



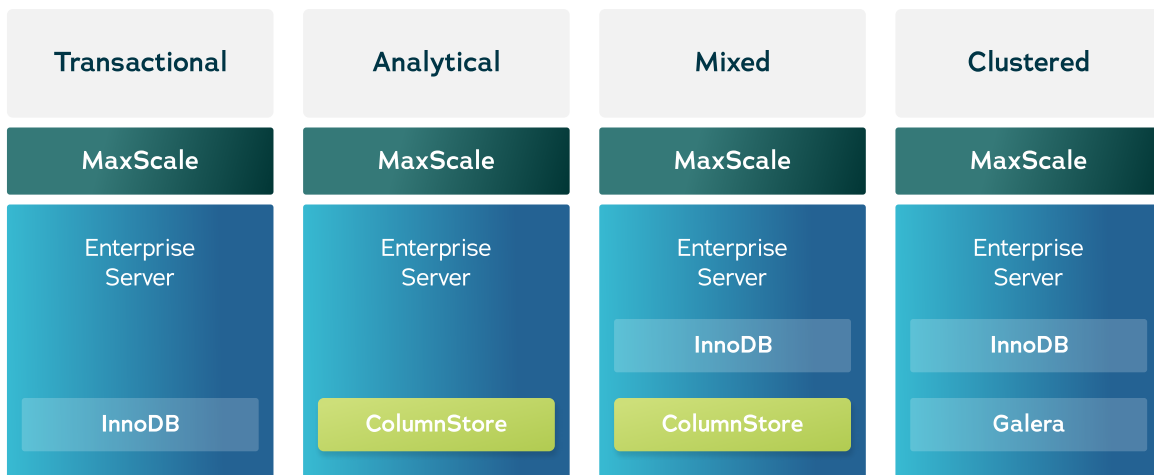
# MARIADB ENTERPRISE SERVER

The complete enterprise-grade database solution - any workload, anywhere

## OVERVIEW

MariaDB Enterprise Server is an open source, hardened version of MariaDB Community Server with additional enterprise features, including MariaDB MaxScale database proxy, Galera cluster, non-blocking backup and advanced audit capabilities. Popular new features are also backported to older MariaDB Enterprise Server versions to bring new features to customers running on older release versions. Enterprise Server uses pluggable storage engines to support transactional, analytical and mixed workloads. Enterprise Server is cloud-native, running on Linux, VMs, or in containers. It can be deployed on premise or self-hosted on any cloud. Companies who switch to MariaDB Enterprise Server from legacy, proprietary databases save up to 90% of total database costs. MariaDB Enterprise Server can start small with a standalone, single-node database, and grow to millions of transactions per second with interactive ad hoc analytics on billions of rows.

### Any Workload, Anywhere



## ENHANCED

MariaDB Enterprise Server includes features engineered for customers deploying and maintaining large databases with strict high availability, disaster recovery and security requirements. For example, to perform frequent backups without impacting applications and enforce full end-to-end encryption.

## HARDENED

MariaDB Enterprise Server undergoes an extensive, thorough and comprehensive QA process to ensure reliability for production deployments. In addition, critical features and bug fixes in future releases are backported to prior non-EOLed versions to ensure long-term stability and support.



## SECURE

MariaDB Enterprise Server is preconfigured for production environments, including default security parameters to enforce account hygiene, remote root and anonymous access, and replication parameters set to enforce durability. In addition, all non-GA plugins are disabled.

## KEY BENEFITS

### Unstoppable Operations

MariaDB Enterprise Server leverages MariaDB MaxScale advanced data proxy to support high availability (HA) features such as massive connection load balancing, primary/secondary node replication, automatic failover, and point-in-time-recovery (PITR) down to last transaction replay with schema rewind. MaxScale provides HA and scalability without the need for developer code, allowing those involved in the DevOps chain to focus on their organization's mission and points of value add.

### Start Small, Scale Fast, Run Anywhere

MariaDB Enterprise Server scales up with more cores, memory and storage on both key architectures, AMD64 and ARM. It can scale out with more database instances – making it easy to start with a standalone instance, run on a bigger server, and add replicas by instance to scale reads. Enterprise Server will run on as little hardware resourcing as a Raspberry Pi, laptop or desktop computer. It will also run on any public or private cloud in VMs or containers starting from 1 vCPU and scaling up to over a thousand.

