

WHITE PAPER

# Quantitative Trading with Pure Storage Solutions

Technology solutions to meet the most  
exacting performance demands

# Contents

**Introduction .....3**

**Inside the World of Quant Trading .....4**

**Essential Technologies from Pure Storage .....5**

**Conclusion .....7**

**Additional Resources .....7**

**About the Author .....7**



## Introduction

By definition, quantitative (quant) trading is a high-stakes activity that is characterized by exacting standards, a need for consistency, and an unrelenting tide of change. Quant trading is constantly evolving as strategies quickly lose their edge and new approaches come to the foreground. This constant change requires solutions that meet the needs of both today and tomorrow, including the adoption of AI and ML capabilities at scale.

Despite this evolution in strategy, quant trading has some needs that remain constant: speed, reliability and flexibility. To meet these demanding standards and evolving needs, the very best all-flash, NVMe storage solutions are required.

---

“We combine the best of high-performance S3, FlashBlade® fast object storage, and flexible grid computing to enable an environment best suited for the data science paradigm of ‘fail often, fail fast.’”

**JAMES LAMING, VP GLOBAL HEAD OF INFRASTRUCTURE, OPTIONS TECHNOLOGY**

---



**Maximum data flexibility:** Solutions handle the fullest range of data types, from time series to billions of small files and metadata to massive throughput for large objects.



**Fast:** With up to 60GB/s of read bandwidth per chassis, FlashBlade//S™ can support the sustained random read throughputs needed to support training and testing workflows.



**Cost effective:** Achieve low TCO with integrated hardware, software, and service, including Evergreen® to stay perpetually up to date.



**Always on:** No downtime during maintenance and upgrades, with no performance loss and no data migration necessary.



**Less space, more power:** With up to 10x more data storage density, FlashArray™ uses up to 85% less energy consumption and 96% less rack space.



## Inside the World of Quant Trading

On its surface, the world of quant trading is simple: deliver profits (or “alpha”) by utilizing data and analysis to generate faster and more accurate insights. Getting there is another story, particularly in an environment where the half-life of winning strategies is constantly shrinking and finding new opportunities is always a challenge. Some of the keys to success include:

- **A focus on AI and ML:** The days when AI was “nice to have” are long gone; now it’s a “need to have.” For success, quant trading with AI and ML requires new types of data, including non-traditional or alternative data sources and additional formats such as unstructured data, smaller files, and metadata. Successful AI projects must also be underpinned by high-performance infrastructure that can simplify and consolidate all this data at scale.
- **New strategies:** Traditional asset classes no longer generate returns like they did in the past, and are now being supplanted by new markets such as carbon, cryptocurrencies and other digital assets, including cross-asset class opportunities that are revealed only through advanced analytics.
- **Faster and smarter ways of working:** Alpha can be fleeting, so it’s essential to have flexibility when performing analytics. Being able to scale quickly to research opportunities is essential, as you’ll need to scale up if a promising new strategy is discovered or scale down if research fails to yield insights or you abandon unsuccessful strategies. In addition, even successful strategies require constant refinement. The idea is to cyclically train, test, tune, and deploy as the model evolves and adjusts to market changes.
- **Feeding the feedback loop:** A good strategy is only part of the picture. Successful quant systems must also incorporate AI-based feedback such as transaction cost analysis (TCA) and pre-trade inputs like market sensing and real-time algo selection in a continuous feedback loop to speed decision making, reduce costs, and improve outcomes. Vigorous analysis and back-testing are also essential and require a “the more the better” approach to data, which can lead to more demand on infrastructure with larger data center footprints that traditionally mean more CAPEX outlay for storage and compute.
- **Data in many forms:** Data is the fuel that drives quant trading and it continues to grow with no end in sight. At the same time, the variety of data available, including unstructured and alternative data—such as social media sentiment and satellite imagery—is expanding exponentially. For successful quant trading, flexibility and adaptability (including multi-protocol capabilities) need to be built into both data processes and infrastructure.

