based filesystem)The VM runs CentOS7.Setup the environment (Docker, Docker compose, and a bunch of the missing Python3 modules)XRootD doesn't work on CentOS9. All services crash<ul style="list-style-type: none">Igor Gaponenko will figure out how to retain the core files in the <code>docker-compose</code> based environment and work with Andy Hanushevsky on the crash.Andy Hanushevsky has suggested upgrading XRootD to a version 5.5 that supports CentOS9. This version will be available end of this month (August). Though 5.4.3 should also work.

✓	DP02	Fritz Mueller	<p>Fritz Mueller on ForcedSourceOnDiaObject :</p> <ul style="list-style-type: none"> the new version of the table has passed the Q&A Fritz Mueller is going to extend TAP to make the table visible through the service. See details in the PR; <ul style="list-style-type: none"> https://github.com/lsst/sdm_schemas/pull/86 <p>Fritz Mueller on the new version of the truth tables:</p> <ul style="list-style-type: none"> the tables are deployed and tested in <code>qserv-int</code> we should wait before Colin Slater finishes Q&A-ing the tables after that Igor Gaponenko will deploy the tables in <code>qserv-prod</code> ... unless there will be any issues with the tables. <p>There are 3 still issues with Qserv that need to be addressed and deployed in IDF:</p> <ul style="list-style-type: none"> Case sensitivity (the directory columns of the <code>RefMatch</code> the table should be treated by Qserv as case insensitive) The connection timeout issue of 8 hours still exists. <code>mysql-proxy</code> is suspected. Or maybe in the front-end MariaDB Cancel queries when the client connection goes off (due to the timeout?) <p>Andy Salnikov on the disconnects:</p> <ul style="list-style-type: none"> The <code>Control-C</code> is intercepted by <code>mysql</code> and the latter is opening a separate connection to <code>KILL <id></code> for the relevant. we may need to fix <code>mysql-proxy</code> to detect (client disconnects) and cancel the queries. The proxy may allow specifying user callbacks for handling such disconnects. <p>Fritz Mueller: eventually we need to replace this aging proxy with a better one. This is one of the high-priority architectural improvements for the "new" Qserv. At some point, we will have an in-person dedicated discussion on the subject between Fritz Mueller, Andy Salnikov, John Gates, and Igor Gaponenko.</p>
—	qserv-operator, qserv-ingest	Fabrice Jammes	Any news?
—	Obscure whole-visit s_regions	Gregory Dubois-Felsmann , Andy Salnikov	

Action items

