

Using ClustrixDB with Multiple Slaves

ClustrixDB provides sophisticated and scalable support for running many replication slaves on a cluster, including:

- Even distribution of replication slaves among nodes in the cluster.
- Even distribution of multiple slaves on the same node among CPU cores.

Slaves are balanced among nodes following two rules:

1. When a new slave is created, it starts on the node with the least currently enabled replication slaves.
2. After a group change, all slaves restart automatically and will be allocated to nodes according to rule 1 as if they were newly created.

Multiple slaves on the same node are balanced among CPU cores following three rules:

1. When a new slave is started, it starts on the CPU core that has the least number of slaves (either started or stopped).
2. Dropping a slave releases an allocated CPU core, allowing a new slave to take its place (an existing slave will not take its place).
3. After a group change, all slaves restart automatically and will allocate CPU cores according to rule 1 as if they were newly created.

These changes will only show performance improvements on clusters with many replication slaves running concurrently. For example, a 3-node cluster with 3 replication slaves will have 1 slave on each node. A 3-node cluster (each node with 4 CPU cores) with 9 replication slaves will have 3 slaves on each node, with one slave running on each of core 1, core 2, and core 3 (core 0 is not used by replication slaves). A 3-node cluster (each node with 3 CPU cores) with 9 replication slaves will have 3 slaves on each node, with two slaves running on core 1 and one slave on core 2.