

www.SingleStore.com

388 Market Street, Suite 860, San Francisco, CA 94111

Tel: +1 855 463 6775

Email: team@singlestore.com


“Once we found SingleStore, there really wasn't a reason to look at anything else. We said, 'Okay, this is awesome. We need this!'”
Siemens

SingleStoreDB

The company

SingleStore was founded by former members of Facebook and Microsoft as MemSQL in 2011. Its product was first released in 2013 under the MemSQL brand. In October 2020, the company rebranded itself to SingleStore and its product to SingleStoreDB to better reflect what it is offering: a single store (database) for all analytical and transactional data.

The company has raised \$464m in venture capital and is headquartered in San Francisco, with other US offices in Sunnyvale, CA, Raleigh, NC, and Seattle, WA, as well as overseas offices in Lisbon, Hyderabad, Dublin, and Singapore, with remote workers in many other countries across the world. It combines a direct sales model with a robust partner programme that includes the big 3 clouds (AWS, Google, and Azure), hyperscalers such as Dell, HPE, and IBM, deployment partners such as Arcion, Datadog, Data-Sleek, and Twingo, and many others. The company particularly focuses on migrations from legacy databases (MySQL, PostgreSQL, Oracle, Microsoft SQL Server, Teradata, and so on) and replacement or augmentation of cloud databases (such as Snowflake, BigQuery, Databricks, and MongoDB), with which it has had significant success. Migration support is currently a service engagement.



The image in this Mutable Quadrant is derived from 13 high level metrics, the more the image covers a section the better. Execution metrics relate to the company, Technology to the product, Creativity to both technical and business innovation and Scale covers the potential business and market impact.

available for download, with the only restriction being the amount of in-memory capacity that is supported.

What is it?

SingleStoreDB is a scale-out distributed database, with a lock-free architecture that supports both row and column storage. Unlike traditional relational offerings, row-based data is processed in-memory for transaction and operational processing, while columnar data (which is compressed) is persisted to disk to support analytics. It uses SQL (which is compiled into machine code) and its SQL support is wire-compatible with MySQL. The product is ACID-compliant. Immediate consistency is supported through synchronous replication, but since this can cause a performance overhead, it is not the default setting. However, synchronous replication is enabled by default for reference databases, ensuring durability against single-node failures.

The product is envisioned as a singular solution to many, or even all, of your data storage needs. This is supported by its multi-model nature: it is capable of handling many different types of data, removing the need to build and manage a complex, expensive, and often fragile architecture of more specialist databases. The product's support for vector data is particularly relevant in light of recent interest in (generative) AI, and important vector-related features include joining vector data with other data, running vector semantic searches and/or keyword matching, and performing hybrid searches.

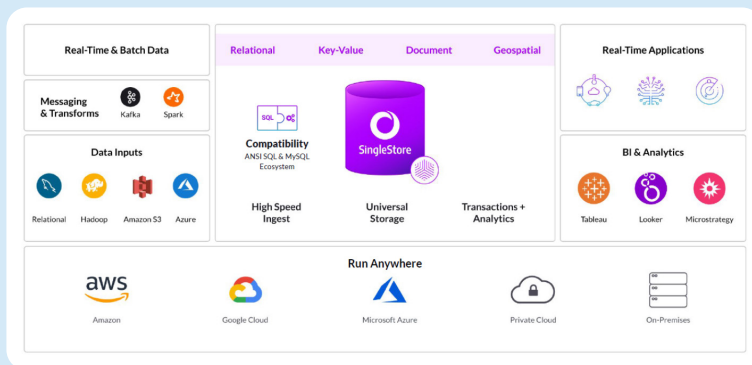


Figure 1 – SingleStoreDB ecosystem overview

SingleStore's customers include fast growing startups/scale-ups as well as some of the world's largest enterprises. Its major clients are comprised of 100+ Fortune 500, Forbes Global 2000, and Inc. 5000 brands. Deployment options include SingleStoreDB Self-Managed (on-premises/containers/Kubernetes) and a fully-managed cloud service, SingleStoreDB Cloud. In addition, a fully-featured version of the product is freely

“With SingleStore, we can build AI and machine learning-powered self-service apps that deliver super-low-latency queries, ultra-fast ingestion, and high concurrency. We are processing millions of real-time queries supporting tens of thousands of concurrent users.”
GE

What does it do?