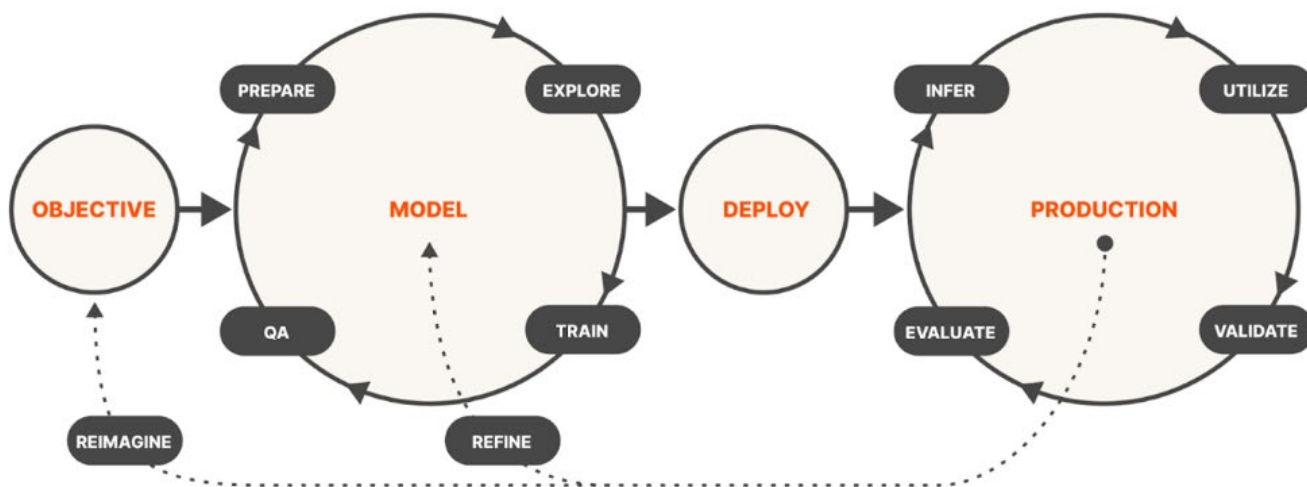


Figure 1. Model development for quantitative trading requires top notch technology for all phases and Pure Storage delivers.



## Essential Technologies from Pure Storage

With unrivaled speed, unparalleled flexibility, and the ability to handle the widest variety of data formats, flash storage is the optimal solution for high-performance data. Pure Storage is an industry leader in delivering all-flash storage solutions that provide unmatched speed, scalability, and simplicity.

“FlashBlade had capabilities we thought were vital: the ability to process large volumes of streaming data in real time, and the ability to make a data set available to multiple workloads simultaneously.”

**MANI VENKATESWARAN, GLOBAL HEAD OF PRODUCT SUPPORT, LIQUIDNET**

- **FlashBlade//S™**

- Distributed, scale-out unstructured data storage provides the flexibility needed to exploit alpha opportunities by effectively balancing loads and processing data at exabyte scale.
- Unique modular architecture allows capacity and performance to be scaled independently, allowing the platform to meet evolving workload storage and I/O needs.
- FlashBlade//S is the Ideal storage platform for unstructured data workloads. It offers multi-dimensional performance for billions of small and large files or objects, delivering both the high throughput and low latency that are critical to AI/ML operations.
- High-performance native NFS, SMB, and S3 protocols along with fully-featured management REST APIs make deployments, management, and upgrades simple and non-disruptive.

Parallel Access	Metadata	Bandwidth	Fast Delete	High File Count	Capacity
Tested in 500k+ Core Environments	>30 million IOPS	300+ GB/second	>50k files/second	2^56	19+PB with 2:1 Compression

- **AIRI//S™**

- A purpose-built full stack pre-integration for AI infrastructure in a cutting-edge, disaggregated, and modular scale-out storage platform for unstructured data
- Recognized by the AI Breakthrough Awards as the “Best AI Solution for Big Data”

- **FlashArray//C™**

- An all-QLC (quad-level cell) array that delivers high endurance, consistent performance, and significantly lower costs for capacity-intensive workloads, offering:
  - o Storage capacity without compromise
  - o All-flash performance at 99.9999% availability
  - o All-flash for all workloads with a no-compromise environment



- **FlashArray//XL™**

- Scale up and scale out with 100% native NVMe (non-volatile memory express) storage.
- Exceptional performance density, allowing you to consolidate more business services—bigger databases, more applications, more users—on fewer arrays
- Optimized for performance at scale with latency as low as 150  $\mu$ s, throughput as high as 36GB/sec, and up to 70% performance boost over FlashArray//X90, all while taking up 90% less space and less power in the datacenter compared to competitive solutions

- **Portworx®**

- The industry's leading Kubernetes (K8s) data platform for building, automating, protecting, and securing cloud native applications
- Manage mission critical applications with zero downtime or data loss while dramatically lowering infrastructure costs
- Reduce the time and resources required to deploy applications via a CI/CD pipeline across any cloud and any infrastructure

- **Evergreen®**

- Evergreen extensible architecture provides non-disruptive upgrades and expansions to enable the best and newest technology removing worry about service interruptions, functionality decreasing over time, or expensive and risky data and hardware migrations.
- You get to choose the best option to consume and deploy your storage: traditional purchase, flexible pay-as-you-go ownership, or true as-a-service with Evergreen//One™.

