

REPORT REPRINT

MemSQL Cloud gets GA nod, positions for data-warehousing opportunity

JAMES CURTIS, JASON STAMPER

28 JUL 2017

In-memory-database firm MemSQL is adjusting its messaging; the company is now referring to its offering as a 'real-time data warehouse'. This coincides with the general availability of its managed cloud service, MemSQL Cloud.

THIS REPORT, LICENSED TO MEMSQL, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, WAS PUBLISHED AS PART OF OUR SYNDICATED MARKET INSIGHT SUBSCRIPTION SERVICE. IT SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



©2018 451 Research, LLC | WWW.451RESEARCH.COM

After releasing an early version of MemSQL Cloud on AWS in November 2016, the company is now releasing it for general availability. In addition to AWS, it will offer MemSQL Cloud on Microsoft Azure. Moreover, the company is evolving its positioning from an in-memory database to a 'real-time data warehouse.' MemSQL is not abandoning its transactional SQL roots; it is merely emphasizing its dual transactional and analytical capabilities, including streaming, that have been part of the database for some time now.

THE 451 TAKE

With the general availability of MemSQL Cloud, the company is establishing a broad portfolio of products that can be deployed in a variety of environments. No doubt the company's messaging adjustment as a 'real-time data warehouse' will get the attention of the players in that space. But the company is not necessarily trying to transition to a full data warehousing offering. Rather, it is emphasizing its multi-modal strengths in offering both transactional and analytical capabilities. The fact that MemSQL is focusing on live data (as opposed to static data), which can then be immediately available for analytics, will continue to appeal to many organizations, particularly as enterprises update or create new applications to handle both workloads, and should outshine any possible confusion the new messaging may generate.

CONTEXT

Founded in 2011 and headquartered in San Francisco, the company reports roughly 200 paying customers, a sizeable jump from January when we reported customer count at a tad over 100. Employee count is well over 100, consistent with our earlier coverage, although an exact number was not provided. Management reports a revenue growth rate of 100% over last year, and the firm has raised \$85m in funding, the most recent a series C of \$36m in April 2015.

Looking back on our previous coverage of MemSQL, we see that the company has been laying the groundwork to position itself as a 'real-time data warehouse.' For instance, in late 2015, MemSQL added Spark support, effectively enabling the ability to push SparkSQL queries into the MemSQL database. It added enhanced pipelining support in April 2016, giving users the ability to quickly create data pipelines from Apache Kafka data that could then be mapped to the MemSQL database. Then, in November 2016, the company released an early version of its managed database as a service (DBaaS), augmenting its on-premises business, as well as its existing IaaS offerings. We note these updates simply to give some context to the path that MemSQL is taking regarding its new messaging, but we also note the changes because the company is not so much changing its database and roadmap as it is looking to highlight existing functionality.

From a technology perspective, MemSQL is a relational database that runs in memory and on disk, and is also referred to as a multi-modal database in that it can handle both OLTP and OLAP workloads. The company touts its ability to ingest live data, such as from Kafka and other sources, along with its ability to carry out analytics via Spark and other embedded algorithms, and integrate with popular BI tools, such as Tableau, Looker and MicroStrategy. The MemSQL architecture means customers can potentially save on an ETL process between transactional and analytical data stores.

While the new messaging is front and center, it need not overshadow the company's announcement on the general availability of MemSQL Cloud, a managed cloud service offered on AWS and Microsoft Azure. Further, this new managed cloud service should not be confused with the company's existing marketplace offerings available on AWS, Microsoft Azure and Google Cloud Platform.

The GA version of MemSQL Cloud is a follow-on from the company's earlier announcement in late 2016, although it has now been hardened for production use. Management notes that with this release, the company can effectively promote its 'everywhere' strategy in that MemSQL can run on-premises, in the cloud, as a managed service, and as a hybrid or multi-cloud setup. Some notable features include the ability to stream in live data on which

analytics can be run, and the ability to scale out the database to enable high concurrency. MemSQL Cloud supports ‘exactly once’ semantics, which avoids duplicates due to network reordering. It also handles geospatial and JSON data types, and it incorporates a consistent security model across on-premises and cloud.

While the managed service runs on AWS and Microsoft Azure on the back end, the company provides a URL for customers to launch the service, and after a quick registration, an instance can be provisioned with a single click. Once set up, users select data to load and can also create data pipelines as needed. While the service is managed, the cloud interface does provide administrator monitoring for scale out, backup updates, and other activities that may occur.

