Configuring Replication

Types of Replication

This page describes ClustrixDB replication, which uses a single stream to write and read binlogs. ClustrixDB also supports parallel replication, which is more performant and can be used to replicate to/from ClustrixDB only.

ClustrixDB supports the MySQL's replication protocol as both a Slave and a Master.

- ClustrixDB supports replication to/from MariaDB 10.3 with no special configuration.
- ClustrixDB supports replication to/from MySQL 5.7 with GTID mode disabled.

Supported Topologies

Here is a sampling of the topologies for replication supported by ClustrixDB:

ClustrixDB can serve as a Master or a Slave to MariaDB 10.4, MySQL 5.7, or ClustrixDB:

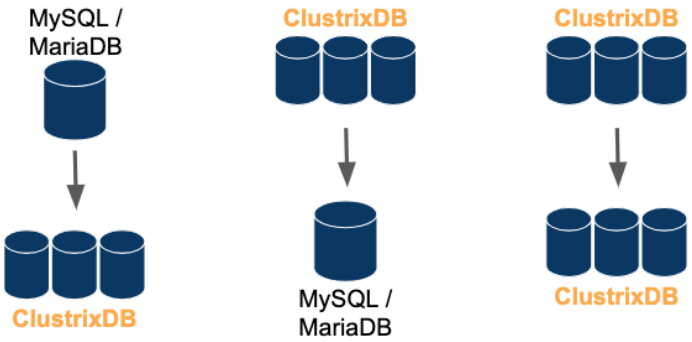

ClustrixDB also supports slaving from multiple distinct binlogs (consolidation) and Master-Master replication:

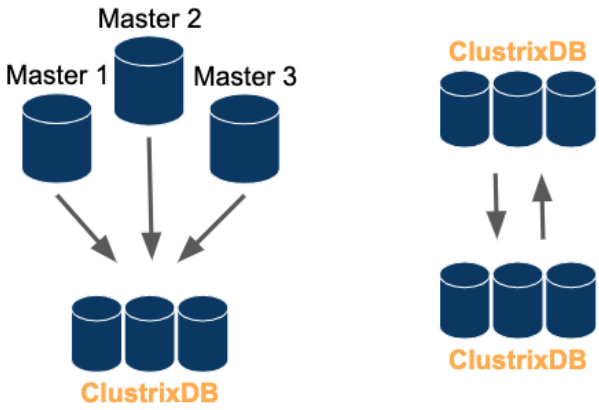

ClustrixDB also supports the creation of multiple binary logs (binlogs), each of which can correspond to specific databases and can be accessed by distinct replication Slaves. When running multiple-Slave configurations, ClustrixDB takes advantage of each node by load-balancing replication connections in a round-robin fashion. For assistance in determining the best topology for your application or whether a specific topology is supported, contact Clustrix Support.

Caveats for Replication:

- ClustrixDB does not support GTIDs

- ClustrixDB does not support ring replication topologies

See also section on Replication concerns when performing Online Schema Changes.

To ensure correct matching of Master and Slave data when replicating between ClustrixDB and a MySQL instance, set both DBMS' to the same time zone. If time zones differ, you risk mismatched data.

The following topics explain how to configure ClustrixDB replication:

- Using ClustrixDB as a Replication Slave
- Using ClustrixDB as a Replication Master
- Configuring Replication Failover
- Parallel Replication