TECHNOLOGY WHITE PAPER

NXON AI FACTORY PTY. LTD

NXON Technical Summary: High-Performance GPU Training Storage Deployment

RoCEv2-Optimized Architecture for Large-Scale AI Training Workloads

Table of Contents

1. USE CASE OVERVIEW — LARGE-SCALE AI TRAINING WITH H200 GPUs

NXON supported a customer environment running large-scale AI training workloads using NVIDIA H200 GPUs. These workloads generated:

- High-frequency checkpointing
- Heavy metadata activity
- Mixed large-sequential and small-random I/O
- Sustained parallel access from over 100 GPU clients
-

The previous NFS-based storage system showed limitations such as:

- Latency spikes during checkpoint events
- Centralized metadata bottlenecks
- Reduced GPU utilization due to I/O stalls
- Limited scalability beyond certain client counts

2. STORAGE SOLUTION DEPLOYED BY NXON

NXON deployed a storage environment incorporating a WEKA® NeuralMesh file system as the high-performance Tier-1 NVMe storage layer.

2.1 Architectural Characteristics

- Single Tier-1 high-performance NVMe storage architecture
- Distributed, parallel file system with no centralized metadata bottlenecks
- User-space client providing kernel-bypass I/O
- Unified POSIX namespace for all datasets and checkpoints
- Approximately 4.7 PB usable capacity deployed in under 2 racks
- Storage network designed for 2 × 400 GbE uplinks per WEKA backend node
- RoCEv2 enabled on NXON switches for lossless Ethernet transport
- No additional tuning or configuration was required on the WEKA platform to support RoCEv2

2.2 Advantages Observed by NXON

- Low and stable latency during metadata-intensive bursts
- High bandwidth per GPU client without manual tuning
- Even distribution of I/O across backend nodes
- Ability to scale performance and capacity without downtime

3. PERFORMANCE OBSERVATIONS (MEASURED BY NXON)

3.1 Aggregate Cluster Throughput

- ~3 TB/s reads
- ~1.5 TB/s writes

3.2 Per-Client Throughput

- Up to ~40 GB/s from a dual-200 GbE GPU server
- Approximately 80% of theoretical link efficiency

3.3 Latency and IOPS Behavior

- Consistently below 0.5 ms
- Sustained millions of IOPS with no tail-latency spikes

3.4 Scalability

- Successfully scaled to over 140 GPU clients
- No observed loss in per-client throughput
- Backend automatically balanced all connections

All performance metrics reflect NXON's findings in this specific project.

4. BENEFITS VS. PREVIOUS PLATFORM

4.1 Checkpoint Performance

- Metadata bursts handled more efficiently
- Predictable checkpoint duration
- Reduced GPU idle time during save events
- Enabled more frequent checkpointing without affecting training cadence

4.2 GPU Utilization Improvements

- Storage delivered sufficient throughput to keep GPUs consistently fed
- Reduced wait times improved total training throughput

4.3 Operational Improvements

- Simple client installation
- No manual filesystem tuning required
- Non-disruptive scaling and backend upgrades
- Zero downtime during capacity expansion

5. NXON ROLE AND CONTRIBUTION

5.1 RoCEv2 Network Engineering

NXON designed and optimized the lossless RDMA fabric:

- Priority Flow Control (PFC) configuration
- ECN tuning
- Queue and buffer profile optimization
- Full validation of RDMA performance under load

Note: RoCEv2 functionality was enabled entirely at the switching layer. The WEKA platform required no additional configurations to operate over RoCEv2.

5.2 Rapid Deployment of Storage Environment

- Completed full deployment in under 4 days
- Expected performance observed during Day-1 validation

5.3 Data Migration

- Migration from the previous storage environment coordinated by NXON
- Parallel POSIX-based data copy approach
- Minimal operational disruption

5.4 Optimization and Tuning

- NIC CPU pinning and affinity adjustments
- Client-side I/O path optimization
- Guidance provided for future scale-out

6. INFRASTRUCTURE EFFICIENCY OBSERVATIONS

6.1 Performance per kW

NXON observed:

- High throughput relative to power consumption
- Improved energy efficiency versus the prior platform

6.2 Compact Datacenter Footprint

- ~4.7 PB Tier-1 NVMe storage
- <2 racks total footprint
- Reduced power, cooling, and space requirements

7. CONCLUSION — NXON'S ASSESSMENT

Based on NXON's deployment and validation, this storage architecture delivered:

- High throughput across a large number of GPU clients
- Stable sub-millisecond latency, even under metadata bursts
- Predictable scaling behavior at cluster sizes exceeding 140 GPUs
- Strong performance-per-kW efficiency
- Minimal datacenter footprint
- Simplified day-to-day operations
- Immediate improvements in AI training throughput and GPU utilization

This design provides a strong foundation for organizations operating modern, large-scale GPU training environments.

8. DISCLAIMER

This document is prepared by NXON based on its work in a specific customer deployment. All performance observations and conclusions reflect NXON's testing and operational experience in that environment alone. This document is not authored by, endorsed by, or intended to represent WEKA or the WEKA platform as a whole. Actual performance and behavior may vary depending on workload, configuration, and deployment conditions.