

Data Warehouse Augmentation with SingleStore

Delivering Real-Time Insights & Lower TCO with SingleStore

Today's data warehouses power business intelligence (BI) and reporting workloads that enable organizations to quickly aggregate and analyze large amounts of data from multiple sources to drive insights. However, data warehouses are not optimized for low-latency analytics — especially when you are dealing with fast-moving streaming data from diverse sources to power your modern applications, or drive insights in real time.

As modern businesses demand speed, scale and agility, traditional data warehouses are unable to handle high throughput streaming ingestion, super low-latency analytics and concurrency needs that enterprises demand. A new approach to powering analytical data applications requires reducing time-to-insights for fast analytics on dynamic data for complex queries, all within sub-seconds.

Common data warehouse solutions provide too many challenges, including poor performance and complexity.

Key Challenges with Data Warehouses

Streaming Ingestion

Inability to ingest, process and analyze fast-moving streaming data necessary to power real-time analytics

Batch Processing

ETL batch windows (anywhere from 2-24 hours) create delays in getting analytics to users and applications

Low-Latency Query Performance

Designed for pre-aggregated data sets, and not fast query performance or ad-hoc analytics

High Concurrency

Not optimized to handle large numbers of concurrent users or queries, unable to easily scale data applications or analytical dashboards

SingleStoreDB surpasses the limits of data warehouses to drive up to **20-100x faster time-to-insights**, powering low-latency analytics and applications. SingleStoreDB offers a distributed cloud-native SQL data platform with ANSI SQL compatibility that allows businesses to achieve ultra-fast query responses with high concurrency on live and historical data. SingleStoreDB supports analytics on streaming data by ingesting millions of rows per second on data-at-rest and data-in-motion. SingleStoreDB's unique architecture is designed to power data-intensive applications because of its unique ability to support **both transactional and analytic workloads** — all while enabling real-time analytics.

Why SingleStoreDB

20-100x Faster Time-to-Insights

High-throughput parallel streaming ingestion together with super low-latency queries to drive faster time-to-insights

Run Apps & Analytics with Efficiency

Offers the ability to run fast analytics and data apps directly on SingleStoreDB while lowering costs and time-to-insights

Up to 50% Lower Costs

Accelerate your data warehouse to deliver better price-performance and overall lower TCO of up to 50% with SingleStore

“SingleStore can process complex queries with large data sets within 1-3ms. The closest Snowflake or BigQuery can get us is in the 200ms range.” — Leading Enterprise Customer

SingleStoreDB for Augmenting Data Warehouses

SingleStoreDB is built from the ground up as a distributed, highly scalable, unified database that delivers maximum performance for both transactional and analytical workloads. It unifies transactional and analytical processing on diverse data (semi-structured, and structured) in a single engine – with the ability to use standard SQL to join these diverse native data types. With **20x to 100x** the performance at one-third the cost of legacy infrastructures, SingleStore delivers unmatched speed, scale and agility in a powerful, cloud-native relational database.

Data Warehouse - Key Challenges

1) **Data Ingestion:** No real-time ingestion, long ETL batch windows & no low-latency upserts



2) **ETL and Batch Processing:** Data is aggregated, transformed and loaded with long batch windows



3) **Low-Latency Analytics:** Designed for pre-aggregated data sets & not optimized for low-latency queries



4) **Price-Performance:** Very costly to scale for performance, especially for workloads that need higher concurrency



Fast Analytics with SingleStore

SingleStore Pipelines: Optimized for fast, real-time ingest – up to millions of events/sec, or batch uploads from external sources including Apache Kafka, Amazon S3, Microsoft Azure Blobs, GCP and HDFS

Real-time transformation and processing with connectors to **Spark** with parallel, distributed lock-free ingestion and **real-time query processing** to drive ultra-fast event-to-insight performance

SingleStoreDB delivers the **world's fastest SQL engine** for both transactions and analytics. Up to **20-100x faster time to insights** and supports multiple data types (JSON, time-series, geo, full-text search, relational)

SingleStore is built to power **both applications and analytics** in the same engine without any data movement or ETL. With **80-90% compression** in the Columnstore together with efficient query compilation and commodity hardware – SingleStore delivers overall **lower TCO up to 50%** compared to pure data warehouse architectures.

Using SingleStoreDB to Augment your Data Warehouse

SingleStoreDB can effectively accelerate or **augment your data warehouse** infrastructure to enable ultra-fast ingestion, unlimited storage and processing with blazing fast queries to drive fast analytics on any data, anywhere.

Customers can deploy SingleStoreDB in any of the leading **cloud environments**.

