

# SingleStoreDB On-Premises

The fastest real-time distributed SQL database built for your data-intensive applications

## Overview

SingleStoreDB is a fast, distributed, highly-scalable SQL database designed to power today's data-intensive applications. It is designed to deliver maximum performance for both **transactional (OLTP)** and **analytical (OLAP)** workloads in a single engine to drive blazing fast speeds and interactivity for your modern applications. With SingleStoreDB, you can ingest millions of events per second with ACID transactions using **Pipelines**, while simultaneously delivering low-latency SQL queries on billions of rows of data.

With our patented **Universal Storage**, which combines the best of rowstore and columnstore, SingleStore offers an innovative breakthrough in database architecture, allowing both transactional and analytical workloads to be processed using a single table type. This, together with separation of storage and compute, Point-In-Time Recovery (PITR) and workspaces, deliver on our vision of being the modern database to power data-intensive applications and real-time analytics.

SingleStoreDB is also designed from the ground up as a **multi-model** database, supporting relational, key-value, JSON, geospatial, time series, and full-text search. With a relational engine at its core, SingleStoreDB delivers effortless scalability and simplicity to modern data architectures, reducing the overall complexity and cost.

And when it comes to deployment options, SingleStore offers both **multi-cloud** and **hybrid options**, ranging from a database-as-a-service, to Kubernetes based hybrid and private deployments, to self-managed installations on VMs or commodity hardware, in the public cloud or **on-premises**. SingleStoreDB provides a resilient, blazing-fast database with cloud-agnostic deployment support on AWS, Azure, and Google Cloud Platforms.

With **10 - 100x** performance at **one-third the cost** of traditional databases, SingleStoreDB is ideal for applications that require fast data ingestion, low-latency queries and elastic scaling with familiar relational SQL.

## SingleStore DB - Product Editions

SingleStoreDB provides tailored offerings to meet the needs of any workload.

### Free

Perpetually Free deployment for developers

#### Features:

- No support

### Standard

Designed for modern data-intensive workloads.

#### Features:

- Standard support

### Premium

Designed for mission-critical workloads.

#### Features:

- Standard support
- PITR and more...

## Benefits

SingleStoreDB enables enterprises and startups to build and scale modern applications enabling faster, more informed decisions, improving customer experience, and benefiting from more cost-efficient operations.



### SPEED

Drive real-time insights with **ultra fast ingest and high performance queries**



### SCALE

**Designed to scale** for today's most demanding hyper-scale applications and analytical systems



### AGILITY

Cloud-native data platform designed for any data, to **run anywhere**: on premises, public, private, or hybrid/ multi-clouds

**Latency-Free Analytics:** SingleStoreDB drives low-latency query responses with high concurrency across both live and historical data using familiar ANSI SQL

**Ultra-fast Event-to-Insight Performance:** Deliver against the toughest service-level agreements using parallel, distributed lock-free ingestion and real-time query processing

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

**Ease of Use and Flexibility:** Powers modern applications and brings simplicity and ease to your data architecture by allowing OLTP and OLAP workloads to be processed using a single table type

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

## Technical Specifications

**Universal requirements:** Each SingleStoreDB On-Premises node requires a host machine with an x86\_64 CPU with at least four CPU cores and 8GB of RAM available per node. When provisioning your host machines, the minimum Linux kernel version required is 3.10 or later.

**Recommended platforms:** RHEL/CentOS 6, 7 or 8, Debian 8 or 9 (version 9 is preferred)

**Hardware recommendation for optimal performance:** CPU: 8 vCPU per host machine, Memory: At least 4GB per core and 32GB minimum per leaf node, Storage: Provide a storage capacity for each node with at least three times the capacity of main memory, and SSD storage is recommended for columnstore workloads.

**Network Hardware Recommendation:** 10 Gigabit ethernet switch or better recommended for cluster node interconnection.

---

## Features

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

**Universal Storage:** A patented new table type that supports both OLTP and OLAP in the same engine. Universal Storage gives you the best qualities of row stores and column stores together while reducing data duplication, data movement, and data latency.

**MySQL Compatibility:** SingleStoreDB is wire-protocol compatible with MySQL/MariaDB making it instantly accessible from any BI tool such as Tableau, PowerBI or Looker and with widely-available bindings to popular programming languages such as Golang, Rust, Python, NodeJS, R, Java, 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, resulting in delivering the capability to perform bulk and/or streaming ingestion online, simultaneously with query workload.

**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**.

**Separation of Storage and Compute:** Offers unlimited storage and 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.

**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.

**SingleStoreDB Replicate:** offers a robust data ingestion tool for replicating data from an external 3rd-party database into SingleStoreDB, and provides a number of features and functions, including both filters and maps, to tailor how the data is replicated.

**Multi-Version Concurrency Control (MVCC) and Lock-Free Data Structures:** With these technologies, readers and writers never block each other, even amidst a high volume of concurrent reads and writes.

**Enterprise Security:** Ensures military-grade security with RBAC, encryption, auditing, password policy management, and strict mode to isolate data from administrators.

**Extensible:** Supports in-database programming via the MPSQL language, which allows the definition of user-defined functions, stored procedures (SPs), table-valued functions, and user-defined aggregates.

