

# MariaDB ColumnStore

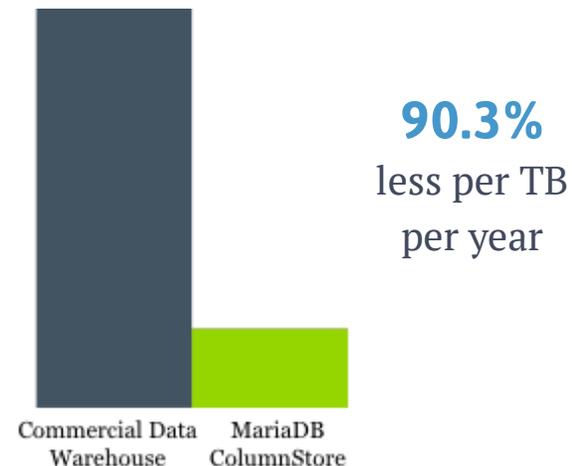
Information is one of the most important competitive advantages in the Digital Economy. Organizations need to better predict what their customers will buy based on past behavior or analyze weather trends so that they can optimize energy consumption. In the midst of major trends like Big Data and IoT, analytics is the critical technology needed to bring value out of data.

The concept of analytics has expanded over the years. It has gone way beyond dashboards and reporting, to predictive insights from data that impact the top and bottom line. Delivering business value is the metric of effective analytics.

Existing approaches to analytics have not yet reached their full potential. Traditional data warehouse vendors pose significant barriers with high costs from licenses and hardware, and rigid architectures that make it difficult for users to justify ROI. Many open source vendors responded with new approaches to analytics such as Hadoop or NoSQL, but those products are niche, supporting only a subset of typical enterprise-class use cases. Also, their limited SQL support makes it hard to find talent capable of supporting and maintaining such specialized environments.

MariaDB is introducing an easier and faster approach to big data analytics. MariaDB ColumnStore is a powerful open source columnar storage engine that supports a wide variety of analytical use cases with ANSI SQL in highly scalable distributed environments.

## Better Price Performance



MariaDB ColumnStore on average costs 90.3% less per terabyte per year compared to commercial data warehouses. Furthermore, ColumnStore provides high compression rates for big data, which allows you to store more data with less hardware. ColumnStore, whether deployed on premise or in cloud, is much more cost effective.

## Easier Enterprise Analytics

Analytic workloads are more read intensive compared to transactional workloads which are more write intensive. With this difference in mind, MariaDB ColumnStore is designed as a pluggable storage engine purpose-built to handle analytic workloads while keeping the same ANSI SQL interface that is used across all the storage engines of the MariaDB portfolio, thereby making it easier for your existing technical resources to adopt and maintain.

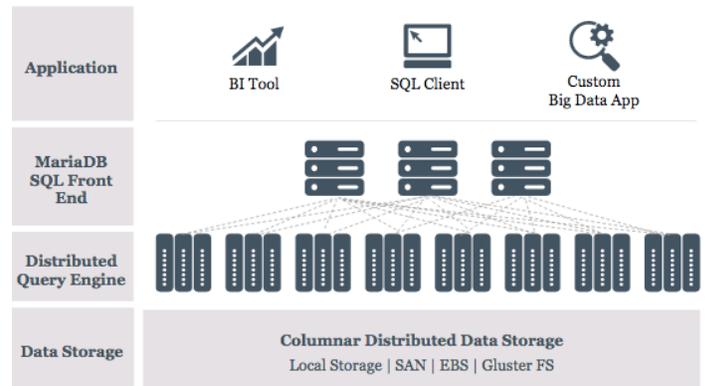
For complex analytics, MariaDB ColumnStore supports queries like complex joins, aggregation, window functions and user-defined functions. To apply these analytics on large datasets, MariaDB ColumnStore's massively parallel and distributed architecture processes millions to trillions of rows in a single query running on gigabyte to petabyte-size datasets. It also eliminates the need to maintain indexes and views and includes automatic partitioning. Its elastic scalability makes it easy to add new nodes to support growth in data volume or analytics needed to run a cluster without interrupting ongoing queries.

By leveraging MariaDB's enterprise-grade capabilities, ColumnStore provides best-in-class security features like encryption for data in motion, role-based access and auditability. ColumnStore provides out-of-the-box connections with the BI tools of your choice.

## Faster, More Efficient Queries

Getting real-time insight is critical in the Digital Economy. However, when those insights are analyzed from petabytes of data, it requires both fast data ingesting and big data processing capability. To provide faster time to insight, ColumnStore optimizes both fully parallel data ingestion, which can process 500K rows per second. MariaDB ColumnStore can also process big data with high speed by storing data in a columnar-oriented format which scans the data directly from the area where the data in a certain column is stored rather than looking up each record row by row.

## ColumnStore Architecture



The columnar format and its built-in compression algorithm are designed to reduce disk I/O for high-performance read queries. MariaDB ColumnStore's distributed query processing further accelerates performance of the read-intensive analytic workloads. As analytics become a core component of data-driven business, high availability of the analytics environment becomes an essential requirement. ColumnStore provides an automatic failover feature to ensure your analytics environment is 100% up and running.

<b>Supported OS</b>	<ul style="list-style-type: none"> <li>CentOS 6, CentOS 7</li> <li>Debian 8</li> <li>Ubuntu 16.04</li> </ul>
<b>Cloud Consumption Options</b>	<ul style="list-style-type: none"> <li>AWS AMI</li> </ul>
<b>Developer Consumption Options</b>	<ul style="list-style-type: none"> <li>Docker</li> <li>Vagrant</li> </ul>