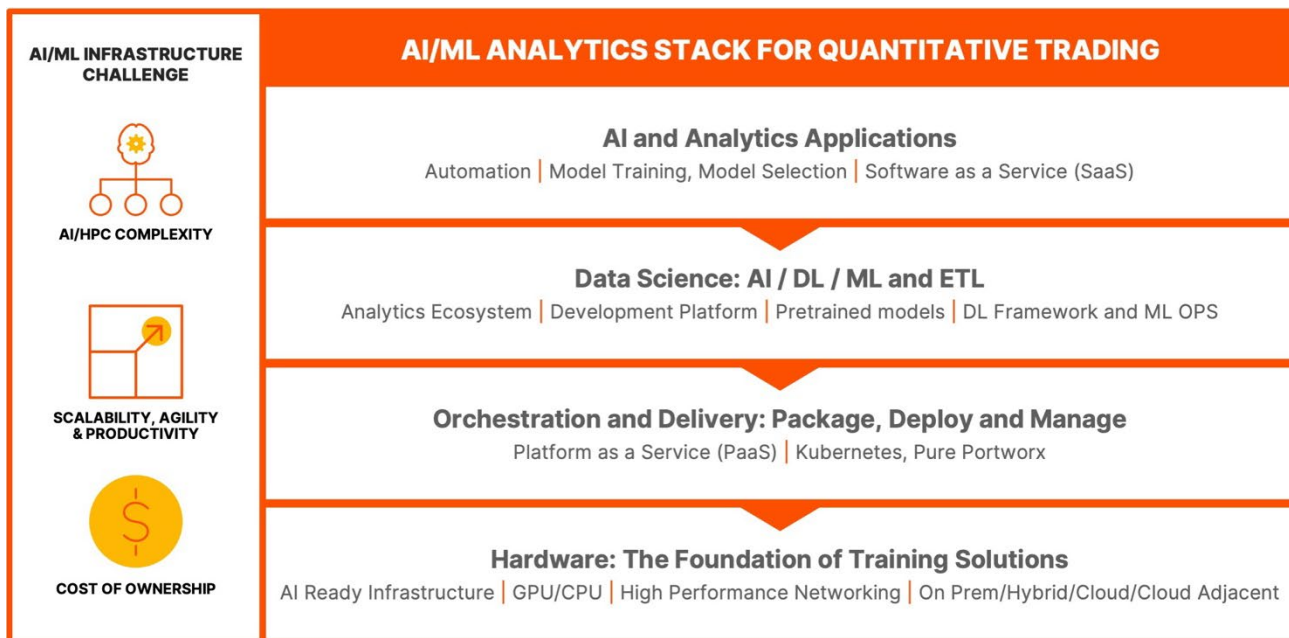


Figure 2. AI/ML Analytics Stack Addressing Infrastructure Challenge in Quantitative Trading



## Conclusion

Even—or especially—in years when traditional investment strategies struggle, quant trading offers traders the opportunity to generate significant profits. In 2022, for example, when the S&P 500 was down more than 19% and global government bonds and credit markets were off by 15-17%, quant strategies (as measured by the SG Alternative Risk Premia Index) driven by various macro factors returned a healthy 3.9%. And that trend is expected to continue in 2023.

Quant trading requires the very best technology to achieve optimum performance both now and in the future. If you can accelerate time to insight, then you can accelerate time to revenue. You can build a consistent, successful track record with a service that can expand to meet your evolving needs and keep pace with new technology solutions.

When microseconds matter, Pure delivers a fully-integrated service with Evergreen that ensures all future needs will be met. Make Pure Storage your trusted partner for quant trading.

---

“...FlashBlade//S...(is) very adept at handling the key trends in enterprise storage going forward: explosive data growth (most of which will be unstructured data); customer concerns about ease of use of their file and object storage platform; increased performance, availability, and capacity requirements; consolidation of workloads onto fewer storage platforms; and an intensifying focus on infrastructure efficiency for both cost and ESG reasons.”

[IDC White Paper: Pure Storage's Next-Generation FlashBlade//S](#)

---

## Additional Resources

### Next Steps

- Learn why TrustRadius ranks Pure Storage as [top rated for scale-out storage](#).
- Find out why Meta chose to partner with Pure for its [AI Research SuperCluster](#).

### Supporting Information

- [Analytics and AI](#)
- [FlashBlade//S](#)
- [Financial Services](#)

## About the Author

Diane Saucier is the Financial Services Director for Pure Storage, leading the marketing efforts in financial services and fintech. Diane has held key roles with global financial institutions and technology vendors, developing multi-asset class product strategy and trading, risk, and compliance solutions. She is a founding Board Member of Women in Listed Derivatives (WILD) and an Advisory Board member for the University of South Florida Fintech Program and John J. Lothian & Company, Inc. Diane holds a bachelor's degree in Economics and Political Science from Northwestern University and a patent for a flexible system for electronic trading.

[purestorage.com](https://purestorage.com)

800.379.PURE

