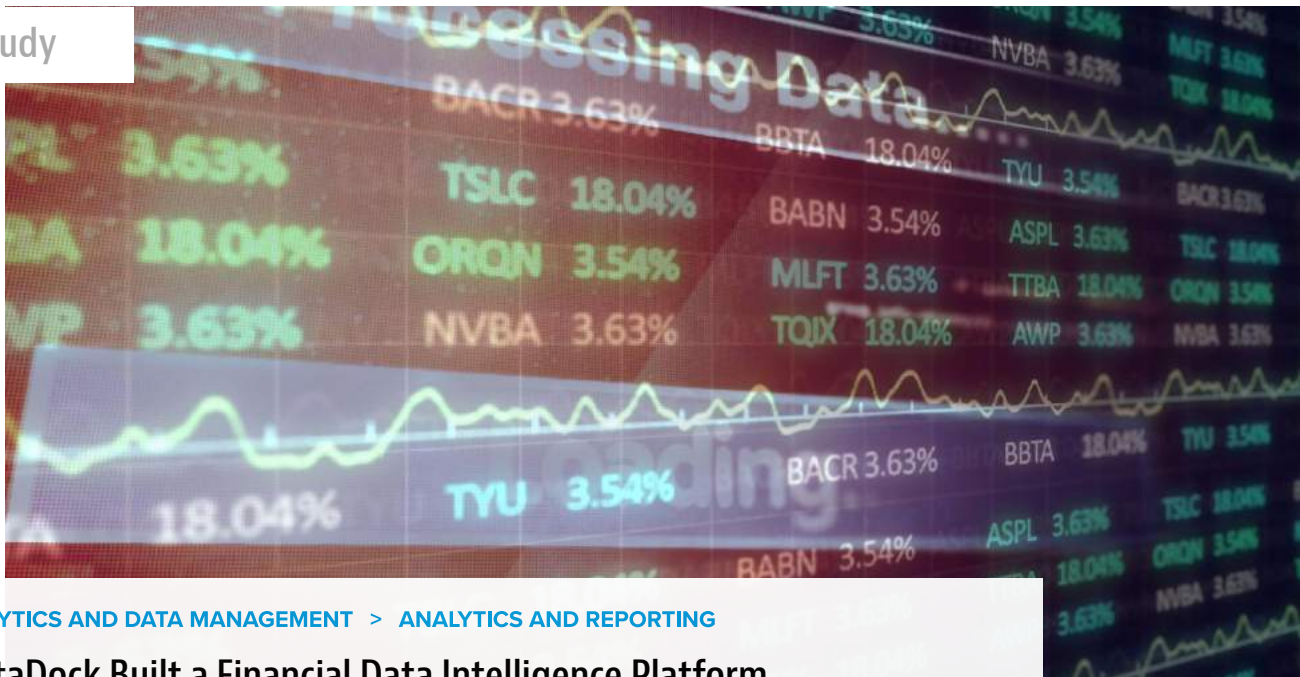


Case Study

Alamy

[DATA ANALYTICS AND DATA MANAGEMENT](#) > [ANALYTICS AND REPORTING](#)

How DataDock Built a Financial Data Intelligence Platform

Startup firm DataDock envisioned data analytics software that would help make data-driven trading decisions. Learn about the software's development.

Karen D. Schwartz | Feb 07, 2023



The average workday for hedge fund managers, traders, brokers, and other financial movers and shakers is long and stressful, with plenty of nail-biting trades, each requiring full analysis. The consequences of making the wrong decision – or having the wrong timing – can be huge. The pressure extends from the moment markets open to the closing bell.

With backgrounds in trading, Kumaran Vijayakumar and Thomas Wadsworth knew the pressures well. They also knew those pressures could be reduced with the right data analytics software for [making data-driven decisions](#). So about five years ago, Vijayakumar and Wadsworth founded DataDock Solutions, a provider of subscription-based analytics, reporting, and intelligence software to the financial services industry.

Related: [Liberty Mutual's Container Deployment – From Docker to Kubernetes](#)

Selecting Databases

ny

gh

g's

Bing Trouble: Google, OpenAI Are Opening Up Pandora's Bots

sti

Domino Code Assist Aims to Bridge Data Science Skills Gap

ata

us

car

th

1 b

Organizations Face IT Challenges in Hybrid Workforce Management

ar

Why Platform Engineering Is the Future of DevOps and SRE

m

n

Object Notation), a format for storing and transporting data.

Bigger and Better

It soon became obvious that despite MariaDB's positive reputation, its columnstore function wasn't stable enough to handle DataDock's database needs. The team had to make a quick pivot. They settled on SingleStore, a relational database that supports columnstore. SingleStore, which previously was called MemSQL, proved to be a good choice.


