

spacetime.io Empowers Investors with Instant Insights to Yield Top Performance on 26.5B Trades per Year with SingleStore

10-100M

10s to 100s of millions of measurements per day

146B + 26.5B

Supports 146 billion quotes and 26.5 billion trades per year

100%

SingleStore is delivering 100% uptime for spacetime.io

-80%

SingleStore data compression is reducing storage by 80%

3/4B

Ingesting 750 millions rows per day

“When you’re talking about a database containing more than 175 billion rows, being able to query a database of that size and pulling responses out in two seconds is very good. Lookup data is nearly instantaneous, which allows us to avoid adding a cached memory product such as Redis.”



Colin Hall

Co-Founder and Director of Engineering, spacetime.io

Today’s equity trading systems struggle to keep pace with both the sheer volume of trading activity and complexity. Managers and traders face tremendous information overload and a loss of situational awareness as they deal with disparate systems and legacy technology stacks that make it difficult to keep up with the pace of the markets. This increases cost and more importantly, increases the risk of poor trading decisions, massive market losses, securities compliance violations, and fines.

spacetime.io, an early stage fintech startup, saw the opportunity to disrupt institutional trading analytics systems by focusing on reengineering trading workflows with an API and cloud-first engineering approach.

spacetime’s trade analytics and recommendation platform, TradeFabric.ai, supports its diversified user base including:

- Institutional traders on the trading floor
- Trade execution committees
- Investment bankers
- Pension funds
- Family offices
- High-net-worth individuals

The platform primarily focuses on equities and also supports cryptocurrency trading.

“The ultimate question we want to answer for users is: did you get proper execution of your trading flow or not? By generating these analytics, what we’re doing is benchmarking their performance against how the market performed during the same timeframe, trading the same kinds of securities in the same region.” — Colin Hall, Co-Founder and Director of Engineering, spacetime.io.

Challenges/Goals

TradeFabric.ai combines intuitive design, real-time analytics, visualization, and machine learning on a secure, scalable, modern architecture. On an annual basis, the platform processes:

- 146 billion quotes
- 26.5 billion trades
- 3 billion foreign exchange rates
- 9.5 billion market snapshots

North American equities trading is concentrated around the opening and closing bells. Peak periods are weekdays from 9:30-10:00am and 3:30-4:00pm Eastern Time; approximately 80% of TradeFabric.ai’s traffic occurs during these timeframes. The platform ingests 100,000 files per day from Amazon S3 buckets. That kind of workload requires ingesting and managing a massive amount of data.

Challenges/Goals (CON'T)

TradeFabric.ai includes many capabilities to deliver critical trading insights to its users as close to real-time as possible, allowing them to act on crucial insights while a trade is still being executed. “Trading happens with the users’ own providers, and we are analyzing trades they’re making with the market data afterward. If their trades are not performing well, it may not be because a broker is performing poorly,” said Hall. “It is often a function of the freshness of the data, the speed limitations of the system, and their trading algorithm. We can help them make beneficial adjustments to all of these things.”

“If their trades are not performing well, it may not be because a broker is performing poorly. It is often a function of the freshness of the data, the speed limitations of the system, and their trading algorithm. We can help them make beneficial adjustments to all of these things.”— Colin Hall, Co-Founder and Director of Engineering, spacetime.io

Investment success requires a range data from historical through real-time

TradeFabric.ai’s reports address performance on a historical, recent, and “right now”/real-time basis, giving traders the flexibility they need for in-depth scheduled and ad hoc analyses. Data pipelines continually bring in real-time market data from the North American equity market and Forex, along with the customer-supplied order information. These analytics are resource-intensive, and meeting the requirements of each report required a sophisticated database underneath the platform.

Mark-to-market pricing and real-time decisions

The foundation of any financial exchange is mark-to-market (MTM) pricing: the current market value of any investment, in real time, rather than book value. Any financial exchange that cannot deliver MTM every millisecond of every day can subject itself and all who trade on the exchange to financial losses fast.

“TradeFabric.ai has two main sets of requirements. One is for information that we’re getting in real time via market and order data. We need to be able to analyze that data in real time, too, and provide analytics to clients at that speed,” said Hall. “For more historical data, we need to be able to work with it in a way that’s performant with reasonable speeds and costs,” said Hall.

spacetime.io needed to ensure that order and client data is ingested in real time, because otherwise, mark-to-market trades could be in violation of strict regulations. TradeFabric.ai users could fail audits or face significant fines. Without real-time capabilities, users would also incur opportunity costs, such as paying too much for securities or selling too low. “When people are interacting with our sites, we want the interactions to be as real-time as possible. So the database has to be real-time,” said Hall.

