



IMPACT REPORT

# MariaDB eyes 'time to insight' with focus on data ingestion for AX database

FEBRUARY 27 2018

BY JAMES CURTIS (/ANALYST-TEAM/ANALYST/JAMES+CURTIS)

MariaDB released an updated version of its analytics database MariaDB AX. Headlining the new version are new data adapters to increase data ingestion and subsequently the ability to lower the window from when the data lands to when users can act on it. The release of MariaDB AX arrived after the company rolled out MariaDB TX (/report-short?entityId=92820) for operational workloads, effectively segmenting its products between analytical workloads (MariaDB AX) and transactional and operational workloads (MariaDB TX). The database vendor also pulled in a new funding round of \$27m in November 2017 and is using this to supplant development efforts.

## The 451 Take

The company's present product strategy of segmenting its offerings for transactional and analytical workloads has been in the works for some time now. The bifurcated product strategy not only gives clarity and eases the adoption for organizations, but it also gives MariaDB a focus for analytics functionality developments. Case in point is the collection of data ingestion adapters, which means the data can be loaded into MariaDB and be available immediately for analysis. When both systems are paired, mixed workloads are an option as well. Also, enhanced high availability and backup are a

good start for driving enterprise functionality. Certainly, an analytics product gives MariaDB broader appeal, but if competition in the operational space wasn't enough, the analytical is just as competitive, if not more so.

## Context

MariaDB can point to MySQL as its genesis. Several people who were previously at MySQL founded MariaDB and the company's team forked the MySQL code base after Oracle bought Sun Microsystems, which had itself bought MySQL. However, while MariaDB offers an alternative to the open source MySQL database, it should be considered a new database completely independent from MySQL, albeit with full MySQL compatibility and also support for the full ANSI SQL standard.

The latter part of 2016 and into 2017 was quite an active time for MariaDB as the company looked to solidify its strategy by simplifying and consolidating its product portfolio. In December 2016, MariaDB rolled out MariaDB ColumnStore (</report-short?entityId=91400>), a pluggable storage engine purpose-built to handle analytical workloads while also maintaining ANSI SQL compliance. Then, in May 2017, MariaDB released an updated version of MariaDB TX (</report-short?entityId=92820>), the company's transactional database offering, which it claimed as a feature-rich open source transactional database and proxy and on par, feature-wise, with more expensive proprietary offerings for most needs. In November, the company announced an update to MariaDB AX, which, when paired with the MariaDB ColumnStore, provides a type of bundled analytical database offering centered around the core piece of MariaDB Server.

The company's current portfolio consists of two primary offerings: MariaDB TX and MariaDB AX. Organizations can choose the type of workload – transactional or analytical – and then choose the database platform. Previously, organizations had to choose among various MariaDB products and self-assemble the database platform to a particular workload.

The new release of MariaDB AX in November 2017 also coincided with the company's announcement of an additional \$27m in funding, taking its total funding haul to \$54m. The recent funding round was led by Alibaba Group and included participation from Intel Capital, California Technology Ventures, Tesi,

SmartFin Capital and Open Ocean Capital.

## Technology

The company's MariaDB AX database product is an amalgam of other MariaDB products that when assembled form the basis for the MariaDB AX analytic database offering. At the core of the offering is the MariaDB Server, an open source database server that handles all the query processing for MariaDB AX. MariaDB ColumnStore functions as the analytic storage engine (as opposed to InnoDB as the transactional storage engine for MariaDB TX) and integrates with MariaDB Server to maintain the storage separately.

Separating storage from the server provides a few benefits. One is the ability to scale both functions separately based on the dataset size and the specific workload. Another benefit is that ColumnStore provides much greater data compression as well as higher overall disk I/O performance.

MariaDB MaxScale makes up the third component of MariaDB AX and functions as the database proxy. The company says that MaxScale offers more than a standard database proxy. MariaDB explains that MaxScale plug-ins can optimize the scalability and availability of a database cluster, secure it, and be used to manage any maintenance downtime. Thanks to the extensible MaxScale architecture, it's said to be easy to extend with custom plug-ins that allow it to handle new tasks.