The new database platform allowed the system to perform as envisioned, handling many calculations on each transaction through a series of lenses and scenarios. "You might start off with 100,000 transactions, but when you do 100 different behavioral patterns with 100 different days, it becomes millions and millions of rows," Vijayakumar explained. "Here's an example: If you do X on Trade Day 1, then might have done Y or Z on Trade Day 2, and if you did Y on Trade Day 2, you might have done A or B on Trade Day 3."

Once DataDock's main analytics platform was running smoothly, the team moved to its next project, a trading platform that aims to replace spreadsheets, chatrooms, and manual processes and uses real-time data and services. The goal was to enable users within a trading or banking environment to share information about trades in real time.

About
Advertise
Contact Us

Sitemap
Ad Choices

Follow us
f y in



Privacy Policy | Cookie Policy | Terms of Use

CCPA: Do not sell my personal info
Privacy Policy
Terms of Service
Content Licensing/Reprints
Cookie Policy

“Our clients were using things like chatrooms to communicate trades, spreadsheets for how they processed and priced those trades, and other spreadsheets for handling the operations of those trades, and they didn’t have any analytics to overlay on top,” Vijayakumar said. “We wanted to replace all of this with a more structured way of processing things, replace calculations that used to be done in spreadsheets, and overlay the data analytics that is our core bread and butter on top so they would know a lot more about their trades as they happened.”

To achieve this goal, Adamec used SingleStore pipelines for batch loading of [data from AWS S3 buckets](#). Pipelines enable developers to create streaming ingest feeds from various sources, including Apache Kafka and Amazon S3. DataDock uses Apache Kafka like a streaming platform for ingesting data throughout the market day.

The screenshot displays the DataDock Unity software interface. At the top, there is a navigation bar with tabs for 'Order Entry', 'Active', 'Watch', 'Interests', 'R/C Active', 'R/C Watch', 'R/C Interests', 'Exec', 'Filled', 'Cancelled', 'Deleted', and 'Compliance'. The main area shows an 'NLP Entry' field with the text 'AAPL jun sep 150 ps 1.00 bid 2k'. Below this, there are several input fields for order details: Ticker (AAPL), Leg 1 (JUN 150 P), Leg 2 (SEP 150 P), Structure (put spread), Direction (Buy), Price (1.00), Size (2,000), Tie (150.00), Structure Delta (4), Shares (-8,000), Premium \$ (200,000), Trader (Barbra Phipps), Desk (Flow Deriv), Customer (Hoopers), Broker (Kirsten Netsrik), and a Comment field. A central table titled 'AAPL Jun Sep 150 PS vs. 150.00 4d' shows the structure of the order with columns for Struct Ratio, Struct Qty, Individual Legs, Size, Bid, Mid, Ask, and T. Delta. The table includes data for JUN 150 P, SEP 150 P, and Theo Shares. On the right side, there are buttons for 'Submit Active', 'Submit Watch', 'Submit Interest', and 'Reset Order'. The top right corner of the interface shows 'DataDock Solutions' and 'Decryption Key'.

DataDock’s Unity software

Another piece of the puzzle is DataDock’s own change data capture approach, a software process that ensures that users see any changes made by other users in real time. This way, Adamec explained, all users see the same market- and order-related data simultaneously.

To ensure that chats ran smoothly, the team implemented the open source RabbitMQ message broker, along with Tornado, a Python-based web framework and asynchronous networking library, and Tornado’s WebSocket technology. The system also includes a heavy JavaScript layer on top of the browser.

Next Steps for the Financial Data Intelligence Platform

Now that DataDock has launched the product, called Unity, the company aims to refine it and launch new products. Since SingleStore supports JSON, the team plans to move away from MongoDB. “We can use those fields in SingleStore tables to store the loosely structured pieces we need,” Adamec said.

The team is also exploring SingleStore’s Bottomless technology, which essentially separates storage and compute, enabling organizations to [use object-based storage](#) like S3 buckets for database data. Although the Bottomless storage seems promising, Adamec wants to test it out first. If it works well, he said he would consider using it to create a more seamless way to replicate data for [disaster recovery](#) purposes.

In terms of a next project, Vijayakumar is working to find a way to extend DataDock’s analytics further. “If our clients can bring their own data into our platform, we could take care of the unification of identifiers so the records match, and then provide that for business analytics like reporting that will work with their CRM and other systems,” he said.

About the author



Karen D. Schwartz is a technology and business writer with more than 20 years of experience. She has written on a broad range of technology topics for publications including CIO, InformationWeek, GCN, FCW, FedTech, BizTech, eWeek and Government Executive.