## Use Cases

**Powering Data-Intensive Applications:** Customers use SingleStoreDB to build and scale their customer-facing applications, powering some of the most mission-critical applications and use cases, including real-time customer analytics, wealth management, risk analytics, cyber security, website analytics, ad analytics, IoT platforms, and more.

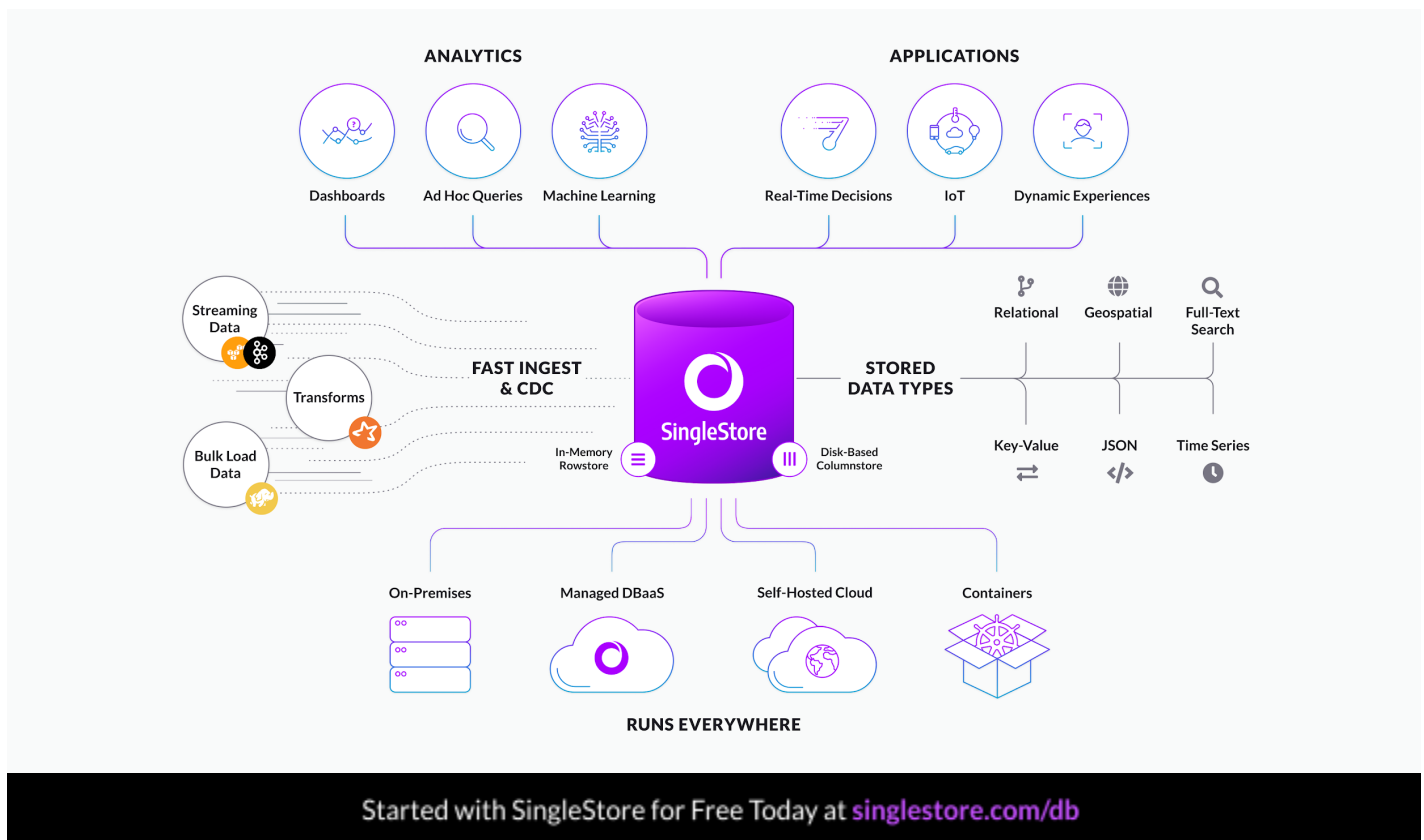
**Driving Fast Analytics:** SingleStoreDB delivers the fastest and most scalable reporting and analytics across all of your operational data; including streaming, real-time, and historical data. Typical scenarios include real-time applications for retail inventory analytics, portfolio analytics, gaming analysis, threat detection, and streaming media quality analytics.

**Operationalizing ML and AI:** SingleStoreDB can power your ML/AI applications with an ultra-fast ingest and query platform that enables real-time model scoring on both streaming and historical data, and vector similarity testing for image matching and other deep learning applications. SingleStore also powers monitoring and anomaly detection in real-time on rapidly changing data through its innovative real-time ingestion, database, and analytics platform. Typical use cases include real-time fraud detection, facial recognition, predictive analytics for IoT etc.

**Accelerating Legacy Data Platforms:** SingleStoreDB can accelerate legacy systems with scalable, rapid data ingestion and fast queries on large data sets, by leveraging the simplicity of SQL. Typical scenarios include modernizing data lakes and data warehouses built on Hadoop, object stores & legacy analytic databases by bringing low-latency queries to the data layer.



## SingleStore Architecture



[Singlestore.com/on-premises](https://singlestore.com/on-premises)