

Modernizing the Wealth Management Experience to Drive Competitive Advantage

10ms

millisecond response time on 40,000 current user queries and zero delays during surge events

40,000

concurrent users can access wealth management dashboards, with high performance even during market events

5X

increase in readily accessible historical data, enabling analysis that provides better answers to user queries

Financial services institutions often drive a large share of their income from wealth management solutions. But wealth management is a crowded market, with key players vying for the same high net-worth customers. To provide differentiated wealth management services, financial services institutions need to deliver premium data experiences. Experiences that provide investors with real-time, easy-to-use interactive dashboards that let those customers view, analyze, model and do scenario planning. Legacy data engines and architectures hinder financial institutions' ability to provide real-time data to tens of thousands of concurrent users, especially amid sudden surges in usage due to market events. SingleStore's modern, cloud-native, distributed data architecture enables fast data ingest along with high-performance queries and elastic scaling to deliver insights and analytics fast. That's why half of the top 10 financial services institutions in American—and more around the globe—run on SingleStore.

Business Goals

A top five North American financial services organization was keenly focused on improving the wealth management experience and building meaningful relationships at scale for its high-net-worth clients. The financial services firm needed a solution that met the demands of its service and customers, including fast, consistent response times and high concurrency in the face of dramatic or unexpected market events. Nearly **40,000 users** access the dashboards, resulting in dramatic spikes in concurrency, especially during critical market events.

But portfolio dashboards for top investors are routinely powered by legacy databases that can't reliably deliver millisecond query responses under load to deliver comprehensive portfolio views when it matters most - during volatile market events. The legacy database powering portfolio dashboards at this company was routinely stressed by fast-changing market events and the tens of thousands of users accessing the system. The existing database infrastructure often struggled pulling disparate market data together fast enough, often requiring batch processing or ingest delays. Whenever the database was being updated with live market data, it struggled to deliver reliable performance for ad hoc queries from user interactions with the dashboards.

To provide real-time wealth management dashboards to clients, the business goals were to:

- Eliminate event-to-insight latency to deliver a **"no-spin", no lag** dashboard to customers
- Improve the customer experience by setting and meeting SLAs
- Invest in a solution that would plug into its **existing tools** and data sources
- Move to **cloud-native** technology to enable scale as users grow

Top Five Financial Services Organization Uses Fast Analytics to Drive Better Wealth Management Experiences

“...Analytics and AI can potentially unlock \$1 trillion of incremental value for banks, annually”

- McKinsey, 2020

Technology Requirements

To eliminate the waiting time between a request for information and updating the dashboards, they needed to move from batch data processing and loading to streaming data as it came in. Batch uploads are the default for Hadoop/HDFS, but during the batch uploads, queries were locked out, which users found unacceptable.

To enable further analysis (whether by data analysts or other tools), they needed data to be in **relational SQL format**. This was critical because its previous query acceleration attempts with in-memory cache solutions resulted in breakage of existing SQL queries and introduced reliability and durability issues adding complexity to the architecture.

Of top priority was parallel, high-scale streaming data ingest and high concurrency to allow millions of real-time queries across 40,000 users, run more than 50 queries in parallel within **10-20 millisecond (ms)** response times, and access to of up to 20 years portfolio history.

See Figure A for the “before” architecture.