In addition to the tabular and vector formats already discussed, SingleStoreDB supports the JSON format, geospatial data types and operations, full text search, and time series. In the geospatial case, the product has data types for points, curves, and polygons, along with geographic indexing. Geographic features understand the curvature of the Earth but planar capabilities, for example mapping a factory floor, are not provided. On the time-series front, the product’s architecture is well-suited to the efficient storage of time-series data, and includes features like time buckets and interpolation.

Moreover, SingleStoreDB accommodates both transactional and (real time) analytical workloads, effectively acting as either an operational database or a data warehouse depending on the task at hand. Analytics in particular are done with minimal, if any, data movement. What’s more, a (patented) single table type supports both transactional and analytical operations, allowing for universal data storage while enabling both types of workload. These operations are performed at speed using the product’s three-tier storage model, shown in *Figure 2*.

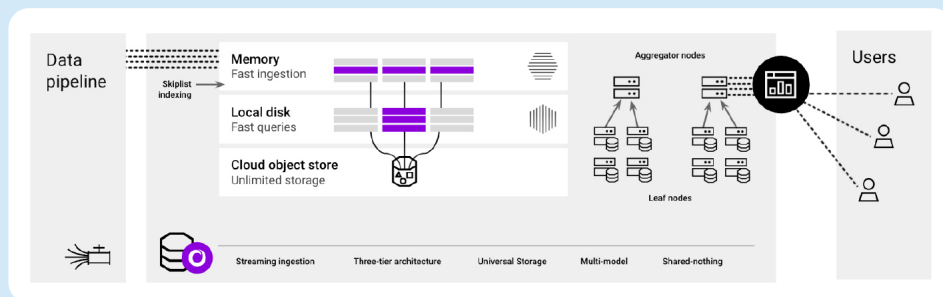


Figure 2 - SingleStoreDB architecture

The product uses a distributed cost-based query optimiser that is aware of how data is stored to direct query operations to in-memory or disk-based storage as required. Note that SingleStoreDB uses skip-lists as its primary indexing technique, as well as supporting secondary indexes. Separation of storage and compute is another relevant scalability feature.

SingleStoreDB’s Workload Manager allows both memory and CPU resources to be allocated by user group, while its distributed ingest supports both batch and stream-based ingestion. Multi-cluster management is provided through SingleStore Studio, along with facilities such as encryption, security, role-based

access control and auditing. There is also a SQL API that serves to support third-party environments such as Tableau and Qlik, and a specialised API – SingleStore Kai – designed to enable fast analytics for MongoDB applications. Finally, SingleStore has recently released SQRl, a chatbot tool that lets you sift through technical documentation or receive code samples via generative AI queries.

Why should you care?

SingleStoreDB is a modern, distributed relational database and brings all the advantages you would expect from the familiarity of SQL. Wire compatibility with MySQL is a major advantage for anybody looking to migrate from an incumbent relational supplier, and the product reverses the use of memory and disk storage with respect to rows and columns, which means that you can get better performance for transaction processing while using fewer hardware resources and maintaining low latency for analytics and query processing. It also has strong geo-spatial and time-series capabilities.

What’s more, replacing a complex system with a simpler and more comprehensible one is often a good idea, and this is what SingleStoreDB allows you to do by subsuming the role of many more specialised offerings. Its support for vector data is particularly notable, as vectors have grown in popularity recently due to their perceived synergy with generative AI engines and the LLMs (Large Language Models) that power them.

Moreover, vectors are useful to LLMs because they can be used to supply important context. It stands to reason that this context is at its fullest, and therefore most effective, when created from a combination of vector and other data. A unified, multi-model database that includes vector data, such as SingleStoreDB, seems the easiest and most direct route to achieving this.

The Bottom Line

In our view, SingleStoreDB starts from the right place without being encumbered by historical design decisions, and the company has proven that it is a serious competitor to the more established relational database vendors. Its proposition of replacing many specialised databases with one, as well as its support for vectors, should both be highly appealing in the current data landscape.

“Our Glovia legacy application on Oracle provided batch 30-minute reports. After implementing SingleStore, analytic reports load within a minute, giving our users on-demand reporting in near real time.”
Dell Technologies

FOR FURTHER INFORMATION AND RESEARCH [CLICK HERE](#)