

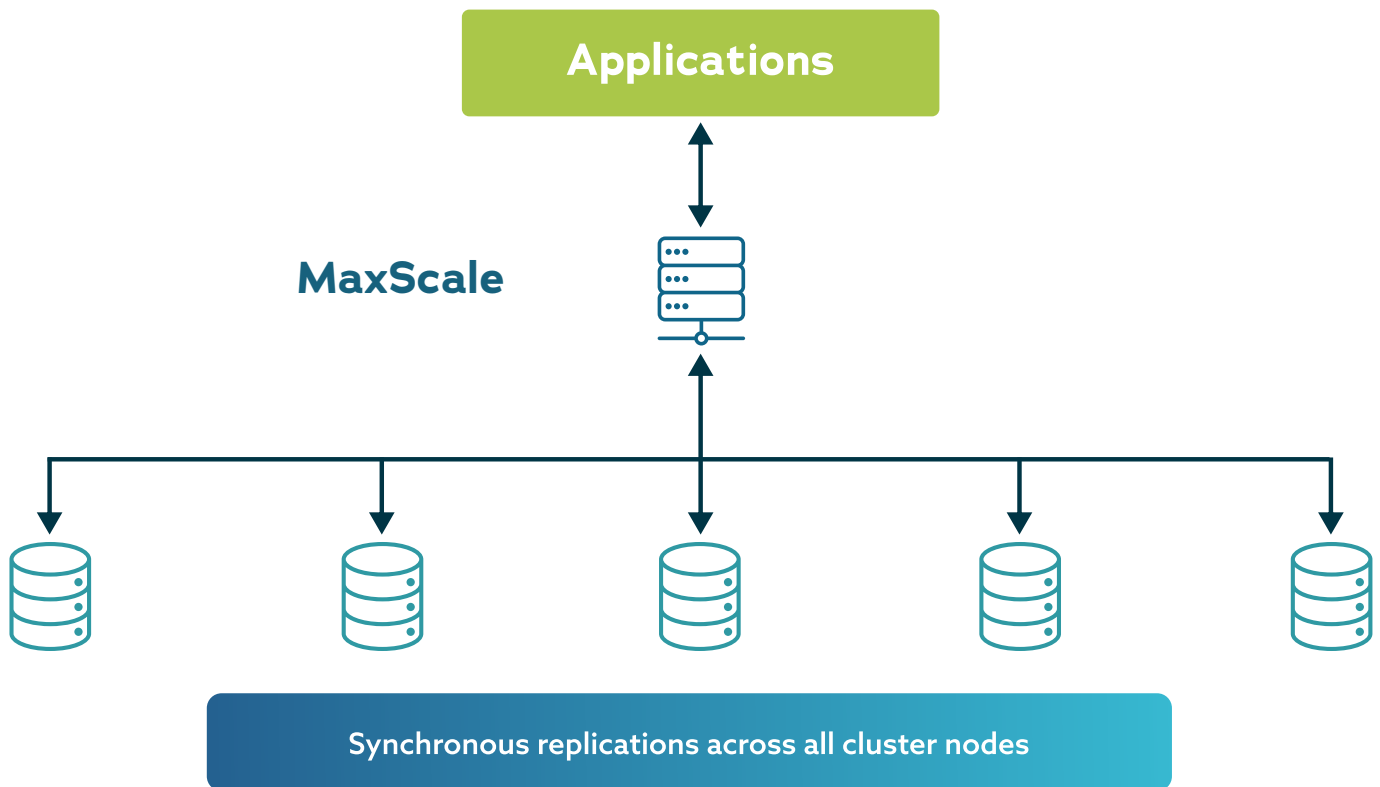
MARIADB ENTERPRISE CLUSTER (POWERED BY GALERA)

Technical Overview

MariaDB Enterprise Cluster offers a multi-primary (active-active) cluster solution for MariaDB, providing high availability, read/write scalability, and true synchronous replication. This means any node can handle read and write operations, with changes instantly replicated to all other nodes, ensuring no replica lag and no lost transactions.

MariaDB Enterprise Cluster requires at least three nodes to avoid split-brain situations. All nodes must be able to communicate with each other over specified ports, and low latency between nodes is ideal. To maintain correct time tracking, all nodes must use the same network time protocol so that they can be properly synchronized.

