SOLUTION BRIEF

Improving GPU Utilization for Autonomous Vehicles

WEKA

Highlights

- 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

Challenge

- Reading millions of files (of various sizes) at a very high bandwidth of up to 10GBytes/second to keep the GPU servers fully utilized
- GPU under-utilization wastes over $2M in GPU server infrastructure per cluster
- Keeping pace with data growth of 100-200PBs per year
- Mounting cost overruns and inflexible procurement options of the underlying storage hardware infrastructure

Solution

To improve GPU Utilization, one autonomous vehicle manufacturer chose the powerful combination of Quantum ActiveScale™ cloud object storage system and WekaIO Matrix™ software.

The combined solution enabled them to acheive the bandwidth needed from the Data Lake to the GPU Cluster to ensure full utilization of GPU resources. The solution also provided the high scalability, data integrity and extreme durability needed to help ensure data is protected and always available for a long time or forever if needed.

Leveraging Two Powerful Technologies Was the Key to Improving Utilization of GPU Resources in Autonomous Vehicles

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 and weather conditions, and according to varied driving laws. Deep learning systems place a significant burden on storage and computational infrastructure because of the rate of data acquisition and processing. A single AV will generate and consume over 40TB of data in 8 hours of driving*, 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. The original all-flash NAS system 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 10GBytes/second per GPU server to meet the ingest demands.

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

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 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: The Powerful Combination of WekaIO Matrix Software and Quantum ActiveScale Cloud Object Storage System

After thoroughly evaluating many options, the team found a clear winner in the combined solution of Quantum ActiveScale cloud object storage system with WekaIO Matrix. WekaIO Matrix with ActiveScale was the only solution that had an architecture that could meet their requirements. Matrix is a fully parallel and distributed file system that distributes both data and metadata across ActiveScale's highly scable storage infrastructure to ensure massively parallel access to data. Matrix delivers low latency and high bandwidth performance for the most demanding data and metadata operations, while ActiveScale delivers extreme data durability and integrity at petabyte scale to help ensure valuable data is protected and always available. ActiveScale delivers up to 19 nines durability and robust data integrity checks occur automatically and transparently to help protect data for a long time or forever if needed.

Conclusion

The enormous size of data repositories and very high bandwidth requirements present a daunting challenge to AV manufacturers. Huge amounts of data need to be quickly stored, processed, analyzed and kept for long periods of time. The software-defined storage solution from WekaIO and Quantum delivers the necessary performance, scale, and efficiency needed to address these challenges with fewer resources and lower cost than traditional approaches.

The combination of Quantum's ActiveScale cloud object storage system with WekaIO's Matrix fast file system enabled one AV manufacturer 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
- Dramatically lower overall costs while improving utilization of GPU resources

To learn more, visit www.Weka.IO and www.Quantum.com.

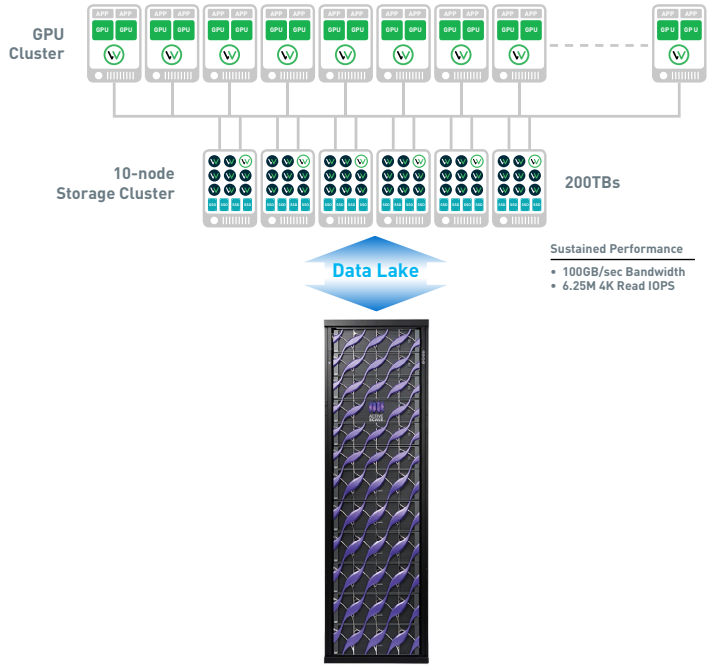

WekaIO Delivers Full Bandwidth from ActiveScale Data Lake to the GPU Cluster

Quantum

Quantum technology and services help customers capture, create, and share digital content—and preserve and protect it for decades at the lowest cost. Quantum's platforms provide the fastest performance for high-resolution video, images, and industrial IoT, with solutions built for every stage of the data lifecycle, from high-performance ingest to real-time collaboration and analysis and low-cost archiving. Every day the world's leading entertainment companies, sports franchises, research scientists, government agencies, enterprises, and cloud providers are making the world happier, safer, and smarter on Quantum. See how at **www.quantum.com**.

www.quantum.com