

Database Meeting 2022-03-30

Date

30 Mar 2022

Attendees

- Igor Gaponenko Fritz Mueller Unknown User (npease) Andy Salnikov Andy Hanushevsky Colin Slater Joanne Bogart

Goals

- Please, register the topics below

Discussion items

Time	Item	Who	Notes
	Project news	Fritz Mueller	
	Progress on topics discussed at the previous meeting Database Meeting 2022-03-23	team	<p>DP02:</p> <ul style="list-style-type: none">• an issue with the schema capitalization changes was resolved in TAP by Fritz Mueller <p>Igor Gaponenko will look at converting Parquet files of the additional metadata tables and get them ingested into Qserv. (Note a discussion on handling date/time types during the conversion into CSV).</p> <p>Colin Slater will finalize the Felis schema for the new tables.</p>
	Upgrade "small" Qserv at NCSA to the latest version		<p>Igor Gaponenko will do it tomorrow during the weekly upgrade</p>
	Problems at IDF after MariaDB upgrade	Fritz Mueller , team	<p>Context:</p> <ul style="list-style-type: none">• Qserv instances qserv-int and qserv-prod were recently upgraded to a newer version of the Qserv operation which was configured with a newer version of MariaDB. The previous version was 10.4, and the newer one is 10.6. There were problems during the database upgrade that were reported at https://lsstc.slack.com/archives/G2JPZ3GC8/p1643844675437739. A solution was deployed to address the issues.• After the seemingly successful start of Qserv, the server (eventually) began experiencing problems (crashes of the Qserv workers and replication services)• Google GKE <i>rolling upgrades</i> seemed to break havoc here by not properly shutting down MariaDB instances. These upgrades cause various stability issues in Qserv. <p>MariaDB posted complaints:</p> <div><pre>2022-03-30 17:32:43 102 [ERROR] InnoDB: Your database may be corrupt or you may have copied the InnoDB tablespace but not the InnoDB log files. Please refer to https://mariadb.com/kb/en/library/innoDB-recovery-modes/ for information about forcing recovery. 2022-03-30 17:32:43 110 [ERROR] InnoDB: Page [page id: space=14, page number=95975] log sequence number 9768674933 is in the future! Current system log sequence number 275818882.</pre></div> <p>These were reported by both Qserv worker databases and the Replication system's database service.</p> <p>Fritz Mueller will disable rolling upgrades in qserv-prod as a short-term solution. Then we'll keep monitoring both prod and int to see if that will make any difference.</p> <p>Fritz Mueller Igor Gaponenko will continue investigating the origin of the problem to understand its root cause.</p> <p>In the long term, we need to find a solution to this situation (of the rolling upgrades). And we do need to develop a policy for MariaDB version upgrades.</p> <p>The worst-case scenario would be to recreate Qserv and to reingest data.</p> <p>Unknown User (npease) will redeploy the Watcher service at IDF to allow seeing problems with Qserv earlier.</p>

	Qserv modality and status in Kubernetes	Fritz Mueller	<p>Context:</p> <ul style="list-style-type: none"> Proposals: Database schema initialization One of the major complications is a requirement that the Replication Worker services shared the same pod with the worker MariaDB services. <p>One option mentioned at the meeting was to allow mounting of the same PV in two pods: as <i>read-only</i> PVC in Qserv workers and <i>read-write</i></p>
--	---	-------------------------------	---

Action items

