MACHINE LEARNING FOR AUTOMOTIVE
IMPROVE UTILIZATION OF GPU RESOURCES

USE CASE
Machine Learning

COMPANY
Leading Automotive Company

BUSINESS CHALLENGE
Improve the utilization of its GPU platform to run more training Epochs faster, accelerating time to market and improving ROI on expensive GPUs

SOLUTION
WekaIO Matrix™

RESULTS
- 7X improvement in performance compared to NFS based ALL-Flash NAS
- 3X improvement in metadata performance improving GPU utilization
- 50% of the cost of high performance storage tier by utilizing standard X86 servers with software
- Improved overall productivity of data scientists by eliminating local-copy to NVMe drives in GPU servers
- Greatly simplified infrastructure with integrated data lake in a single namespace

Autonomous vehicles (AV) promise to change the face of driving, but the success of the industry rides on the ability to train AVs to operate flawlessly in all road conditions, weather types, and according to varied driving laws. Deep learning systems place a significant burden on storage and computational infrastructure because the rate of data acquisition and data processing mask any prior workloads. A single AV will generate over 40TB of data in 8 hours[1], a burden that multiplies with a fleet of training vehicles. Daily data collection is measured in multi-petabytes and must be ingested to the training data lake for pre-processing to support the training sets.

THE CHALLENGE: KEEPING EXPENSIVE GPU SERVERS FULLY UTILIZED

A leading AV manufacturer developed its training models on a small cluster of ten GPU servers and a shared data repository built on a commercial all-Flash NAS appliance. This solution worked well for initial model development, but it became clear that it could not scale to a full production system of 50 to 100 GPU servers that had to handle petabytes of new data daily. The challenges faced with the early development model had to be addressed before moving to the production phase.

Poor Performance to GPU Clusters

The data sets used to train the AI models were comprised of millions of tiny 4K image files interspersed with some medium sized 100MB files and a few very large multi-gigabyte files, all of which needed to be read at very high bandwidth (up to 10GBytes/second) to keep the GPU servers fully utilized. An NFS based All-Flash NAS could only achieve 1-1.5GBytes/second, leaving the GPU servers starved of data. Scientists had struggled with the limitations of NFS and had devised a method of copying training data sets into the GPU server local NVMe drives to improve overall training times, but this solution would not scale to production-sized clusters. In addition, the training systems were required to routinely traverse the file system to choose new files for the training runs. A single file system traverse was taking hours to execute, while the GPUs remained idle. Just 10% GPU under-utilization would result in over $2M in wasted GPU server infrastructure on a production cluster of 50 GPU servers.

The production solution had to have a **high bandwidth, low latency** ingest rate that could saturate a 100Gbit network link–up to 10 GBytes/second per GPU server to meet the ingest demands.

Exascale Data Growth

Data growth was anticipated to be 100-200PBs per year and far exceeded the namespace available on commercial scale-out NAS appliances. Tiering to a third-party solution was

[1] https://www.networkworld.com/article/3147892/internet/one-autonomous-car-will-use-4000-gb-of-dataday.html

not an acceptable option because the training models needed to have perpetual access to the catalog. Moving data back-and-forth between tiers would prove impossible and would invariably result in manufacturing inefficiency. The solution had to have a **single namespace** that could scale to exabytes and beyond without human intervention.

Mounting Cost Overruns

Early development systems had utilized all-flash appliances from traditional storage vendors for the training catalog. But the team realized that they could not afford to scale to production levels with this architecture because it was cost prohibitive to save petabytes of data on Flash. Flash was only required to deliver the performance into the GPU clusters but did not make sense for the data catalog. The solution had to integrate a **disk-based architecture** with its much lower cost structure and be **software-defined** for flexible procurement options of the underlying hardware infrastructure.

THE SOLUTION: WEKAIO MATRIX SOFTWARE ON COMMODITY SERVER INFRASTRUCTURE

After thoroughly evaluating the entrenched storage suppliers, no commercial vendor could solve all the challenge of performance, scale and cost. The team researched new vendors in the storage space and WekaIO Matrix was the only solution that had an architecture that could meet their requirements.

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.

The software's ultra-low latency network stack runs over Ethernet or InfiniBand, delivering the lowest latency and highest bandwidth performance for the most demanding data and metadata operations. The software has an integrated tiering mechanism that leverages S3 object storage to allow a single namespace to scale to exabytes while maintaining full data locality to the training applications. And because it is software only, customers can negotiate directly with infrastructure suppliers and Flash vendors to get the best possible economics at scale.

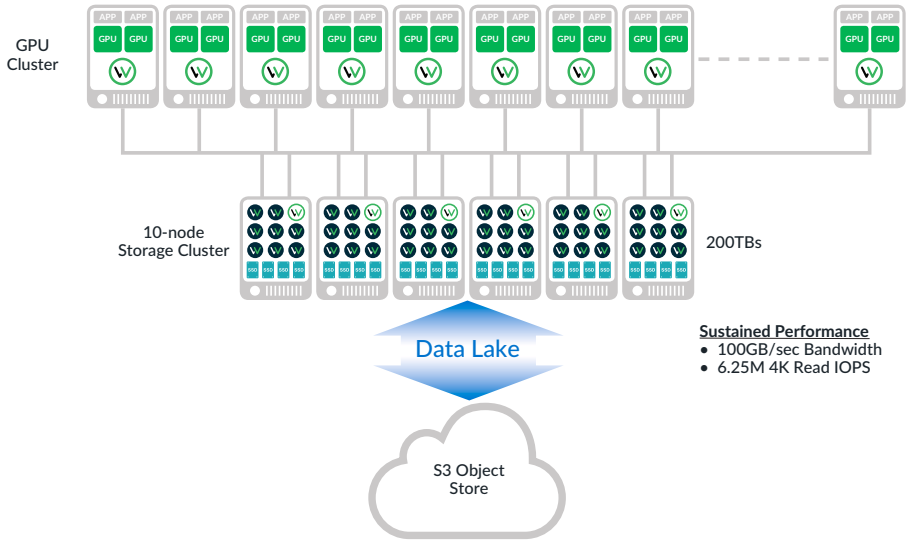

WekaIO Delivers Full Bandwidth from S3 Data Lake to the GPU Cluster

THE RESULT: 7X PERFORMANCE IMPROVEMENT AT 50% OF THE COST OF ALL-FLASH APPLIANCES

By choosing WekaIO Matrix, the world's fastest file system, the team was able to:

- Improve performance to the product GPU cluster by 7X compared to legacy NFS based All-Flash NAS solutions
- Improve performance by over 2X compared to locally attached NVME drives
- Improve metadata performance by 3X resulting in better GPU utilization
- Reduce performance tier costs by 50% compared to All-Flash appliances
- Provide fully integrated "hands-free" access to training catalog with integrated tiering
- Achieve the lowest overall cost with disk based object storage solution

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

W01UC201803