## Products

For its latest release of MariaDB AX, the company beefed up its ingestion capability. More specifically, MariaDB added bulk data adapters for C++, Python and Java, with others in the works. The adapters are able to collect and write data directly to ColumnStore, bypassing the database server. Besides providing the ability to take in data faster (no need to go through the SQL parser), the data can be available almost immediately for analytical operations, which is further aided by the fact that data doesn't need to be indexed.

Further, MariaDB has added two streaming adapters. The first one leverages the company's recently developed change data capture client that integrates with MaxScale. Like the bulk adapters, the stream adapter is written directly to ColumnStore but the data is being streamed directly from MariaDB TX (InnoDB) to MariaDB AX. As such, the approach requires both systems but also provides automatic and continuous data updates as there is no need to worry about manually syncing the data from MariaDB TX.

The second streaming adapter is specific to Apache Kafka. Likewise, MariaDB created a Kafka client connector that will attach to Kafka topics and then insert those topics into the ColumnStore.

In addition to the data adapters, the company has improved the analytical capabilities within MariaDB AX. Specifically, MariaDB has an API that can be leveraged to write user-defined aggregate and window functions. The new API essentially enables customized functions to be developed and carried out within ColumnStore.

Enhanced high availability and disaster-recovery functionality have also been added to MariaDB AX with support for GlusterFS. Because ColumnStore works with folders, previously, if a node failed, all the folders and subsequent files would be lost. With GlusterFS support, files can be attached from one to another, thus eliminating the need to set up a separate SAN. Moreover, MariaDB developed a new tool that automatically determines the system topology and can restore data to available nodes, where before it was carried out manually to find any available storage nodes.

## Competition

Given MariaDB's historical roots, it makes sense that MySQL, which is owned by Oracle, would be the firm's closest competitor at least when MariaDB TX is deployed. With MariaDB bifurcating its product line into in transactional and analytical systems, the competitive field has expanded.

Oracle, IBM, and Microsoft all provide relational systems that can address analytical workloads. Specifically, these offerings include Oracle Exadata, IBM Db2 iterations such the IBM Integrated Analytics System, and Microsoft's Analytics Platform System. Still others with relational roots that

provide analytical capabilities include Actian with its Vector database, MemSQL and PostgreSQL, including PostgreSQL supporter Crunchy Data.

There is also a host of other vendors that offer analytic-based systems, sometimes referred to as data warehousing systems. These include Teradata, Micro Focus with Vertica, and Pivotal with Greenplum, as well as cloud offerings such as Snowflake Computing, Google BigQuery, Amazon Web Services with Redshift and Athena, Microsoft Azure SQL Data Warehouse and IBM Db2 on Cloud.

With MariaDB's data adapters, the notion of mixed workloads may also come into play and something that we have written on and refer to as hybrid operational and analytics processing or HOAP (/report-short?entityId=93844), for short. Some vendors driving these systems include SAP HANA, MemSQL, VoltDB, NuoDB, Clustrix and Splice Machine.

## SWOT Analysis

### Strengths

---

MariaDB has grown and fostered a strong community around its open source offering and with the new data adapters, organizations can address additional use cases, including the ability to handle mixed workloads.

### Opportunities

---

With its bifurcated product line that leverages MariaDB Server, the company is establishing a good position to address several different use cases and workloads for organizations. For instance, the new

### Weaknesses

---

While MariaDB AX gives the company an analytics product, it is still relatively new and maturing, which includes improved high availability and backup, as well the new data adapter tools.

### Threats

---

The relational database market is quite competitive but the analytics systems space is just as competitive if not more so. Further, MariaDB faces challenges with organizations that may prefer commercial software instead of open source.

data adapters not only improve steaming functionality but can also provide the ability to run mixed workloads.