“When people are interacting with our sites, we want the interactions to be as near-real-time as possible. So the database has to be real-time.” — Colin Hall, Co-Founder and Director of Engineering, spacetime.io

Technology Requirements

The data and speed requirements TradeFabric.ai’s reports vary:

- **Historical reports** reveal how brokers and their trading strategies do over time, ranging from one fiscal quarter to two years. The system needs to accommodate 10,000+ client orders and millions of order fills, with a configurable timeframe that’s evaluated on-demand.
- **Recent reports** look at the past few days of trading performance, from one day to one week. The system needs to be able to handle the tens of thousands of orders and 50+ benchmarks per order assessed against market performance, that occur in that timeframe.
- **“Right Now” reports** provide real-time analytics showing how orders are doing in a snapshot of time. Market data is used immediately as it’s ingested and stored in the database, and an order may be re-evaluated multiple times within a minute. Traders need this data to be available “right now” to optimize their performance on every trade.



IMAGE / Example of a Recent report in spacetime.io’s TradeFabric.ai

As a platform built with a scalable, modern architecture in mind, support for the cloud was a must, but not every TradeFabric.ai user can work in the cloud. Some have on-premises support requirements, so the underlying database infrastructure of TradeFabric.ai also needs to offer this flexibility. Other essential elements for spacetime included being able to easily add nodes and scale the architecture, and the ability to pull in historical market data, reprocess orders, and look at specific segments, all of which must load within seconds.

Why SingleStore

As a startup company, spacetime.io was starting from scratch when it came to the database selection for TradeFabric.ai, so the team considered a range of options, including SQL Server, MySQL, and Amazon Redshift. The company has relationships at Microsoft and Amazon and each of the competing databases is good in its own right.

SQL Server was unable to ingest data fast enough to meet spacetime.io's needs, and spacetime.io also could not swallow the cost. "I'm a born-and-bred SQL Server guy," explained Hall, "but when we looked at the on-premises version of the database software, it gets more expensive as nodes are added. It was also complex because we didn't have Microsoft Active Directory in our environment." spacetime.io looked at the hosted version of SQL Server with Amazon RDS and found that while it could spin up that database very quickly, it was complicated to time and execute against.

Hall has also worked in MySQL, so that was the next database up for consideration. However, the lack of support in what is inherently an open source product, where support is left to the community in most cases, was a concern.

spacetime.io's director of analytics was familiar with AWS Redshift, but it lacked an on-premises deployment option and offered limited personal support from AWS.

"That's when we came across SingleStore," said Hall. "We did a thorough evaluation and it wasn't that the other databases were so ineffective at what they were doing, it was that SingleStore was the best product for our needs. Kudos to SingleStore for winning out in a competitive field."

"It wasn't that the other databases were so ineffective at what they were doing, it was that SingleStore was the best product for our needs. Kudos to SingleStore for winning out in a competitive field."

— Colin Hall, Co-Founder and Director of Engineering, spacetime.io

"We have a unified database environment with SingleStore because we can install it locally, use the managed service, and host it ourselves. The coding environment doesn't change as we go from platform to platform. If a client has a requirement to install it on-premises, we can facilitate that," he explained. "You couldn't effectively do that with a database provider that only offers a cloud solution." Primary reasons spacetime.io chose SingleStore to power real-time and historical trade analyses included:

- **Performant database:** SingleStore has in-memory tables for the data that needs to be cached with fast access times, along with sorted compressed data storage for market data.
- **Flexible deployment models:** the database engine can be easily deployed in the cloud, as a managed service offering, and on-premises.
- **Extensible architecture:** SingleStore provides simple, incremental scaling as needs grow.
- **Built-in data ingestion:** native ingestion capabilities are available via SingleStore pipelines. spacetime.io did not need to set up third-party tools or services to ingest data into TradeFabric.ai.
- **Reliable support:** SingleStore offers strong commercial support as well as community options, with advanced administration training choices available and open communication with the engineering team and database specialists.

“We have a unified database environment with SingleStore because we can install it locally, use the managed service, and host it ourselves. The coding environment doesn’t change as we go from platform to platform. If a client has a requirement to install it on-premises, we can facilitate that. You couldn’t effectively do that with a database provider that only offers a cloud solution.” — Colin Hall, Co-Founder and Director of Engineering, spacetime.io

spacetime.io is also using SQL Server as a smaller footprint in an area of its data infrastructure that is more natively suited for its needs.

