



Flash Memory Summit

Analyzing & Characterizing AI Workloads versus MLPerf Storage

John Mazzie, Principal Systems Performance Engineer

Sujit Somandepalli, Principal Systems Performance Engineer

Ravuf Syed Abdul, Sr. Systems Performance Engineer

Wes Vaske, SMTS Systems Performance Engineer

- What does an AI workload look like?
 - Run MLPerf Image Segmentation (Unet3D)
 - Analyze a workload trace
- MLPerf Storage Benchmark
 - What is it?
 - Run and trace the benchmark
- Comparison of real-world workload to GPU-less emulation

Trace Analysis of an MLPerf Training workload

- Coming Soon!


MLPerf Storage Benchmark

- Using the tool “dlio” from Argonne Leadership Computing Facility (ALCF)
- DLIO:
 - Uses the same data loaders as the real workload
 - Implements a sleep in the execution loop for each batch
 - Sleep time is computed from running the real workload
 - A sleep time and batch size effectively defines an accelerator

Characteristics of Unet3D workload


- **Myths:**

- Large Block Random with high Queue Depth represents AI workloads
- We can divide the throughput requirement of a workload by the max throughput of our storage to get number of supported devices / servers



Queue Depth

- Queue Depth is low and frequently IOs are inserted into an empty queue.
 - Plot on the right



System Configuration

CPU - AMD EPYC 74F3 24-Core Processor

Memory – Micron 32GB DDR4 (256GB)

OS - AlmaLinux 9.2

Benchmark – MLPerf™ Storage Suite v0.5 (Unet3D)

Drive – Micron 9400 Pro (7.68TB)

I/O Throughput




Flash Memory Summit

- For 1, V100-32gb Emulated Accelerator with DLIO – the Micron 9400 drive throughput,

- Max 1542.5 MiB/s
- Mean 562.95 MiB/s


1 Acc



- For 15, V100-32gb Emulated Accelerators with DLIO – the Micron 9400 drive throughput,

- Max 6560.39 MiB/s
- Mean 5742.08 MiB/s

15 Acc




Queue Depth Vs Time



Flash Memory Summit


- For 1 emulated Accelerator - the drive never goes beyond 10QD and doesn't fully utilize NVMe SSD.

1 Acc




- For 15 emulated Accelerators, the drive hits around 120QD and NVMe SSD been fully utilized for max throughput.

15 Acc



% of Read at Queue Depth

1 Accelerator



15 Accelerators