## James Curtis (/analyst-team/analyst/James+Curtis)

Senior Analyst, Data Platforms & Analytics

### M&A ACTIVITY BY SECTOR

Information management / Data management / Data warehousing (97)  
 (https://makb.the451group.com/results?basic\_selected\_sectors=117)

### M&A ACTIVITY BY ACQUIRER

Action Corp [fka Ingres] [Garnett & Helfrich Capital] (6)  
 (https://makb.the451group.com/results?basic\_acquirers=Action+Corp [fka Ingres] [Garnett & Helfrich Capital])

Alibaba.com [aka Alibaba Group] (10)  
 (https://makb.the451group.com/results?basic\_acquirers=Alibaba.com+ [aka Alibaba Group])

Amazon.com Inc. (51)  
 (https://makb.the451group.com/results?basic\_acquirers=Amazon.com+Inc.)

Amazon Web Services Inc. [aka AWS] [Amazon.com Inc.] (7)  
 (https://makb.the451group.com/results?basic\_acquirers=Amazon+Web Services Inc. [aka AWS] [Amazon.com Inc.]

### COMPANY MENTIONS (PRIMARY)

MariaDB Corp (/search?company=MariaDB+Corp)

### COMPANY MENTIONS (OTHER)

Action Corp , Alibaba Group , Amazon , American National Standards Institute , Amazon Web Services , California Technology Ventures , Clustrix , Google , Greenplum , IBM , Intel Capital , MemSQL , Micro Focus International , Microsoft , NuoDB , Open Ocean Capital , Oracle , Pivotal , SAP , SmartFin Capital , Snowflake Computing , Splice Machine , Sun Microsystems , Teradata , Tesi , Vertica Systems , VoltDB (/search?company=VoltDB)

### CHANNELS

Data Platforms & Analytics , Systems & Software Infrastructure (/dashboard?view=channel&channel=3)

### SECTORS

All / Information management / Data management / Data warehousing (/search?sector=117)

Google Inc. (190)  
([https://makb.the451group.com/results?basic\\_acquirers=Google+Inc.](https://makb.the451group.com/results?basic_acquirers=Google+Inc.))

Greenplum [fka Metapa] (1)  
([https://makb.the451group.com/results?basic\\_acquirers=Greenplum+\[fka Metapa\]](https://makb.the451group.com/results?basic_acquirers=Greenplum+[fka+Metapa]))

IBM Corporation (169)  
([https://makb.the451group.com/results?basic\\_acquirers=IBM+Corporation](https://makb.the451group.com/results?basic_acquirers=IBM+Corporation))

Intel Corporation (83)  
([https://makb.the451group.com/results?basic\\_acquirers=Intel+Corporation](https://makb.the451group.com/results?basic_acquirers=Intel+Corporation))

Micro Focus International plc (15)  
([https://makb.the451group.com/results?basic\\_acquirers=Micro+Focus International plc](https://makb.the451group.com/results?basic_acquirers=Micro+Focus+International+plc))

Microsoft Corporation (165)  
([https://makb.the451group.com/results?basic\\_acquirers=Microsoft+Corporation](https://makb.the451group.com/results?basic_acquirers=Microsoft+Corporation))

Oracle Corporation (127)  
([https://makb.the451group.com/results?basic\\_acquirers=Oracle+Corporation](https://makb.the451group.com/results?basic_acquirers=Oracle+Corporation))

SAP SE [fka SAP AG] (61)  
([https://makb.the451group.com/results?basic\\_acquirers=SAP+SE \[fka SAP AG\]](https://makb.the451group.com/results?basic_acquirers=SAP+SE+[fka+SAP+AG]))

SkySQL Ab (1)  
([https://makb.the451group.com/results?basic\\_acquirers=SkySQL+Ab](https://makb.the451group.com/results?basic_acquirers=SkySQL+Ab))

Sun Microsystems, Inc. (23)  
([https://makb.the451group.com/results?basic\\_acquirers=Sun+Microsystems, Inc.](https://makb.the451group.com/results?basic_acquirers=Sun+Microsystems,+Inc.))

Teradata (17)  
([https://makb.the451group.com/results?basic\\_acquirers=Teradata](https://makb.the451group.com/results?basic_acquirers=Teradata))

Figures shown indicate number of transactions