ANALYTICS FOR FINANCIAL MODELING
FAST WALL CLOCK TIME FOR BETTER TRADING

INDUSTRY
Financial Services

USE CASE
Quantitative Analytics
High Velocity Analytics

BUSINESS REQUIREMENT
Deliver best possible wall clock time for new latency sensitive simulation workloads

SOLUTION
WekaIO Matrix™

BUSINESS BENEFITS
- Delivers equivalent performance to local SSD drives
- Meets, or exceeds, wall clock time service level agreements (SLAs) to ensure best trading strategies
- > 65% storage cost savings by leveraging shared storage instead of individual drives across each node in the HPC cluster
- Greatly simplified storage infrastructure
- Integrated cloud tiering and cloud bursting for better economics and capacity on demand

High frequency trading has undergone a revolution through the application of predictive financial models to trading strategies, and success is measured by the fast execution of scientific model results. Leading quantitative trading companies are constantly seeking out new strategies to gain better market insight and improve trading outcomes. The teams who work in quantitative trading are among the best minds in math, computer science and engineering. As academics from major institutions they look for the most advanced technologies and techniques to develop a competitive edge. The constant evolution of new models presents a high level of unpredictability to the underlying infrastructure and requires a modern IT architecture that can handle the most demanding of storage workloads—data intensive and latency sensitive applications.

THE CHALLENGE: MINIMIZING WALL CLOCK TIME FOR TRADES
Given the academic nature of quantitative trading, it is common to leverage the tools used in high performance computing (HPC) simulations for financial modeling. Many financial trading institutions have enjoyed many years of success with HPC parallel file storage infrastructures, like those from IBM that were designed around large block sequential I/O. But today's trading workloads require writing millions of tiny files out to persistent storage at a very high rate and performing near real-time analytics on the data.

Poor Shared Storage Performance to HPC Clusters
Legacy parallel file systems were developed and optimized decades earlier for spinning media and cannot easily leverage low-latency storage media based on modern flash technology. HPC systems buckle under the small file, latency sensitive workloads and wall clock times can drop as much as 75% compared to local SSDs inside the high-performance compute cluster. The poor performance—and associated increase in wall clock time to complete a market simulations—means that legacy HPC systems are not well suited for the new workloads common in financial services. As a work around, many teams are electing to run their workloads on local-drive SSD storage inside the HPC nodes to restore an acceptable trading window.

Challenging Management of Local Storage Flash
Local SSD storage inside the individual HPC compute notes can deliver great performance, however over time it can introduce new headaches due to the nature of the financial services workloads. Flash technology provides great random, small file performance but unlike disk drives, it has limited endurance and is particularly vulnerable when a workload is write intensive. As drives begin to fail the daunting challenge of how to keep the application environment productive arises. At extreme scale, replacing all the SSD would be a prohibitively expensive proposition as all HPC nodes would have to have local NVMe
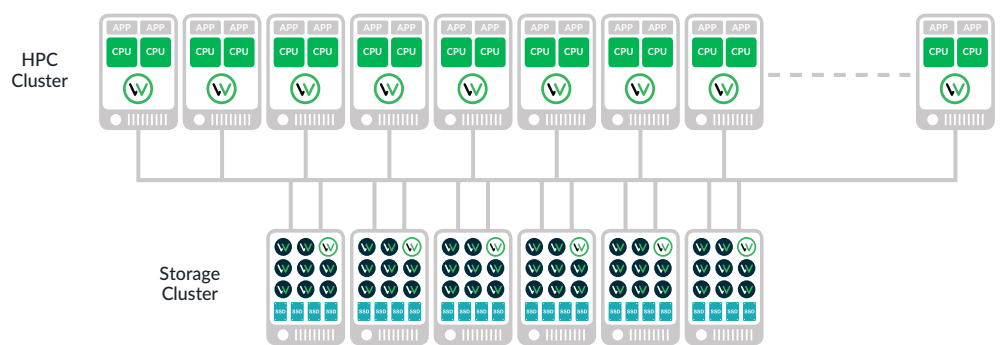
storage even if the application did not need that much physical space. In addition, as drives continue to fail further application disruption is common and evenutally the decision regarding total replacement arises.

THE SOLUTION: WekaIO MATRIX SOFTWARE ON COMMODITY SERVER INFRASTRUCTURE

Many organizations consider SAN-based solutions utilizing iSCSI for all the storage nodes, however these solutions incur their own challenges and performance is not guaranteed. WekaIO Matrix, the world's fastest filesystem, has broken many industry standard benchmark records, and has become the de facto storage standard for Quantitative and High Velocity Analytics use cases.

Matrix is a fully parallel and distributed file system that has been designed from scratch to leverage high performance Flash technology. Both data and metadata are distributed across the entire storage infrastructure to ensure massively parallel access to NVMe drives. With a single namespace that spans local storage and the cloud, Matrix delivers performance that is 3x that of local file systems and 10x that of traditional NAS or SAN.

The Matrix software is highly resilient parallel file system that delivers the highest bandwidth, lowest latency performance to any InfiniBand or Ethernet enabled GPU or CPU based cluster. The POSIX-compliant file system leverages an innovative, highly optimized software stack that leverages modern architectures such as low-latency NVMe-oF and massively distributed metadata. Built for enterprise grade HPC,

WekaIO Matrix Delivers Local Disk Performance to the HPC Cluster

it tightly integrates object storage for best economics at scale and supports features such as snapshots, cloud backup and multiprotocol support including NFS and SMB.

Financial Services customers who use WekaIO Matrix like the fact that the software is fully cloud integrated and has the ability to leverage public cloud compute resources which affords the ability to integrate tiering and cloud bursting in to the workflow. The added bonus of Matrix is the simplicity of the production solution. Many storage alternatives require complex performance tuning to achieve the performance characteristics required, WekaIO Matrix is simple to manage out-of-the-box with no specialized tuning.

BUSINESS BENEFITS

Implementing WekaIO Matrix in your Financial Services environment:
- Delivers equivalent performance to local SSD drives
- Improves wall clock time to help ensure service level agreements (SLAs) are met or exceeded
- Delivers > 65% storage cost savings by leveraging shared storage instead of individual drives across each node in the HPC cluster
- Dramatically simplifies infrastructure management with no HPC compute cluster downtime
- Enables leveraging public cloud elastic compute for new application workload support

To find out more or to arrange for a free trial, go to https://www.weka.io/get-started or contact us at info@weka.io.

WEKA.IO

2001 Gateway Place, Suite 400W, San Jose, CA 95110 USA   T 408.335.0085   E info@weka.io   www.weka.io