DELL Technologies

Solution brief

Speeding trades with AI and HPC

In-depth design guidance for algorithmic trading architectures

Algorithm optimization cycle

1. Decide what product to trade
2. Determine strategy
3. Select exchange
4. Design the algorithm
5. Back-test
6. Put in production
7. Refine the algorithm

Make faster, data-driven trading decisions.

Technology has driven innovation in the financial industry for several decades, and today there is more data and computing power available than ever before. In the digital economy, data — and the IT solutions used to harness it — are often a financial services company's prime source of competitive advantage.

This is especially true for algorithmic trading, a highly automated investment process where humans train powerful software applications to select investments and implement trades automatically. The ultimate evolution of algorithmic trading is high-frequency trading, where the algorithms make split-second trading decisions designed to maximize financial returns. Automating and removing humans from trading has several advantages, such as reduced costs and greater speed and accuracy.

Developing trading algorithms requires a proprietary mix of data science, statistics, risk analysis and DevOps. Then the algorithm is back tested, which involves running it against historical data and refining the algorithm until it produces the desired profits. The algorithm is then put into production, making trades in real time on behalf of the firm. The real-world yields produced by the algorithm produce even more data, which is used to continually train the algorithm in the back end and improve its performance.

This training feedback loop is a data-intensive process. More recently, developers have taken up machine learning, a subset of artificial intelligence (AI), to improve predictive capabilities, using deep neural networks to find trends that trigger buy or sell decisions.

In addition to automation and intelligence, high-frequency trading platforms deliver competitive advantage by placing thousands of trades before the market can react. Therefore, high-frequency trading has led to competition in computational speed, automated decision making, and even connectivity to the execution venue to shave off microseconds and beat other traders to opportunities.

What's more, financial trading firms are continually developing, implementing and perfecting algorithmic trading strategies to stay a step ahead of the competition. This puts significant stress on infrastructure because the algorithm must continuously adapt to new input to remain relevant. As such, the back-end infrastructure must accommodate for live-data feed and quick processing of large amounts of data. Databases must be able to feed the compute engine in real or near-real time to update the algorithm.

The data-intensive training requirements and the need for high speed and low latency mean that these sophisticated algorithms are typically trained and run on High Performance Computing (HPC) systems to provide the rapidity and accuracy required to dominate the market. An HPC system that supports algorithmic trading should be able to accommodate current workloads seamlessly and provide the flexibility, performance and scaling required to continually train and update algorithms to stay ahead of the market.

Learn more

Read the Guide
DellTechnologies.com/HPC

A custom approach to simplifying solution stacks

Dell Technologies has the expertise and experience to design and implement HPC and AI solutions optimized for algorithmic trading. A Dell Technologies HPC & AI Reference Guide, "Algorithmic Trading," provides an in-depth technical analysis of a range of solutions for financial trading firms. It takes a deep dive into the various options available for specific use cases and workloads, including an exploration of emerging trends. It also includes considerations for software, services and infrastructure design with complete architectural design examples, such as:

- **Data lake** configurations for data ingestion using streaming tools such as Boomi, Apache® Kafka® and StreamSets® aimed for low-latency real-time data feed with the Ready Solution for Data Analytics Real-Time Data Streaming. Or Apache Hadoop® with Cloudera® and Greenplum® supported by Dell EMC Ready Solutions for Hadoop, Dell EMC Ready Solutions for AI & Data Analytics with Cloudera CDP Data Center and the Dell Greenplum Reference Architecture with VMware Tanzu.
- **Containers for data analytics** using Kubernetes®, the founding blocks of VMware® Tanzu enabling fast-deployment of data analytics models and applications via Dell EMC Ready Solutions for Data Analytics with Spark® on Kubernetes and Data Science and Advanced Analytics with VMware Tanzu.
- **Compute** using the Dell EMC HPC Ready Architecture for AI and Data Analytics.
- **Scratch storage** with Dell EMC Ready Solution for HPC PixStor™ Storage and Ready Solutions for HPC BeeGFS® Storage.

Dell Technologies and NVIDIA

Dell Technologies and NVIDIA work together closely to deliver unprecedented acceleration and flexibility for AI, HPC and data analytics to help our customers tackle some of the world's toughest challenges. NVIDIA® GPUs are the accelerator of choice for algorithmic trading since they have obvious logic for parallelizing with a straightforward code development and mature numerical libraries. Dell Technologies collaborates directly with NVIDIA to provide you with an integrated and successful build for GPU-enabled solutions.