

SingleStore targets data-intensive applications with unified database for analytical and transactional workloads

Analysts - Matt Aslett

Publication date: Monday, August 16 2021

Introduction

We noted earlier this year that operational databases need to evolve to support emerging operational application functionality, including automation and proactive intelligence, such as recommendations driven by machine learning models. SingleStore refers to this new class of applications as being 'data-intensive' – a phrase that it is increasingly emphasizing as it looks to highlight the use cases that are the best fit for its database, which is available both as a managed service and for self-managed deployment. With its latest update – version 7.5 – the company announced support for the separation of compute and storage, as well as the default use of its Universal Storage architecture, which is fundamental to its ability to support what we call hybrid operational and analytic processing.

The 451 Take

Although the company was founded a decade ago, it could be argued that it is still relatively early stages for adoption of SingleStore. Rebranding both the company and its storage architecture in 2020 was effectively a reset, and the announcement that Universal Storage is now considered functionally complete (at least from a user-experience perspective) and is now the default table type in both SingleStore DB and the SingleStore Managed Service is a key milestone. Wider industry trends, such as the drive toward intelligent applications, are playing to the company's strengths, and the focus on data-intensive applications should help the company highlight its strengths in supporting a new class of applications that require both transactional consistency and high-performance analytics.

Details

SingleStore was founded in 2011 as MemSQL with a plan to create a new high-performance memory-optimized operational database to serve the next generation of online businesses. While the company's database was always designed to support transactional operational applications, it quickly became clear that the next generation of operational applications involved the delivery of functionality driven by real-time analytic processing.

Having initially added a columnar storage engine alongside its original row-store engine, the company set to work on a storage architecture to support both OLTP and OLAP workloads. Initially known as SingleStore and introduced in version 7.0, that storage architecture is now known as Universal Storage, and with the recent launch of version 7.5, is now both functionally complete and the default table type. The SingleStore brand, meanwhile, was adopted by the company as its official name in October 2020.

The company has continued to grow since our [previous update](#) and now boasts 250 employees (up from more than 200), 150-200 customers (compared to 128), and roughly \$50m in annual recurring revenue (up from approximately \$40m). About 40% of that revenue is now generated by its SingleStore Managed Service (previously known as Helios, launched in 2019), which is available via Amazon Web Services, Google Cloud and Microsoft Azure.

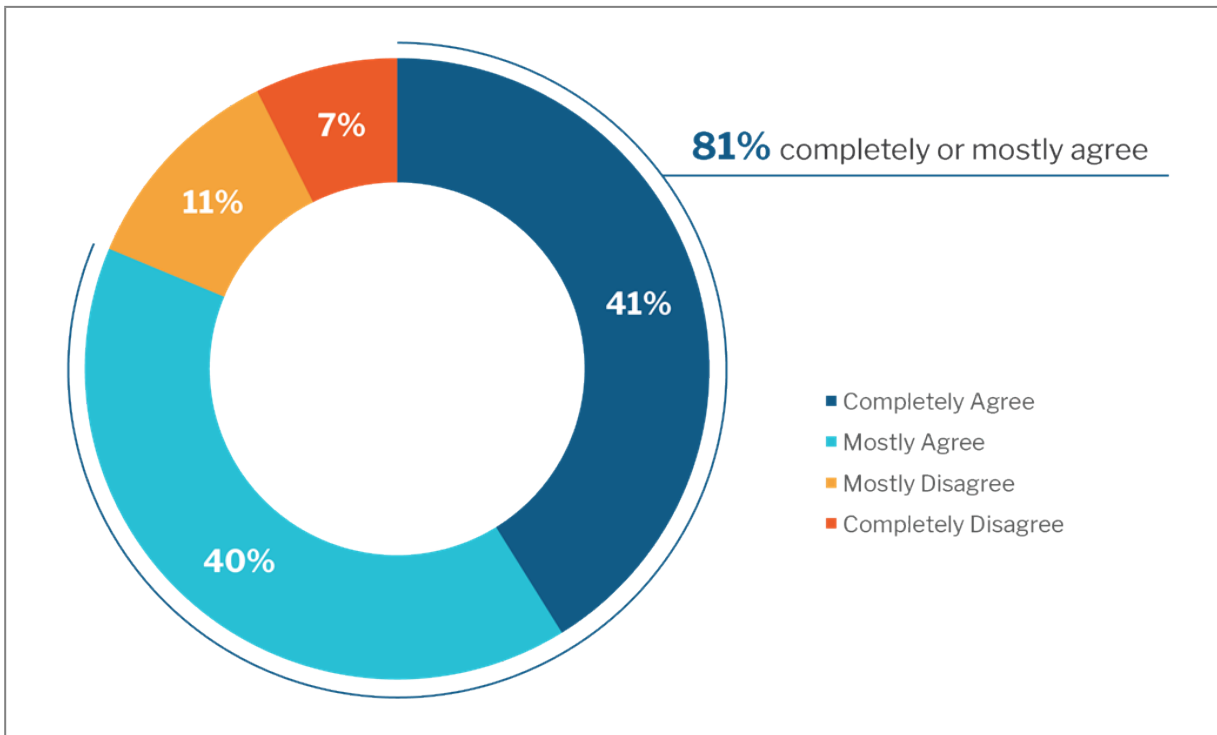
In addition to being included in version 7.5 of SingleStore DB, the functionally complete Universal Storage is also available via SingleStore Managed Service, along with other 7.5 enhancements such as the separation of compute and storage, point-in-time recovery, the ability to suspend and resume resources as required for transient workloads, and a new flexible pricing model.

As noted above, the Universal Storage architecture has been a work in progress for SingleStore for several years, but it is now considered functionally complete from a user-experience perspective. It is now the default table type in the database, supporting both transactional and analytic workloads by mixing the use of memory for transactional workloads and compressed columnar disk storage for historical data and analysis.

While support for operational and analytic workloads is what many people would think of when the company refers to its technology as a unified database, the reference also pertains to the database's support for multiple data models, including relational, document, key-value, time-series, geospatial, full-text and streaming.

This is functionality that is increasingly expected by database purchasers – 81% of respondents to 451 Research's Voice of the Enterprise: Data & Analytics, Data Platforms 2020 agreed that their organization has a preference for databases that can simultaneously support multiple data models rather than using multiple dedicated databases.

Preference for Databases that Support Multiple Data Models



Source: 451 Research's Voice of the Enterprise: Data & Analytics, Data Platforms 2020

With the latest update, SingleStore has also introduced support for the separation of compute and storage – enabling each layer to be scaled separately. The company notes that it has approached this slightly differently than other providers: ingesting data into the SingleStore cluster and then into object storage enables immediate analysis of the data (rather than ingesting directly into object storage and then accessing the data).

Also new is point-in-time recovery of the database across on-premises, cloud and hybrid environments, as well as the ability to suspend and resume compute resources as required in order to avoid unnecessary costs when workloads are inactive. The latter is relevant in relation to the company's new pricing model, which enables customers to use the product on demand or commit to an annual subscription based on estimated monthly usage (with bursting above the monthly rate charge using on-demand pricing, and unused credits being rolled over).