### Run More Workloads Yet Reduce Database Sprawl

MariaDB Enterprise supports transactional, analytical and mixed workloads for structured and semi-structured JSON data, thereby supporting a vast range of different workloads with pluggable “engines,” making it possible to simplify database infrastructure and management by using a single solution for everything – interactive analytics with columnar storage and schema flexibility with JSON.

### Developer-Centric

MaxScale abstracts away the underlying database infrastructure so developers don't have to worry about handling failures or changes in the topology (planned or not). Instead, developers can leverage modern SQL features in MariaDB Enterprise Server, such as JSON functions, bi-temporal tables, window functions, common table expressions (CTEs) and user-defined aggregate functions. It also supports smart transactions by combining row storage optimized for fast transactions with columnar storage optimized for fast analytics. With smart transactions, developers can enrich web and mobile applications with real-time analytics and historical data to create more insightful customer experiences driven by compelling opportunities – and, for SaaS customers, to deliver self-service analytics.

### Cloud-Native

Enterprise Server with pluggable storage engines runs in containers. It has a small, portable footprint, allowing distributed applications to run across the edge and multicloud, avoiding multiple database architectures. This also reduces the number of separate databases servers needed for a variety of workloads which streamlines configuration and management complexity when using tools like Terraform and Chef. Instant schema changes allow developers to iterate much faster and provide them with greater flexibility while at the same time ensuring data consistency and integrity. Further, Enterprise Server implements a cloud-native storage architecture and can optionally use S3-compatible object storage to lower costs and take advantage of unlimited capacity.

### Ease of Migration Off Expensive Proprietary Platforms

With MariaDB Enterprise Server, you can replace legacy, proprietary databases and harness cloud economics without lock-in to a single cloud service provider or SaaS offering. It supports ANSI SQL, MySQL,

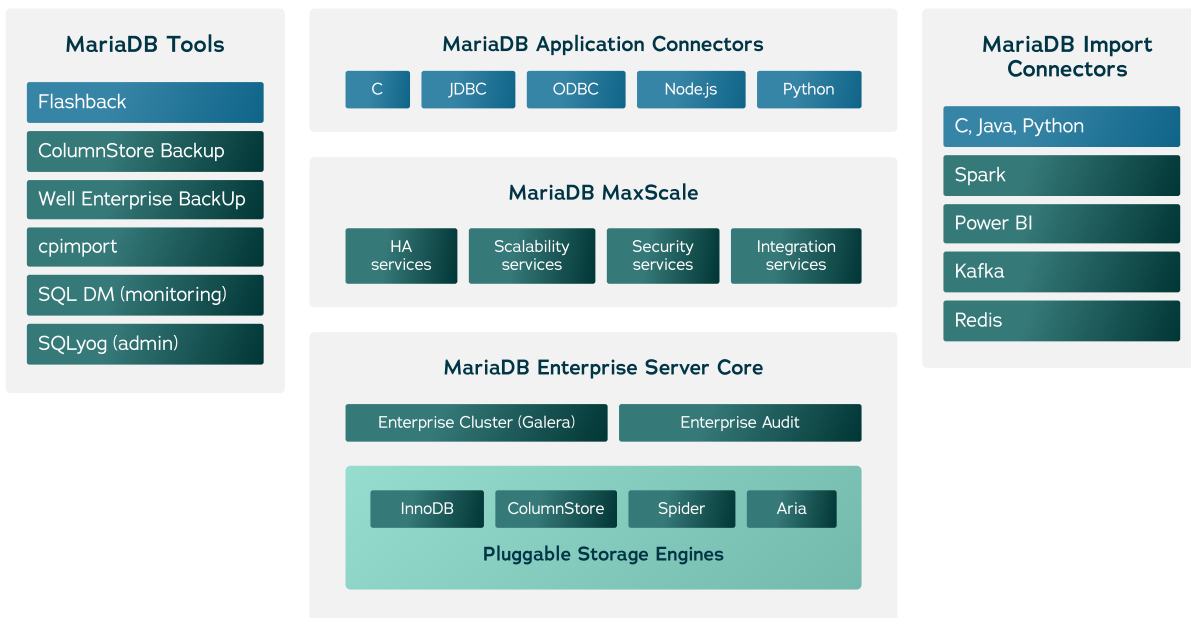


Oracle PL/SQL, SQL Server and Sybase SQL dialects, making it easy for developers to port existing applications with fewer code and schema changes. MariaDB customers have imported hundreds of thousands of lines of legacy database code. MariaDB is the only enterprise database solution with a 100% open source Oracle PL/SQL implementation. It supports standard Oracle database packages and procedures, including cursors, loops, variables, exceptions and functions such as CAST, LENGTH and SUBSTR.

