



w: www.mariadb.com

e: info@mariadb.com

MySQL Cluster

1. What is it?

MySQL Cluster meets the needs of next-generation web and telecommunications database challenges.

MySQL Cluster is a write-scalable, real-time, ACID-compliant transactional database. It is open source and extensively deployed providing 99.999% availability.

MySQL Cluster has an in-memory, distributed, shared nothing, multi-master architecture with no single point of failure. It horizontally scales on commodity hardware, natively providing auto-sharding to serve read and write intensive workloads as well as offering built-in synchronous replication using a 2-phase commit protocol.

The main features are:

- Auto-sharding for write scalability
- 99.999% availability
- SQL and NoSQL APIs
- Real-time performance
- Multi-site clusters with geographical replication
- Online scaling and schema upgrades

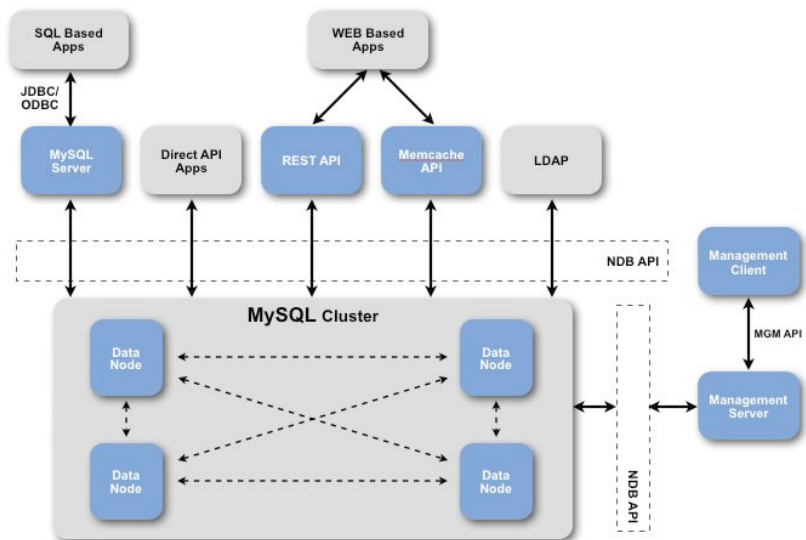
2. How does it work?

MySQL Cluster has a two layer architecture: the interface layer (SQL and NoSQL) and a storage layer. The nodes in the interface layer are called API nodes and the those in the storage layer are called data nodes.



The data is stored in the data nodes which also execute all transactions and associated database operations.

When an API node receives an operation to execute, the node translates it's own protocol to that of the NDB API which is the interface exposed by the data nodes.



One (or more) management server(s) acts as master of the cluster configuration and as arbitrator in the event of failure of one or more nodes.

3. When to use it?

- When you have a huge amount of small transactions
- When you don't have complex queries
- When you need predictable <1ms response times
- When you need 99.999% availability

4. Resources

- Download MySQL Cluster: www.mariadb.com
- Wikipedia: en.wikipedia.org/wiki/MySQL_Cluster