

SingleStoreDB Cloud

Fully managed cloud database service

Overview of SingleStoreDB

SingleStoreDB is a real-time distributed SQL database, designed for **modern, real-time, data-intensive applications**. It delivers extreme performance for both transactional applications (OLTP) and analytical workloads (OLAP).

With support for **streaming data ingestion** (Pipelines), a unique table type that supports **both transactions and analytics** (Universal Storage), **separation of compute and storage** (Unlimited Storage), limitless **point-in time recovery** (PITR) and a **distributed** (shared-nothing), **MySQL-compatible** architecture, SingleStoreDB provides fast-growing companies the ability to build and scale modern, real-time applications and analytics.

All this comes with **compelling price-performance**. SingleStore customers often report **2-3x better TPC-H performance** at a **fraction of the cost** of multiple technologies for OLTP, OLAP, ETL, NoSQL, etc.

SingleStoreDB Cloud

SingleStoreDB Cloud is a fully managed database service available on AWS, Azure, GCP, or in the customer's data center with the *Dedicated* edition. SingleStoreDB Cloud removes the need for users to manage infrastructure, provision clusters, handle upgrades or troubleshoot failures. This greatly reduces operational expenses, while accelerating organizational agility and allowing enterprises to focus on customers — rather than on managing their data infrastructure.

SingleStoreDB is architected from the ground up as a multi-model database, supporting relational, key-value, document, geospatial, time-series data, and full-text search to simplify modern data architectures while reducing overall complexity and costs.

SingleStoreDB – Product Editions

SingleStore provides multiple offerings to meet the needs of any workload:

Standard

For most workloads.

Features:

- Standard support
- 99.9% uptime

Premium

For mission- critical workloads.

Features:

- Standard support
- 99.99% uptime

Dedicated

For stringent Governance, Risk & Compliance needs

Features:

- **Silver** support
- 99.99% uptime

SingleStoreDB Cloud may also be trialed with \$600 in [free cloud credits](#)

Benefits



SPEED

Drive **real-time analytics** with streaming data ingestion and low latency query performance



SCALE

Scale-out architecture for the most demanding applications and analytics



RUN ANYWHERE

On-premises, public, private, or hybrid clouds

Analytics performance: SingleStoreDB drives low-latency (low double digit ms) query responses with high concurrency across both live and historical data using familiar ANSI SQL

Streaming data pipelines: Deliver real-time analytics with parallel, distributed lock-free ingestion and real-time query processing

Scale-out architecture: Elastic scale-out architecture with distributed massively parallel data processing delivers consistent, predictable response under high ingest and user concurrency

HTAP/ Translytical workloads: Supercharge the development of modern, data-intensive applications with a unified table type for transactions and analytics

Drop-in compatibility: MySQL wire-protocol-compatibility, enables easy plug-in directly with existing technologies and skills with support for standard SQL, commonly used BI tools and technologies like Amazon S3, Spark, Kafka and Hadoop

Why choose Cloud

Cloud-only features. As a fully managed cloud service, SingleStoreDB Cloud offers unique features including Workspaces, Management API, usage based billing, continuous updates and service assurance.

Easy to operate. The Cloud offering helps customers accelerate development with its simple management, integrated tools, and autonomous operation at scale.

Enhanced security: SingleStoreDB Cloud secures customer data using defense in depth with three security layers: Data Security, Access Control and Network Isolation. In addition, SingleStoreDB adopts strong operational security best practices and helps customers keep up with evolving data compliance and certification needs.

Features

Workspaces is a Cloud-only feature that enables independent deployment and scaling of compute resources. Each workspace can attach and detach databases as needed, allowing data to scale across workloads without the complexity and cost of managing data across different data silos. Users can create and terminate workspaces directly using the cloud portal, or through our scalable **Management API**.

SingleStoreDB Pipelines: Built-in parallel data ingestion technology natively ingests high-throughput, real-time data from external sources like Apache Kafka, Amazon S3, Azure Blob, Filesystem, Google Cloud Storage and HDFS data sources.

Universal Storage: A unique table type that supports both OLTP and OLAP in the same engine. Universal Storage gives you the best of row stores and column stores in one table type thereby reducing the latency and costs of duplicating and moving data.

Unlimited Storage (Separation of Storage and Compute): Allows users to effortlessly scale compute resources to meet the needs of any workload while managing the storage needs completely independently. Workspaces lets you isolate workloads and power multiple applications on the same database:

MySQL Compatibility: SingleStoreDB is wire-protocol compatible with MySQL/MariaDB — making it instantly accessible from any BI tool like Tableau, PowerBI or Looker. It also includes widely available bindings to popular programming languages including Golang, Rust, Python, NodeJS, R, Java, .NET and C++.

Distributed Ingest (Bulk or Streaming) with Concurrent Non-Blocking Reads: SingleStoreDB offers a lock-free architecture based on the skip list index that efficiently processes transactions and updates without locking or blocking concurrent reads. This results in the capability to perform bulk and/or streaming ingestion online, simultaneously with query workload.

Limitless Point-In-Time Recovery (PITR): With PITR users can restore the database to any specified date and time, down to the microsecond. PITR enables customers to power mission-critical applications, and provides peace of mind that application or user errors can be quickly and easily remediated.

Cloud observability and monitoring: Built in real-time and historical monitoring using Grafana dashboards to track key resources for SingleStoreDB Cloud deployments.

SingleStoreDB Replicate: Offers a robust data ingestion tool for replicating data from an external third-party database into SingleStoreDB, and provides a number of features and functions — including both filters and maps — to tailor how the data is replicated.

Compiled, Vectorized Query Execution: Built-in distributed query optimizer evenly divides the processing workload to maximize the efficiency of CPU usage. Query plans are compiled to machine code and cached to expedite subsequent executions. Vectorized operations can reduce per-row CPU overhead by more than **100x**.

Enterprise Security: Ensures military-grade security with RBAC, encryption, auditing, password policy management and strict mode to isolate data from administrators. OAuth (Cloud) to simplify federated authentication.

Programmability: SingleStoreDB offers **extensibility** for in-database programming via its MSQL language, allowing users to write user-defined functions (UDFs), stored procedures (SPs), table-valued functions (TVFs) and user-defined aggregates. SingleStoreDB's **Code Engine, Powered by Wasm** supports creating functions (UDFs and TVFs) using C/C++/Rust code compiled to WebAssembly.