Solution

spacetime.io began by using [SingleStoreDB Self-Managed](#) on AWS with a MariaDB connector to ingest 100+ TB of market and order data annually and compress it down to 16 TB in columnstore. The solution powering TradeFabric.ai uses AWS S3 combined with SingleStore [Pipelines](#) to pull in real-time market data. Data in the columnstore uses in-memory and caching, while rowstore is entirely in-memory.

“Being able to use SingleStore Pipelines instead of a third-party ingestion application was appealing. We didn’t have to set up a Kafka bus. We set up an Amazon S3 bucket and pipelines to ingest the data, with no software in between,” said Hall. “Less software equals fewer problems.”

“We have a unified database environment with SingleStore because we can install it locally, use the managed service, and host it ourselves. The coding environment doesn’t change as we go from platform to platform. If a client has a requirement to install it on-premises, we can facilitate that. You couldn’t effectively do that with a database provider that only offers a cloud solution.” — Colin Hall, Co-Founder and Director of Engineering, spacetime.io

Amazon S3 is also used as part of TradeFabric.ai’s recovery model. “Since all the data is in Amazon S3, theoretically, if the database blew up, we could reingest the data in the S3 buckets if we really needed to,” said Hall. spacetime.io still does use a few small SQL Server instances alongside SingleStore for tenant and user configuration, and API management.

spacetime.io had a year’s worth of market data to incorporate into TradeFabric.ai to start. It relied on SingleStore’s scalability to ramp nodes up and down to accommodate this batch ingestion process in a timely manner. Initially, the process was estimated to take five days, but scaling from four nodes up to 16 reduced it to just one day.

After testing and finding an additional 10-15% performance improvement, spacetime.io moved to [SingleStoreDB Cloud](#). “The ability to add nodes and be able to scale the architecture up and down almost on demand was very appealing. We can conserve costs, hardware, and even credits if we’re on the managed service,” explained Hall.

Outcomes

A growing number of large institutions use TradeFabric.ai concurrently, and its SingleStore-powered architecture can support massive concurrency as spacetime.io's user base grows:

Real-time insights drive top broker performance

spacetime.io improves time-to-insight for traders using TradeFabric.ai, allowing them to make informed business decisions in the moment. Ingestion is near-real-time, bringing in 750 million rows every day, with data imported into SingleStore almost simultaneously as it is written to Amazon S3.

spacetime.io's SingleStore-powered TradeFabric.ai platform empowers traders by drilling down into individual orders, aggregating orders, creating report cards, and generating other data visualizations that answer the question of whether they're getting good prices on their trades and executing compliantly.

“Lookup data is nearly instantaneous, which allows us to avoid adding a cached memory product such as Redis.” Colin Hall, Co-Founder and Director of Engineering, spacetime.io

For market data stored in columnstore tables and used for order benchmarking, queries return responses in less than two seconds. “When you're talking about a database containing more than 175 billion rows, being able to query a database of that size and pulling responses out in two seconds is very good,” Hall said. “Lookup data is nearly instantaneous, which allows us to avoid adding a cached memory product such as Redis.”

For computed analytics, such as millions of orders stored in-memory in rowstore, TradeFabric.ai has sub-second retrieval times. Cached market data comes in under three seconds, and historical analysis on large data sets takes just 5-10 seconds.

SingleStore is delivering 100% uptime

“We actually have 100% uptime in our database environment with SingleStore,” said Hall. Having a database that's available to you and reliably ingests the data when you need it is essential when you do a lot of analytics,” said Hall.



SingleStore 80% compression reduces 100+TB storage to 16TB

SingleStore's compression enables spacetime.io to massively reduce its storage needs for market and order data from 100+TB uncompressed data to 16TB. By controlling the storage size for these large databases, the company can scale cost-effectively without needing to compromise on the amount of data used by TradeFabric.ai for historical and real-time analytics.

“Compressing” overnight analysis to near-real-time optimizes trader experience

“Traders are used to overnight analysis and we're offering something that's actually near-real-time,” said Hall. “For most queries, the pages come up within a few seconds. If traders had to wait 30-40 seconds for a promised real-time experience, they would use something else.”

Hear Colin Hall speak about how spacetime.io is innovating with SingleStore to deliver business results in this [webinar](#): Accelerating Fintech Applications at Scale



SingleStore is helping companies compete and win across every vertical. [Learn More >](#)