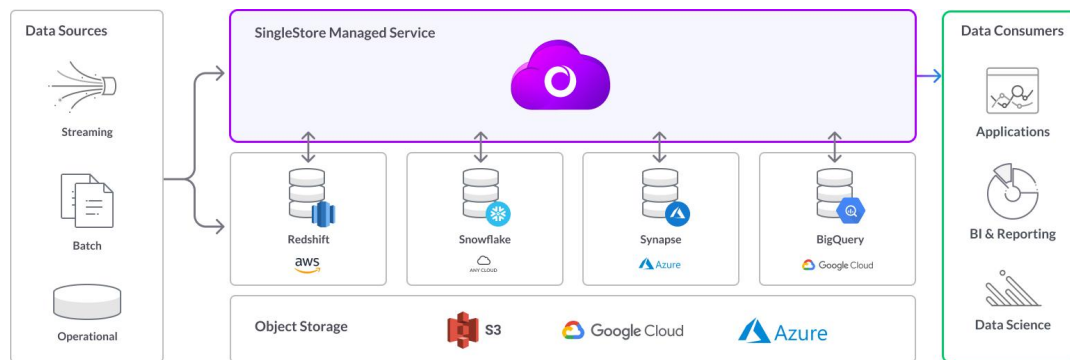


Figure 1: Data Warehouse Augmentation with SingleStore

Key Outcomes with SingleStoreDB

- ✓ Fast analytics on dynamic data for millions of real-time complex queries across 1000s of users
- ✓ High-scale streaming ingest at millions of events/ sec
- ✓ Blazing fast queries with sub-second latencies
- ✓ Power real-time applications and fast analytics within the same engine



SingleStoreDB - Key Features

Patented Universal Storage: Both large-scale OLTP and OLAP are supported on this single, default table type. Universal Storage gives you the best qualities of row stores and column stores while reducing data duplication, data movement and data latency.

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 source.

MySQL Compatibility: SingleStoreDB is wire-protocol compatible with MySQL/MariaDB which offers access to hundreds of languages, 100% compatibility on data types and 95% coverage of built-in functions to ease migrations.

Security & Compliance: Delivers enterprise-grade security with integrated user authentication, full encryption of data in transit and at rest and SOC2, ISO27001, HIPAA, GDPR and CCPA compliance.

Separation of Storage and Compute for Transactions and Analytics: Allows users to effortlessly scale compute resources to meet the needs of any workload, while managing storage needs completely independently.

Distributed Ingest, Bulk or Streaming, with Lock-free/ Non-Blocking Reads and Concurrency: Offers a lock-free architecture 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.

Suspend & Resume Workloads Effortlessly: Clusters can be suspended and resumed nearly instantaneously, making all of your data available when you need it, and minimizing costs when workloads are inactive.

Flexible Credit Pricing Model: Provides flexibility of on-demand, or with monthly credit bundles to handle dynamic and growing compute workloads at reduced TCO.

Latency-Free Analytics: SingleStoreDB lets you achieve ultra-fast query response 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 responses under high ingest and user concurrency.

Ease of Use and Flexibility: SingleStoreDB brings simplicity and ease to your data processing by allowing OLTP and OLAP workloads to be processed using a single table type.

Tiered Storage: Three-tiered storage including in-memory, SSD Cache and the Cloud object store with separation of storage and compute (unlimited storage) and 80-90% data compression

Customer Case Studies



Leading Electronics Manufacturer: SingleStore enables real-time insights for one of the largest mobile phone and electronics manufacturers. The manufacturer needed to monitor real-time sales and market movements of mobile devices which requires visualizing data across multiple dimensions including device type, region, price, product attributes and more. Before SingleStore, this manufacturer was dealing with slow and lagging analytical dashboards powered by Teradata. After **augmenting Teradata** with SingleStoreDB, the company could **ingest 12 million rows per second** and **4 billion+ rows of new data per day**, and **perform 160K+ queries in 100ms**.



Leading Media Company in North America: SingleStore helped a leading media company drive **fast analytics** for their **ad sales inventory application**. Prior to SingleStore, the application was powered by RedShift and the media company was dealing with 2+ hours to ingest new data from S3 into Redshift — resulting in sales teams double booking ad spots, ultimately leading to poor customer experiences. Also, Redshift provided 5-minute query response times and poor concurrency just with 100+ users. With SingleStore's augmentation of Redshift, the media company's dashboards saw a **99% improvement** in speed of ingestion with **data ingested in less than 2 seconds**. Dashboard latency had **300x improvement in query latencies from 5 minutes to > 1 second** with **high concurrency for 1000s of users**.



One of the Tier-1 Banks in North America: SingleStore helps a Tier-1 bank in NA with real time ad serving and marketing analytics. They wanted to identify customers in real time across their websites, and place relevant ads in their customer journey for higher campaign conversions. Before SingleStore, they were using **Snowflake and Redis** for their marketing analytics which required manual work and delayed their end-to-end data availability from 4-6 weeks. The Tier-1 bank **augmented Snowflake** with SingleStore to drive **real-time customer ad targeting**, and to **adjust campaign effectiveness in real time**.



Leading Cybersecurity Solution Provider: SingleStoreDB helped a tier-1 cybersecurity and threat detection provider maintain its market position by enabling faster threat detection. Prior to SingleStoreDB, there was a significant lag — up to five minutes — between potential threat detection and reporting, eroding the firm's competitive position. This latency was driven by a combination of factors, including difficulty with streaming data ingestion and supporting a growing volume of queries. With loads upwards of 1,000 concurrent queries per second, Snowflake just couldn't keep up. With **SingleStore augmenting Snowflake**, the company was able to **reduce data ingestion latency by 15x**, lower time to insights by **180x** and handle queries in **under 500ms — even with 1000s of concurrent users**.