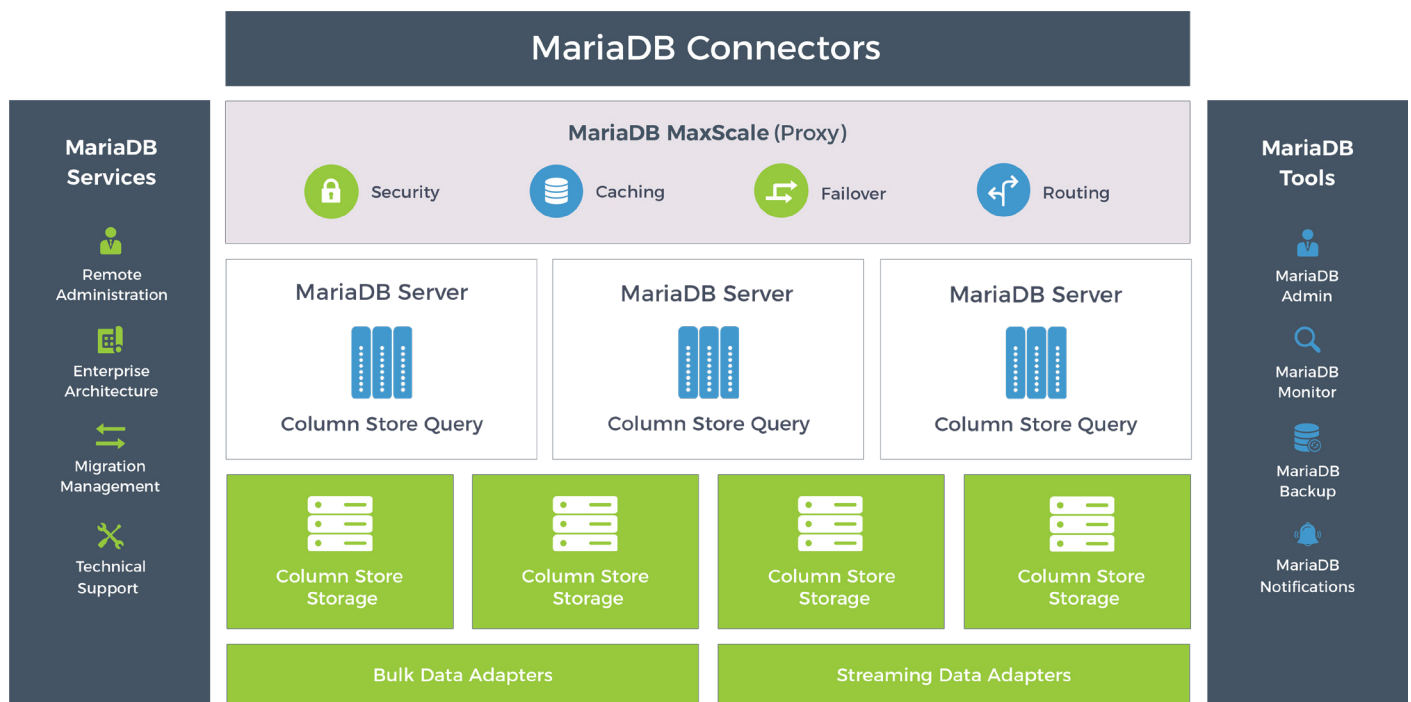


MariaDB AX

Born of the community. Raised in the enterprise.

MariaDB AX is modern analytics and data warehousing for everything from business intelligence and reporting to predictive analytics made easy – and is accessible to everyone who needs to get from data to insight faster.

MariaDB AX is an enterprise open source solution for scalable, high-performance analytics and faster time-to-insight, simplifying the process of ingesting and analyzing data – at any scale, and without the need for complex, time-consuming batch processes, schemas and indexes optimized for a small number of predefined queries or expensive, proprietary appliances.



Core components

MariaDB Server: A modern, general-purpose database engineered for performance, reliability, security and extensibility with a modular architecture.

MariaDB ColumnStore: A specialized storage engine for scalable, high-performance analytics without the need for complex schemas and indexes – distributed, columnar storage with parallel query processing.

Bulk data adapters: C++, Java and Python libraries for enabling developers and data scientists to collect and write data to storage on demand, bypassing the SQL interface for higher throughput.

Streaming data adapters: Apache Kafka* and MariaDB MaxScale services for importing data from MariaDB TX or Apache Kafka in near real time – automatically, transparently and continuously.

NOTE *The Apache Kafka streaming data adapter is beta

Benefits and features

Analytics made flexible

The use of traditional data warehousing models (e.g., star/snowflake) is no longer required. The columnar storage engine enables the same data to be queried in different ways, and without indexes

Analytics made fast

There is no need to wait when millions or even billions of rows can be queried in a matter of seconds. The columnar storage engine is optimized for analytical queries – queries on most, if not all, data.

Analytics made simple

The use of batch processing to import data is no longer required – developers and data scientists can collect and write data from C++, Java and Python applications using bulk data adapters, bypassing the SQL interface.

Analytics made easy

There is no need to learn a new language or become familiar with a new type of database. MariaDB AX is built on MariaDB Server. The data is analyzed with standard SQL, and with the full power of SQL – no limitations.

Analytics made affordable

There is no need for expensive hardware or proprietary appliances. The database and storage nodes are engineered to run on existing commodity hardware, on premises or in the cloud – public, private or hybrid.

Analytics made scalable

The storage and processing capacity can be increased by scaling up or scaling out. The data and queries are automatically partitioned and distributed when running multiple nodes – to query more data, add more nodes.

Analytics made current

There is no need to wait for new data. The streaming data adapters for Apache Kafka* and MariaDB MaxScale change-data-capture streams can be used to automatically and continuously import new data.

Analytics made powerful

The analytical capabilities can be extended to support advanced analytics. The user-defined aggregate and window function API enables developers and data scientists to create custom analytical and statistical functions.

Innovative use cases

Finance	Identify trading patterns, detect fraud and anomalies
Healthcare	Identify population cohorts, predict health outcomes
Telecommunications	Network and customer behavioral analysis
Manufacturing	Detect production anomalies, predict machine failures

NOTE

*The Apache Kafka streaming data adapter is beta

Website: mariadb.com

Americas: sales-AMER@mariadb.com

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

Asia Pacific: sales-APAC@mariadb.com

© Copyright 2017 MariaDB Corporation Ab, Tekniikantie 12, 02150 Espoo, Finland. MariaDB is a trademark or registered trademarks of MariaDB Corporation.