

REPORT REPRINT

DataStax shifts focus to its enterprise edition, steps back from Apache Cassandra

JAMES CURTIS

12 SEP 2017

The company is making a significant pivot in its go-to-market strategy, concentrating primarily on the customer, which means diverting development efforts to the DataStax Enterprise Edition instead of the Cassandra open source community.

THIS REPORT, LICENSED TO DATASTAX, 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.



©2017 451 Research, LLC | WWW.451RESEARCH.COM

Changes are afoot at DataStax. The company that has been the primary commercial entity supporting the open source Apache Cassandra project is now stepping back from that effort and is instead diverting development resources to its own offering – DataStax Enterprise Edition. Inclusive of this change are the company's new corporate objectives that include a customer-first approach, along with simplifying the company's tooling, and a go-to-market focus on customer experience, otherwise known as CX.

THE 451 TAKE

DataStax is a well-known and respected player in the NoSQL space. Certainly, with DataStax's history in contributing to and supporting Apache Cassandra, it no doubt raises eyebrows when the company announces that it is diverting its development resources to its enterprise distribution. And while DataStax will continue to keep Cassandra at its core, the gap between open source Cassandra and DataStax Enterprise Edition will likely widen. Apache Cassandra will be affected to some degree, although not significantly, given the widespread use of Cassandra and that it is relatively mature technology. But other entities and community members will need to step up to continue to support Apache Cassandra. With that as a backdrop, DataStax's timing and strategy pivot makes some sense because it positions the company to stay competitive in what is shaping up to be a highly competitive NoSQL space.

CONTEXT

We have covered DataStax since 2011, shortly after the company's founding in 2010. It started out as the primary supporter and sponsor of Apache Cassandra, which was initially developed by Facebook. As a NoSQL database, Cassandra is categorized as a column-oriented or wide-column store. However, as DataStax has supported Cassandra, the company has progressively added new components and data models to it, such as search, security and management tooling, as well as Spark integration. Our most recent coverage noted the availability of graph in DataStax Enterprise Edition (DSE) based on the company's acquisition of Aurelius. And most recently, the company launched its DataStax Managed Cloud with limited availability in May, with broader availability slated for late summer.

DataStax notes that its recent strategy pivot has been an eight-month effort and has certainly recently affected the company's bottom line. The last quarter, according to management, has been its best quarter ever, picking up several new logos, although no specifics were provided. From our most recent coverage, DataStax reported between 500-600 paying customers. The company now reports more than 450 but points out that the way it counts customers has changed and the fact that the company is phasing out its startup program, where DataStax provided young companies with DSE for free. However, the company does offer its DSE distribution as a free download for development work, although a subscription is required for production deployments.

Employee count has stayed consistent at 450, and DataStax has secured \$190m in funding to date.

STRATEGY

The significant news coming from DataStax is its strategic pivot. Management laid out three core objectives it is working on through to 2020. One objective is the company's future focus on its DSE instead of the open source Cassandra distribution. DataStax notes it is diverting development resources from the open source Cassandra project to DataStax's own customers, and specifically to the company's DSE offering. DataStax will continue to contribute to Cassandra to some extent but not at the level it was before, where it was contributing over 85% of the updates to Cassandra. Any new enterprise additions DataStax makes, however, will remain with DSE.

Another objective is a focus on simplification, such as updating tools and their integration with DSE as well as the ability to use DSE. Currently, DataStax offers DSE to be run in a customer's own datacenter that it then manages. DataStax is also able to augment customer deployments with DataStax resources, which can also include cloud deployments. And with DataStax Managed Cloud, the company provides a fully managed service for customers that do not have the resources or would prefer to outsource the management of the infrastructure to another entity.

Last, DataStax is pursuing a CX strategy and is beefing up its sales and marketing efforts to align with this. CX has become a catchphrase for many NoSQL vendors to 'see the world through the eyes of one's customers.' And this makes sense, given the architectural properties many NoSQL data models provide in terms of low latency, scalability and flexible data schemas, particularly when applied to web- and mobile-centric applications. For DataStax, its CX approach includes focusing more on a business-centered audience instead of a developer audience. Moreover, the company is driving a number of initiatives to support field teams, such as developing qualification criteria, use cases and presentations, as well as setting up war rooms for collaboration.

PRODUCTS

In April, DataStax released DSE 5.1. A few noteworthy updates are the addition of row-level access control along with improved replication to multiple destinations. For improved analytics, Spark 2.0 is certified and search functionality, based on Solr, has been added to CQL (Cassandra Query Language).

Since the release of DataStax Graph in fall 2016, the company notes both growth and technical improvements, as well as alignment with the company's CX strategy. DataStax reports 36 customers using graph over 13 industry verticals. DSE Search has been added to DSE Graph, integrating more deeply with Spark, and the bulk load has been updated.

In terms of graph, DataStax sees an ideal match with CX. Early on, the company was driving awareness and general adoption of graph. But now DataStax is looking to drive more use case scenarios in which graph plays a critical role but is not necessarily top of mind from a customer perspective. For instance, DataStax might advance customer 360, fraud detection and recommendation use cases, but the customer may not understand that graph is a key component.

COMPETITION

Even though DataStax has had a history of supporting Apache Cassandra, open source Cassandra is still considered a competitor, and even more so now that DataStax is diverting its attention to its enterprise edition. Customers would, of course, need to build on the enterprise features similar to what DataStax has added.

However, there are a number of NoSQL vendors worth noting as well. MongoDB is often noted as a primary competitor to DataStax and offers multiple data models, including graph. The company also continues to drive its Altus cloud service. Couchbase is another, also multi-model, and the company has been quite active in positioning its database as an 'engagement database,' that likewise targets web and mobile scenarios with a CX focus. MarkLogic positions itself as an operational and transactional database and touts its semantics (graph) capabilities as complementary to its document and XML data models. Redis Labs, known for its in-memory capabilities based on a key-value model, drives a flexible platform offering with its Modules strategy, where graph is available as a module. Aerospike similarly positions itself as a system of engagement based on a key-value data model and is differentiated with its SSD, flash-based architecture.

A few other NoSQL vendors offer multiple data models and are potential competitors to DataStax. These include ArangoDB, MapR-DB, OrientDB, and InterSystems Cache, a non-relational database. And then there is ScyllaDB, which has been quite active in looking to displace Apache Cassandra.

Also worth a mention is Instaclustr, which offers a managed cloud service for Apache Cassandra and has recently increased its commitment to the open source project. Additionally, on the cloud front, there is Amazon Web Services' DynamoDB, Microsoft Azure Cosmos DB and IBM Cloudant.

SWOT ANALYSIS

STRENGTHS

DataStax has taken Apache Cassandra and added a number of new data models - graph, for instance - as well as a plethora of features on top to make it enterprise-grade.

WEAKNESSES

DataStax has kept a somewhat lower profile while it was working out its new strategy and perhaps that led to some confusion in terms of the company's future engagement to open source Apache Cassandra community.

OPPORTUNITIES

With the focus now on a more business-centric audience as opposed to a strict developer audience, DataStax should be able to initiate new discussions that then lead to driving a fruitful 'land and expand' strategy to existing customers.

THREATS

There are several NoSQL vendors that, to some degree, are positioning to so-called modern applications with a mobile and web focus, which can lead to some confusion in the market regarding differentiation between vendors.