

Managing CPU Resources

As a distributed MPP database, ClustrixDB is able to leverage multiple nodes and cores to process queries faster than single-node databases. Two tunable processes are available that facilitate this.

Clustrix has chosen a set of default parameters to control those processes that serve most common workloads. Depending on your specific workload, cluster specifications, and the latency requirements of your application, you may want to tune some of these variables.

Fair Scheduler

The ClustrixDB Fair Scheduler ensures that long-running queries do not monopolize CPU resources. It does this by prioritizing queries that return fewer rows ahead of queries that return a larger number of rows. This helps balance workloads across nodes and is particularly useful for clusters that mix long-running Online Analytical Processing (OLAP) and shorter-running Online Transaction Processing (OLTP).

Fanout

Fanout is the process by which fragments of a query are run simultaneously on multiple CPUs per node. This allows for expanded parallelism and much shorter execution times for ALTER TABLE and large OLAP queries.