### Innovation without Disruption

Select enhancements from MariaDB Enterprise Server (e.g., non-blocking backups) and MariaDB Community Server (e.g., optimizer trace) are backported to previous versions of Enterprise Server to provide greater long-term stability and support, making the most critical enhancements available without the need to perform major upgrades on an annual basis, avoiding stagnation and facilitating innovation without disruption.

### Product Architecture



### Pluggable Storage Engines

MariaDB Enterprise supports multiple storage engines to optimize for transactional, analytical, mixed and high-throughput workloads. Different storage engines, including InnoDB, ColumnStore, Aria, MyRocks, and Spider, support different use cases, such as federated table access, table sharding and table archiving in the cloud. InnoDB is the most popular general-purpose storage engine for OLTP or transactional workloads, with ColumnStore being the most popular for OLAP or analytical processing workloads. Different tables can use different storage engines, and joins can be performed between tables across storage engines.

### High Availability and Enterprise Cluster

MariaDB Enterprise Server leverages MariaDB MaxScale to transparently run in a primary/secondary node (also known as a replica) configuration for high availability. With read/write splitting, replicas also support a performance boost for read-intensive applications. MaxScale’s management of replicas prevents primaries from being shut down until all replication has been completed to avoid a switchover or manual shutdown from causing data loss. In addition to MaxScale, MariaDB Enterprise Cluster offers a multi-master (also known as multi-writer) option based on Galera Cluster that provides synchronous multi-writer, multi-node replication with built-in monitoring for cluster, node and session status. Enterprise Cluster encrypts transaction buffers and performs non-blocking DDL replication for clusters greater than three nodes, preferably in an odd-numbered configuration.



### **Automatic and Instant Failover**

Enterprise Server with MaxScale is the only database with advanced features comparable to Oracle Application Continuity. Together, they hide infrastructure and database failures from applications, resulting in zero-interruption failover and online disaster recovery, and ensure the database is always available and its data is always protected resulting in continuity for your applications. As part of performing automatic failover, MaxScale migrates backend database connections to available server(s) – applications don't have to create new connections because of a failover. MaxScale also restores database sessions on a newly promoted primary after it has performed an automatic failover – applications don't have to restore sessions because of a failover. Finally, MaxScale replays in-flight transactions after performing an automatic failover so they can continue – apps don't have to retry transactions because of a failover.

### **Analytics and Mixed Workloads**

MariaDB Enterprise Server is the only major open source offering with a pluggable storage engine optimized for ad hoc queries and OLAP, MariaDB ColumnStore. ColumnStore delivers data warehousing capabilities without requiring a separate enterprise data warehouse, ETL, complexity or expense and delivers an order of magnitude faster queries on large datasets. It stores data on Amazon S3 compatible object storage, on-premises, or in the cloud with up to 90% compression, reducing disk I/O and dataset size and eliminating the need for shared storage. Compared to InnoDB-based analytics operations, ColumnStore has fast and direct data ingestion, removing the need for indexes and views. ColumnStore delivers up to orders of magnitude better performance on queries to large sets of data focused on specific ranges of specific columns, depending on the type of SQL operation run; for example, it can perform INSERT cache operations 4700x faster. Cross-engine updates can be performed between ColumnStore and InnoDB. Table partitions can be set up in the latest version of Enterprise Server to support a single table view for these updates.

### **Enterprise Security**

MariaDB Enterprise Server goes above and beyond standard security features such as roles and auditing, with advanced data protection and security layers such as complete, end-to-end encryption. Enterprise Server encrypts all data in motion with secure connections (TLS) and at rest, including logs, with transparent data encryption (TDE). Enterprise Server uses Hashicorp Vault to delegate encrypted table key management, thereby managing encryption keys outside the database.

