KIOXIA

White Paper

[KIOXIA Case Study] How NVMe™/NVMe-oF™ Technology Helps to Shorten Machine Learning and Deep Learning Workflow

Background

As flash memory has increased in density through SLC->MLC->TLC->QLC, the storage capacity of each memory cell has increased. Furthermore, the multi-layering technology has greatly increased the density of flash memory, leading to higher capacities. However, various technologies have to be accumulated to maintain and improve flash memory reliability as it continues to grow in capacity. For example, ECC (Error Correcting Code) is one of the key technologies to ensure data reliability in flash memory. ECC corrects and recovers from various errors that can occur during memory operation, making users unaware the error occurred.

KIOXIA has developed a highly reliable ECC algorithm by using improved algorithms in addition to existing ECC technology. KIOXIA is further developing more accurate reliability algorithms based on non-linear neural network models to bring memory performance and reliability to higher levels.

Current Problems and Directions for Solution

To develop and optimize neural net reliability algorithms, it is necessary to accumulate as much data as possible (actual patterns of memory usage) and different algorithms for each phase by analyzing the data in a database. KIOXIA's reliability algorithm development team built a 1Gb Ethernet network of popular x86-based CPU servers and NAS-type file servers as infrastructure for creating databases required to both analyze the data and optimize the algorithms. Jobs have been submitted over the network for analysis work.

However, with this infrastructure, data transfer time becomes a bottleneck, slowing down improvement in analysis accuracy. In addition, multiple analysis passes of large amounts of data per pass was necessary, while traditional methods had limitations in improving algorithm accuracy, even with large amounts of data.

KIOXIA's research infrastructure consists of GPUs to implement Deep Learning, which has been rapidly developing in recent years, to speed up training over x86-based CPU servers. In Deep Learning, accuracy improves significantly with increasing amounts of data. For training data collection, a process of analyzing as many as 80-100 different types of data is executed in parallel, generating the same number of databases, resulting in an overall database size of 1 TB, which stores millions of different parameter samples.

In addition, Deep Learning accesses the entire database for training iteratively, thus generating accurate algorithms through Deep Learning training depends on how efficiently a GPU with high processing performance can be operated - how efficiently data can be continuously fed to the GPU.

KIOXIA's reliability algorithm development team first tried to solve this by connecting NVMe SSDs directly to the GPU server as high-speed block storage. However, NVMe SSDs alone have limited capacity, thus training with larger data required interrupting GPU processing and dividing the analysis data according to the SSD capacity. In addition, network transfer speeds had become a bottleneck in transferring data from NAS to GPU servers and storing databases. Consequently, it has not yet been achieved a significant reduction in algorithm optimization time due to the difficulty of running expensive GPUs 24/7.

Solution Based on NVMe SSDs and KumoScale™

KIOXIA's reliability algorithm development team identified an ability to operate high-speed block storage and shared file system simultaneously as key requirement and cited NVMe performance as the storage technology to solve these problems. We were also considering NVMe-oF technology, which allows more flexible and efficient allocations of storage to GPUs than DAS-type attached storage, and the operation of a shared file system to allow simultaneous access to data from multiple GPUs.

Block Storage Requirements

- Ability to load a large amount of data into training applications at the same speed as SSDs
- Flexibility to allocate the necessary virtual block storage depending on data set size
- Available for pre- and post-analysis processing

Shared File System Requirements

- High speed performance required for training processing
- Shared file system accessible from all GPU servers
- No need to duplicate data used for distributed Deep Learning processing among multiple GPU servers
- Available for pre- and post-analysis processing

In addition, we focused on accelerating network connectivity with NVMe-oF technology, with the requirement to use 100Gb Ethernet networking and high-speed network transfer protocols.

Key Components and Technologies

- GPU server for Deep Learning (100Gb dual port Ethernet, RoCEv2 network protocol)
- Network switch (100GB/s redundant data path, 1GB/s control path)
- NVMe-oF storage system

  KumoScale storage node

  » NVMe-oF connection (100GB/s dual port Ethernet, RoCEv2 network protocol)
  » Virtual Block Storage
  » Cross Domain Data Resiliency (data protection) 2 replicas
  » Shared File System (NFS - GlusterFS®)

  KIOXIA NVMe SSD 3.2TB (Total of 150TB)