MariaDB MaxScale plays a crucial role in an enterprise cluster, as it serves as a load balancer and provides automatic failover. In a mission-critical cluster of five nodes, the application first communicates with MariaDB MaxScale, which sends the initial transaction to a node in the cluster. That node processes the transaction through the write set, which it transmits to the other nodes in the cluster. Once the other nodes have processed and acknowledged the transaction, it is committed to all nodes.





FEATURES

Multi-Master

MariaDB Enterprise Cluster is a multi-master system, as all nodes can accept read and write operations and synchronize to ensure data consistency across all nodes.

Fault Tolerant

If one or more MariaDB Enterprise Server nodes fail, the remaining nodes continue to operate. Additionally, the failed nodes are automatically taken offline, repaired, and reintroduced into the cluster, without intervention from the application level.

Geographically Distributed Data

Replicated nodes can be placed across a wide area network (WAN), and, combined with asynchronous replication, improve performance and reduce response times for end-user web services.

Security

By default, Enterprise Cluster replicates data between nodes without encrypting it. This is acceptable when the cluster nodes run on the same host or in networks where security is guaranteed through other means. However, in cases where the cluster nodes exist on separate networks or are in a high-risk network, Enterprise Cluster allows you to encrypt replication data in transit using the Transport Layer Security (TLS) protocol. Enterprise Cluster also supports data-at-rest encryption.

HOW IT WORKS

Enterprise replication works through the following processes:

Write Set Broadcasting

When a client commits a transaction on any node in the cluster, that node, identified as the donor for that specific transaction, captures the changes in a write set associated with that transaction. This write set is then broadcast to all other nodes in the cluster.

Certification and Application

Each receiving node performs a certification test to ensure that the incoming write set does not conflict with any concurrent transactions being committed locally.

If the write set passes certification, it is applied to the local database, and the transaction is committed on that node.

Alternatively, if a conflict is detected, the conflicting transaction, which is usually the one executed locally, is aborted, ensuring data consistency across the cluster.

Virtually Synchronous

Virtually synchronous means that the commit order is globally consistent, and all successful transactions are guaranteed to be applied on all active nodes, even though the actual data application might happen slightly after the commit on the initiating node. A transaction is not truly considered committed until it has passed certification on all nodes.

wsrep API

This API defines the interface between the replication library, known as the wsrep provider and the MariaDB database server.

The API allows the database to expose hooks to capture and apply transaction write sets.

ROLE OF MARIADB MAXSCALE

In combination with MariaDB MaxScale, Enterprise Cluster provides high availability, scalability and automated failover for mission-critical transactional workloads. MariaDB Enterprise Cluster provides parallel replication and data consistency across all nodes, automatically managing the identification and removal of failed nodes, as well as recovering and rejoining new nodes.



USE CASES

Read Master

A traditional primary and replica(s) topology, but with Enterprise Cluster, all nodes are capable of being primary at any time. Replication can guarantee zero replica lag for such installations and, due to parallel replica writes, much better throughput for the cluster.

WAN Clustering

Synchronous replication works fine over the WAN network. There will be a delay, which is proportional to the network round-trip time (RTT), but it only affects the commit operation.

Disaster Recover

Disaster recovery is a subclass of WAN replication. Here, one data center is passive and only receives replication events, but does not process any client transactions. Such a remote data center is always up to date, ensuring that no data loss can occur. During recovery, the spare site is nominated as primary, and the application continues as usual with a minimal failover delay.

Latency Eraser

With WAN replication topology, cluster nodes can be located close to clients. Therefore, all read and write operations will be extremely fast with the local node connection. The RTT-related delay is experienced only at commit time and can generally be accepted by the end user. In this environment, read operations are as fast as possible.

ABOUT MARIADB

MariaDB seeks to eliminate the constraints and complexity of proprietary databases, enabling organizations to reinvest in what matters most – rapidly developing innovative, customer-facing applications. Enterprises can depend on a single complete database for all their needs, that can be deployed in minutes for transactional, analytical, hybrid use cases. Trusted by organizations such as Deutsche Bank, DBS Bank, Red Hat, ServiceNow and Samsung – MariaDB delivers customer value without the financial burden of legacy database providers. For more information, please visit mariadb.com.

mariadb.com

Americas: sales-AMER@mariadb.com

Europe, Middle East, Africa: sales-EMEA@mariadb.com

Asia Pacific: sales-APAC@mariadb.com

© Copyright 2025 MariaDB plc