

# Database Meeting 2022-01-05

## Date

05 Jan 2022




## Attendees




- Igor Gaponenko Fritz Mueller Andy Salnikov Fabrice Jammes Unknown User (npease) John Gates Kian-Tat Lim

## Goals

- Discuss Fritz Mueller's proposal for configuring Qserv containers.

## Discussion items

Time	Item	Who	Notes
	A proposal for configuring Qserv containers	Fritz Mueller	<ul style="list-style-type: none"><li>Fritz Mueller presented a proposal: <a href="#">DRAFT: configuring Qserv containers</a></li><li>Kian-Tat Lim prepared an updated version of the proposal based on the subsequent discussion among the team members: <a href="#">DISCUSSED: configuring Qserv containers</a></li></ul> <p>Ongoing tickets that are relevant in this context:</p> <ul style="list-style-type: none"><li> <a href="#">DM-33080</a> - Jira project doesn't exist or you don't have permission to view it.</li><li> <a href="#">DM-33082</a> - Jira project doesn't exist or you don't have permission to view it.</li><li> <a href="#">DM-33081</a> - Jira project doesn't exist or you don't have permission to view it.</li></ul>

Configuring and adding workers to Qserv cluster	<a href="#">Fabrice Jammes</a> <a href="#">Igor Gaponenko</a>	<p><a href="#">Fabrice Jammes</a> raised the topic of updating a configuration of the Replication/Ingest system at run time. This is needed for two reasons:</p> <ul style="list-style-type: none"> <li>• for registering workers at the startup time of Qserv</li> <li>• for scaling up an existing cluster</li> </ul> <p><a href="#">Igor Gaponenko</a> reported that there is an ongoing effort to improve the situation here. The first step is to migrate worker services (specifically <b>cmsd</b> and the replication system's worker) to self-configure themselves (learning their identities) from the unique (UUID-generated) dataset identifiers stored in the corresponding Qserv worker databases. For further details and the current status of this development see:</p> <div> <ul style="list-style-type: none"> <li>•  <a href="#">DM-32810</a> - Jira project doesn't exist or you don't have permission to view it.</li> </ul> </div> <div> <ul style="list-style-type: none"> <li>•  <a href="#">DM-32885</a> - Jira project doesn't exist or you don't have permission to view it.</li> </ul> </div> <p>The second step will be to make changes within the Replication system's communication network to allow workers to log into a (yet to be implemented) <i>redirector</i> service. This will reverse dependencies within the system and eliminate a need for the explicit configuration of the workers. A preliminary plan for this development was discussed between <a href="#">Igor Gaponenko</a>, <a href="#">Fritz Mueller</a> and <a href="#">Andy Salnikov</a> before the Winter break. This project is still at an early stage. The actual work on it will start after <a href="#">Fabrice Jammes</a> will finish migrating Qserv to the <i>lite</i> containers and their entry points.</p>
Schema initialization and migration	<a href="#">Fritz Mueller</a> <a href="#">Fabrice Jammes</a>	<p>The topic was just briefly mentioned in the context of the Qserv configuration discussion as there is an overlap between both. It was decided to postpone the discussion till the next meeting.</p>
Status report on testing	<a href="#">Igor Gaponenko</a> <a href="#">John Gates</a>	<p>Lockups are seen in the latest version of the branch when testing mixed query loads in the <b>large</b> Qserv cluster at NCSA. Two types of queries are launched simultaneously in this round of tests:</p> <ul style="list-style-type: none"> <li>• one or two unconditional queries like <code>SELECT * FROM database.table LIMIT 1</code></li> <li>• 100 or 200 of the near neighbor queries, each covering from 1 to 7 chunks</li> </ul> <p>The lockup is happening shortly (a few minutes) after launching the queries. The problem is reproducible.</p> <p>Details were posted in the last comment to the ticket:</p> <div>  <a href="#">DM-31537</a> - Jira project doesn't exist or you don't have permission to view it. </div> <p>The direct link to the most relevant comment: <a href="https://jira.lsstcorp.org/browse/DM-31537?focusedCommentId=443619&amp;page=com.atlassian.jira.plugin.system.issuetabpanels%3Acomment-tabpanel#comment-443619">https://jira.lsstcorp.org/browse/DM-31537?focusedCommentId=443619&amp;page=com.atlassian.jira.plugin.system.issuetabpanels%3Acomment-tabpanel#comment-443619</a></p>

## Action items

- ☐ [John Gates](#) will work with [Igor Gaponenko](#) (if needed) to investigate the lockups.
- ☐ [Fritz Mueller](#) will lead a discussion for initializing and upgrading Qserv schemas at the next meeting. This will be preceded by a discussion among interested members of the group at the team's Slack channel.
- ☐ [Igor Gaponenko](#) will be looking at migrating the configuration system of the Replication/Ingest system from the database tables to a more conventional technique.
- ☐

- [Fabrice Jammes](#) will work on finalizing migrating the operator-based Qserv deployment tools to the *lite* containers and the new configuration model.
- ☐ [Unknown User \(npease\)](#) will finish improving the parameter handling in the entry points as per [Fritz Mueller](#)'s proposal.