MemSQL Cloud was pre-released in November 2016 followed by general availability, and no exact customer counts have been provided. However, a few noteworthy customers are worth a mention. One is DirectEmployers Association, a non-profit consortium of 800+ member companies. DirectEmployers helps organizations stay in compliance with recruiting practices, particularly as they pertain to US governmental regulations, such as the Vietnam era Veteran’s Readjustment Assistance Act. The organization developed its own technology leveraging MemSQL where HR personnel can manage job postings across various channels and then conduct real-time analytics to ensure compliance.

Another customer is Thorn , which works to fight the sexual exploitation of children. The organization uses MemSQL Cloud, including machine learning for image recognition analytics. As part of the deployment, Thorn leverages the MemSQL dot product feature, a matching and scoring algorithm that incorporates vector mathematics. The application scans the web for ads and then correlates that with an image database from law enforcement.

Beyond these examples, MemSQL sees its database ideally catering to primarily active data use cases. While static data can be handled (assuming structure added), the company envisions organizations creating new or updating existing applications that can take advantage of the database’s streaming and analytical capabilities, thus the ‘real-time data warehouse’ messaging.

MemSQL will be available first in North America on AWS and Microsoft Azure, with Google Cloud Platform to follow later, along with other worldwide regions. Pricing is based on what the company calls a MemSQL Cloud unit. Customers can purchase any number of units that are then mapped to server instances for either hourly or annual subscriptions. MemSQL Cloud starts at 99 cents per hour per unit, and customers receive a \$300 credit for signing up.

COMPETITION

Despite the company’s messaging adjustment to a ‘real-time data warehouse,’ MemSQL will continue to compete with many of its traditional in-memory and multi-modal competitors. These include VoltDB, Altiibase, Pivotal (GemFire and Apache Geode), GridGain Systems, Hazelcast, Software AG and ScaleOut Software. The traditional relational database vendors – Oracle, IBM and Microsoft – will also continue to compete with MemSQL. SAP HANA likewise pitches transactional and analytics capabilities, and MemSQL notes that it sees HANA in some competitive situations. We also know that SAP HANA customers use the database for data warehousing or analytics purposes, predominantly with the SAP Business Warehouse.

In terms of data warehousing vendors, given the company’s new messaging, we would expect a handful of competitors, particularly on the cloud front. It should be noted, however, that the data-warehousing players clearly pitch an analytics vs. a transactional or operational strategy. Snowflake and Amazon Redshift come to mind as vendors that provide cloud-only offerings. Teradata is another that offers both an on-premises and an array of cloud offerings and promotes a ‘Teradata Everywhere’ strategy. The traditional relational database vendors (Oracle, IBM, Microsoft) all have cloud DBaaS offerings and all manage their cloud platforms. Note that this month, IBM renamed its cloud data warehouse – formerly dashDB for Analytics – to Db2 Warehouse on Cloud.

It’s also worth noting some big-data vendors, although MemSQL points out it doesn’t often see these vendors, but that could change as the market matures. MarkLogic comes to mind because it positions itself as a transactional and operational database and provides some analytical capabilities. MongoDB and Redis Labs are others. Splice Machine and Esgyn are part of a group pitching multi-modal (OLTP and OLAP) in a single platform and are built leveraging Spark, Hadoop and other open source tools.

SWOT ANALYSIS

STRENGTHS

MemSQL brings an in-memory, high-performance strategy with full ACID (atomicity, consistency, isolation, durability) compliance while offering strong streaming capabilities for real-time analytics.

WEAKNESSES

The scaling benefits afforded by cloud infrastructure apply to the MemSQL Cloud offering, whereas the on-premises deployment leverages a scale-up or scale-out architecture. MemSQL Cloud is also not yet available beyond North America.

OPPORTUNITIES

While MemSQL could up pick some customers with its new data-warehousing messaging, the potential larger opportunity will be with the cloud as customers continue to have an interest in analytics, particularly in the cloud where the economics tend to be quite favorable. MemSQL's architecture means that customers can potentially save on a separate ETL process.

THREATS

Moving to or positioning for the data-warehousing market, as well as maintaining a strong presence on the transactional side of the house, has its benefits and drawbacks. Certainly, a larger competitive field will be at play, but it could also send a mixed message to potential customers.