System Configuration

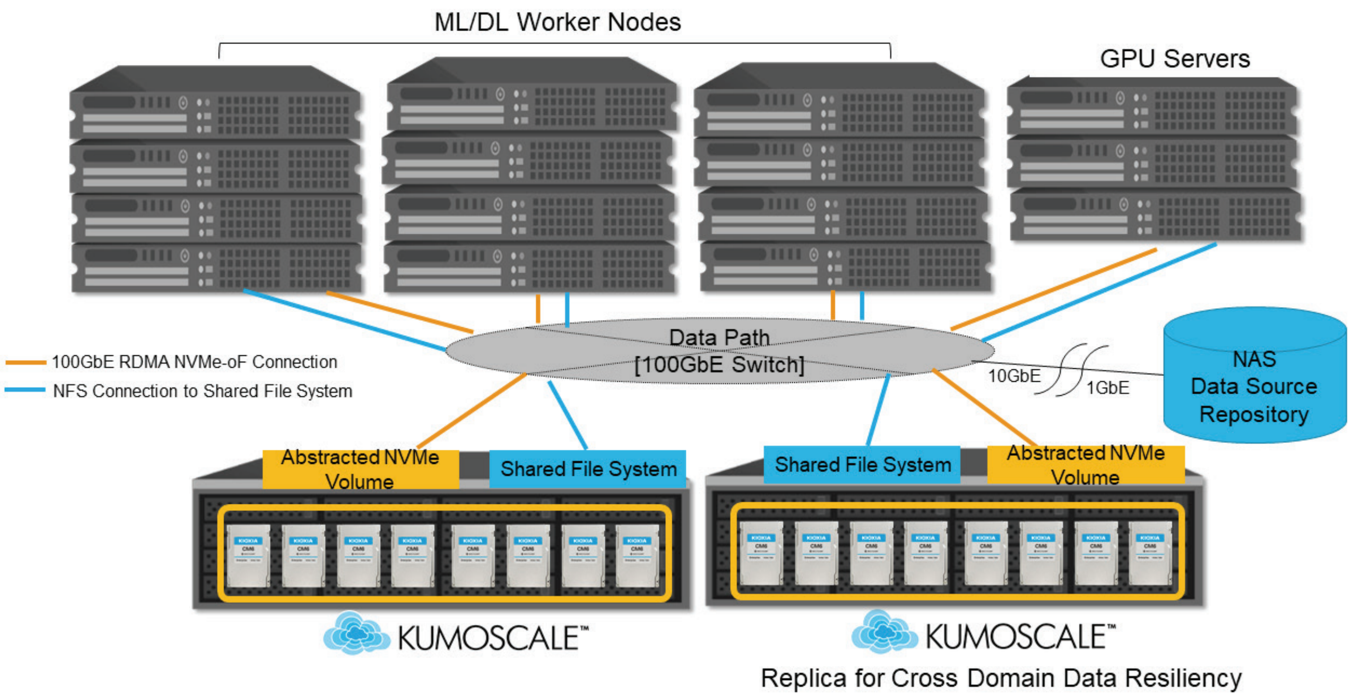

Implementation Benefits and Future Outlook

According to KIOXIA's NAND reliability development team, significant time has been saved in feeding data used for training processing into the GPU system compared to the previous system using DAS-attached NVMe SSDs and NAS. Furthermore, the KumoScale storage system completed the cycle in almost one day, which used to take about one week, including the database creation for training.

In addition, the training process using multiple GPUs has been performed without data replication by building a shared file system on the KumoScale storage system, making the process as efficient as originally intended.

On the other hand, increasing data sets to analyze have required an average quarterly increase in storage capacity of approximately 15 TB. More frequent capacity adjustments to the virtual block storage and shared file system carved out from the KumoScale storage system to the worker nodes are also necessary to accommodate the size of the data sets to be analyzed.

The analyzed data and training data are archived on a NAS, and the future challenge for us is to utilize these archived data flexibly as new NAND devices are developed. KIOXIA's NAND reliability development team and KumoScale software development team are working together to address these challenges.

NVMe SSD & NVMe-oF optimize ML/DL with NVIDIA GPUDirect® Storage

- GPUDirect Storage and NVMe SSDs dramatically reduce turnaround time for ML/DL.

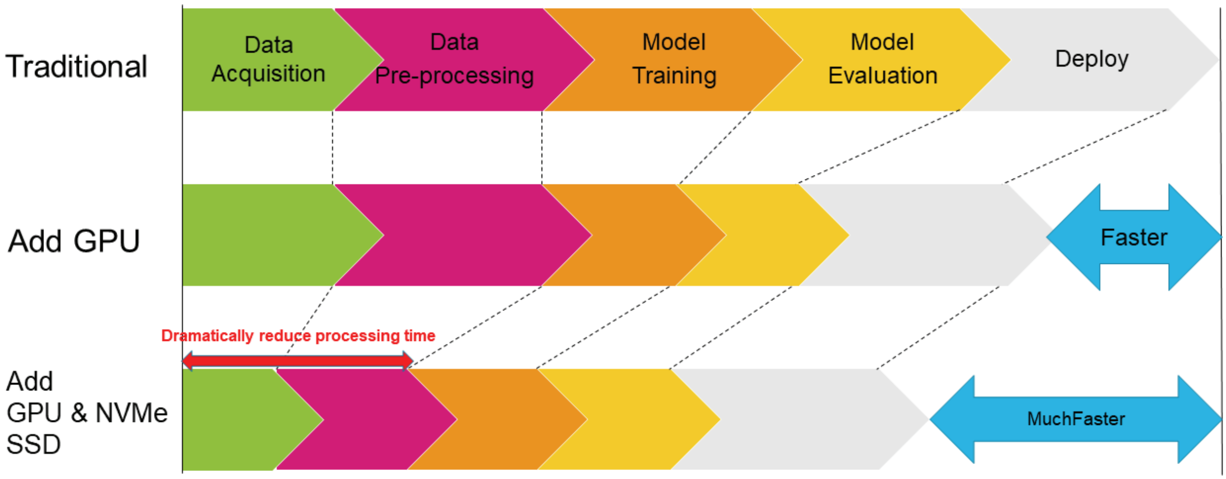

*NVMe and NVMe-oF are registered or unregistered marks of NVM Express, Inc. in the United States and other countries.
*Other company names, product names, and service names may be trademarks of their respective companies.