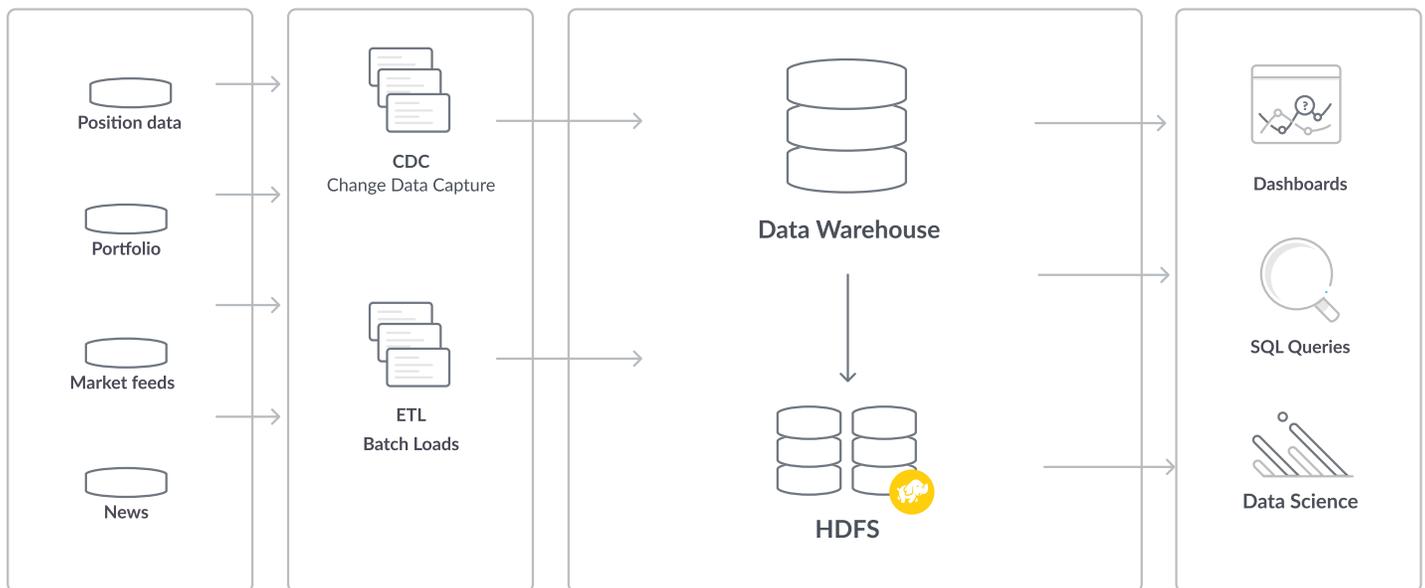


Fig A - “Before” Architecture: Legacy architecture with batch ELT and data warehouse

Why SingleStore? The Unified Database for Fast Analytics

SingleStore offered an innovative, modern and easy-to-use solution that accelerates query performance, driven by durable performance, standard relational SQL, and fast and scalable data ingestion with full ACID compliance. The fact that SingleStore can run anywhere—on-premises environments, in containers, or on any of the leading Cloud platforms—was also key to the financial services organization’s selection of this unique solution.

More importantly, SingleStore’s ability to easily scale without performance degradation was critical. Scaling the ingest and query performance to support more than 40,000 users simultaneously was also a key criterion for selection.

One of the key differentiators is that SingleStore delivers on the customer’s “no spin” requirement - meaning that end customers do not have to wait to run a dashboard, as they did before, especially during turbulent market conditions with higher numbers of concurrent users. This faster service provides them with competitive differentiation by meeting growing customer demands.

The financial services organization **augmented its Hadoop/HDFS** with SingleStore to power its wealth management solution and leaving Hadoop/HDFS as a data lake for long-term data storage.

Ingest and query processing run lock-free, simultaneously and there is no query downtime during batch updates. Data from multiple sources, such as position data and portfolio, market feeds, and news, is ingested into SingleStore in real-time, and queued with a central Kafka cluster that integrates directly with the rowstore tables of SingleStore.

SingleStore’s native Kafka integration ensures each message achieves exactly-once semantics. Lastly, the standard ANSI SQL interface maintains portability of the existing reporting queries to accelerate time to market while utilizing the existing skills and analytics tooling already in place.

See Figure B for the “after” architecture.