### **Enterprise Backup and Disaster Recovery**

MariaDB Enterprise Server includes several tools and features for implementing comprehensive disaster recovery plans, including non-blocking backups, and point-in-time recovery, and online point-in-time rollback. In addition, delayed replicas (e.g., an hour, 12 hours and one day) can be used as part of an advanced strategy to prevent accidental or malicious data loss/corruption from reaching backups. Backups for InnoDB include support for full, partial, and incremental backups.

### **Enterprise Audit**

MariaDB Enterprise Audit provides server-based auditing for MariaDB Enterprise connections, queries and objects (databases, tables, etc.), logging activity in a granular and comprehensive fashion. Advanced filtering features enable narrow definitions of which information is logged, giving DBAs broad yet granular control over what gets audited. Auditing is configured with templates stored as JSON documents with options to filter templates per user for databases and tables.

### **Open and Extensible Third-Party Integrations**

MariaDB Enterprise Server provides enterprise federation by enabling high-performance access to tables in remote databases using ODBC connections. In addition to ODBC, MariaDB Connectors are available for C, C++, Java, JavaScript, and Python. Enterprise Server also includes MariaDB Client, which can be used to run queries interactively or within scripts.



Enterprise Server integrates with Apache Kafka and Redis as the foundation of a modern data infrastructure for enterprises, publishing changes to Kafka for stream data processing and integration and caching query results in Redis. This not only improves query performance but reduces the load on the database, so fewer resources are required or can be freed up for other queries. In addition to more efficiently utilizing database resources, enterprises can get more from their existing Kafka and Redis deployments.



## CAPABILITIES

### HIGH AVAILABILITY

- Multi-writer clustering (Galera)
- Lossless replication (MaxScale)
- Automatic failover
- Last transaction replay
- Transparent query routing

### SCALABILITY

- Sharding
- Compression
- Read-write splitting
- Distributed

### ADVANCED SECURITY

- Encrypted transaction buffers
- Password expiration and reuse prevention
- Disabled user accounts
- Pluggable authentication
- Roles and user resource limits
- Transparent data encryption (TDE)
- Auditing with configuration logging
- Query result limiting (i.e., DoS protection)
- Connection attempt throttling
- Hashicorp Vault plugin
- Enforced SSL/TLS connections

### SCHEMA

- Bi-temporal modeling for historical data
  - Application-period time, system versioned and combined bi-temporal tables
- Instant ADD/DROP/MODIFY COLUMN
- Invisible columns
- CHECK constraints
- Default value functions/expressions
- Multiple triggers per event
- Virtual column indexes
- Spatial indexes
- Descending and composite indexes
- Decimal scale of 38
- Federated tables via ODBC

### DISASTER RECOVERY

- Non-blocking backups
- Point-in-time rollback
- Delayed replication

### PERFORMANCE

- Query optimizer, cost-based with trace
- Non-locking ANALYZE TABLE
- Thread pool
- Query result caching
- Bulk insert streams

### STANDARD SQL

- Common table expressions (recursion option)
- Extensive JSON functions
- Window functions
- Ordered-set aggregate functions
- User-defined aggregate functions
- Percentile and median window functions
- Table value constructors
- INTERSECT/EXCEPT ALL
- INSERT/REPLACE... RETURNING

### POWERFUL ANALYTICS

- Distributed, columnar storage
- Massively parallel processing with AMD64/ARM
- Statistical functions (e.g., CORR)
- Apache Spark connector
- Apache Kafka connector
- Pentaho connector
- C, Java and Python import clients
- Shared-nothing storage
- Object storage support (Amazon S3 compatible)
- Amazon S3 bulk high-speed imports

### ORACLE DATABASE COMPATIBILITY

- Data types and sequences
- Dynamic SQL and cursors with parameters
- Stored procedures and packages

---

[mariadb.com](https://mariadb.com)

Americas: [sales-AMER@mariadb.com](mailto:sales-AMER@mariadb.com)  
Europe, Middle East, Africa: [sales-EMEA@mariadb.com](mailto:sales-EMEA@mariadb.com)  
Asia Pacific: [sales-APAC@mariadb.com](mailto:sales-APAC@mariadb.com)

© Copyright 2023 MariaDB plc  
699 Veterans Blvd.  
Redwood City, California, USA.