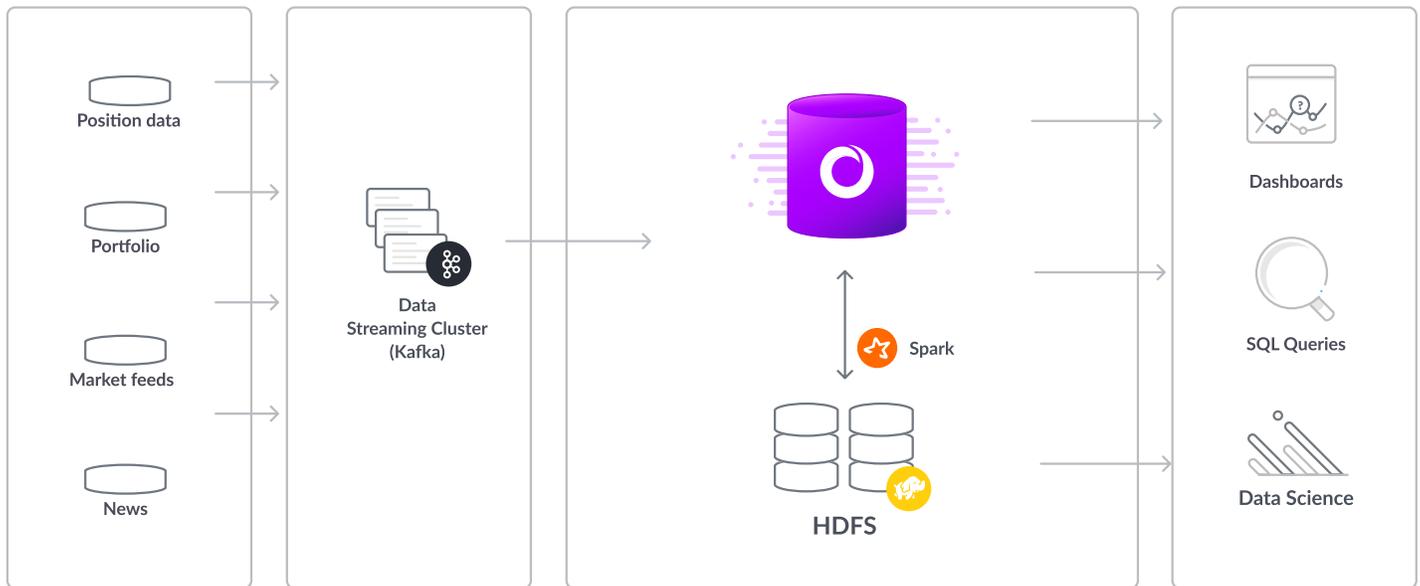


Fig B - Architecture with SingleStore: Modernized wealth management architecture

Business Outcomes

10 ms queries enable modern portfolio management

Low query latency was achieved within **10–20 milliseconds** across all clients and under a variety of highly concurrent workloads, down from 100ms of latency with the previous architecture. SingleStore's distributed, memory-optimized architecture delivers low-latency queries while its durable transaction logging model ensures each event is logged to disk.

Fast access to its premium data services enables them to focus on delivering a larger share of income from its wealth management solutions. As such, the financial services organization was able to improve customer engagement. It now has the ability to deliver modern portfolio management, building meaningful relationships at scale for its high-net-worth clients.

40,000 concurrent users access premium data services

The financial services organization achieved predictable sub-second query responses for more than 40,000 clients across a variety of intensive workloads. High concurrency allows **millions of real-time queries** across tens of thousands of users with zero delays in queries if market events occur. This has enabled them to set and meet aggressive SLAs and eliminate event-to-insight latency to deliver faster data-driven actions.

Now, the financial services institution's high-net-worth clients have fast access to comprehensive portfolio views when it matters most.

5x increase in historical data available for fast analysis

Since the customer modernized its architecture from batch loading to streaming data into SingleStore DB, they can retain much more historical data—**five times more**—in SingleStore for fast access.

Its legacy architecture could only be able to keep one year of historical data in Hadoop, and it had to archive 19 years of historical data in HDFS. Now it stores five years of historical data in SingleStore DB and archives the remaining 15 years of data in Hadoop/HDFS. This enables its investors to deliver more data-driven actions to customers faster.

SingleStore helps half of the top 10 financial service institutions in North America and a number of fast-growing FinTechs around the globe accelerate the performance of their data infrastructure to enable support real-time decision making and deliver exceptional customer experiences.

With a cloud-native, distributed data architecture, SingleStore offers a unified database for fast analytics powering applications that require fast data ingest, high performance queries, and elastic scaling. Its familiar relational SQL interface also enables quick adoption by all teams across the entire enterprise.

With **SingleStore Managed Service**, the fully-managed, on-demand cloud database service you can get started in just a few clicks—on any cloud of your choice. **Test drive now.**

Check out singlestore.com for more details

