
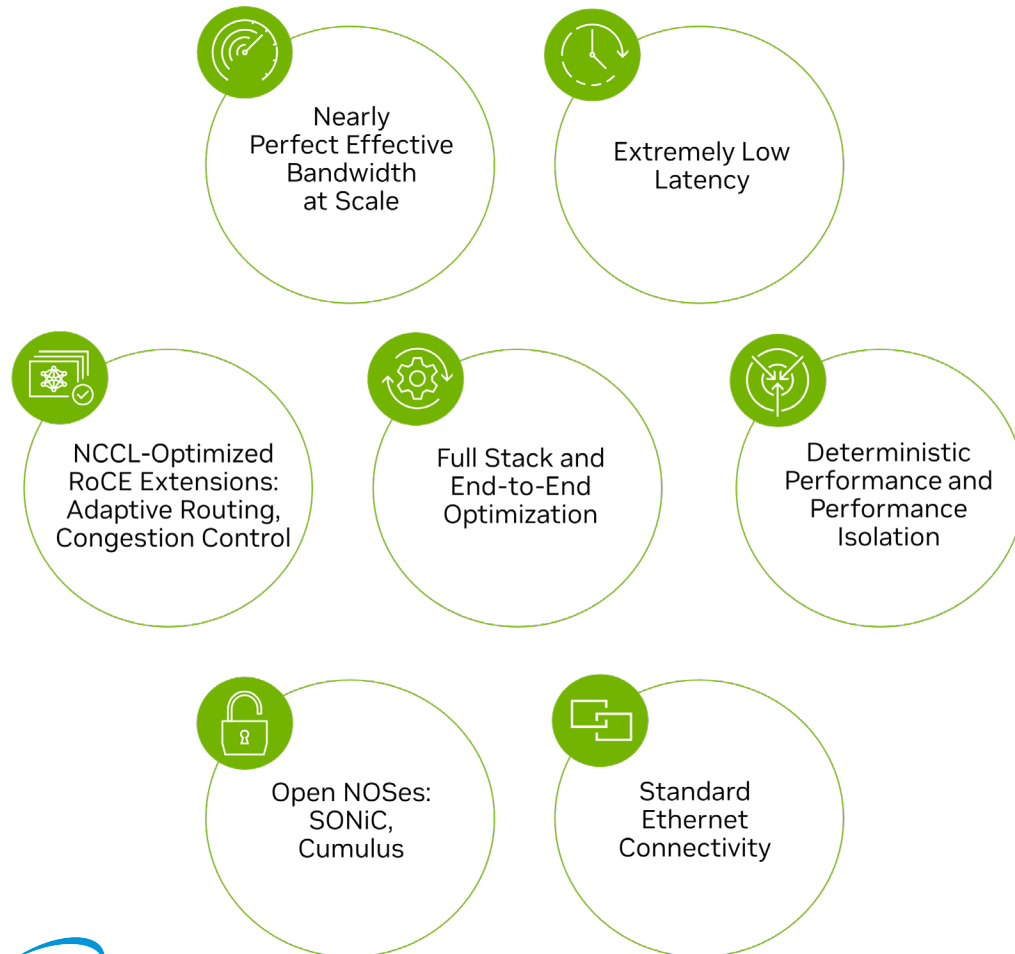


# NVIDIA Spectrum-X Technology for AI Storage Fabrics

Reggie Reynolds

# NVIDIA Spectrum-X: World's First Ethernet Platform for AI

Combining Specialized High-Performance Architecture with Standard Ethernet Connectivity




# AI Storage Optimization Use Cases

Foundations to Build and Run Your AI Storage

## Storage Matters for High Performance AI Factories


### Multi-modal Data Fetching



### Checkpointing



### Distributed Inference



Storage network performance is key to every stage of the AI model lifecycle

Data Ingest • Training • Fine Tuning • Inference

# NVIDIA Spectrum-X Offers High Performance for Storage

Storage Ecosystem Partners Standardize on Spectrum-X Ethernet Technology

- Spectrum-X for storage workloads delivers up to:

48%

Higher storage  
bandwidth per GPU

1.2X

Higher performance in  
noisy neighbor scenarios

- Spectrum-X Ethernet is being adopted by key players in the high-performance storage ecosystem:



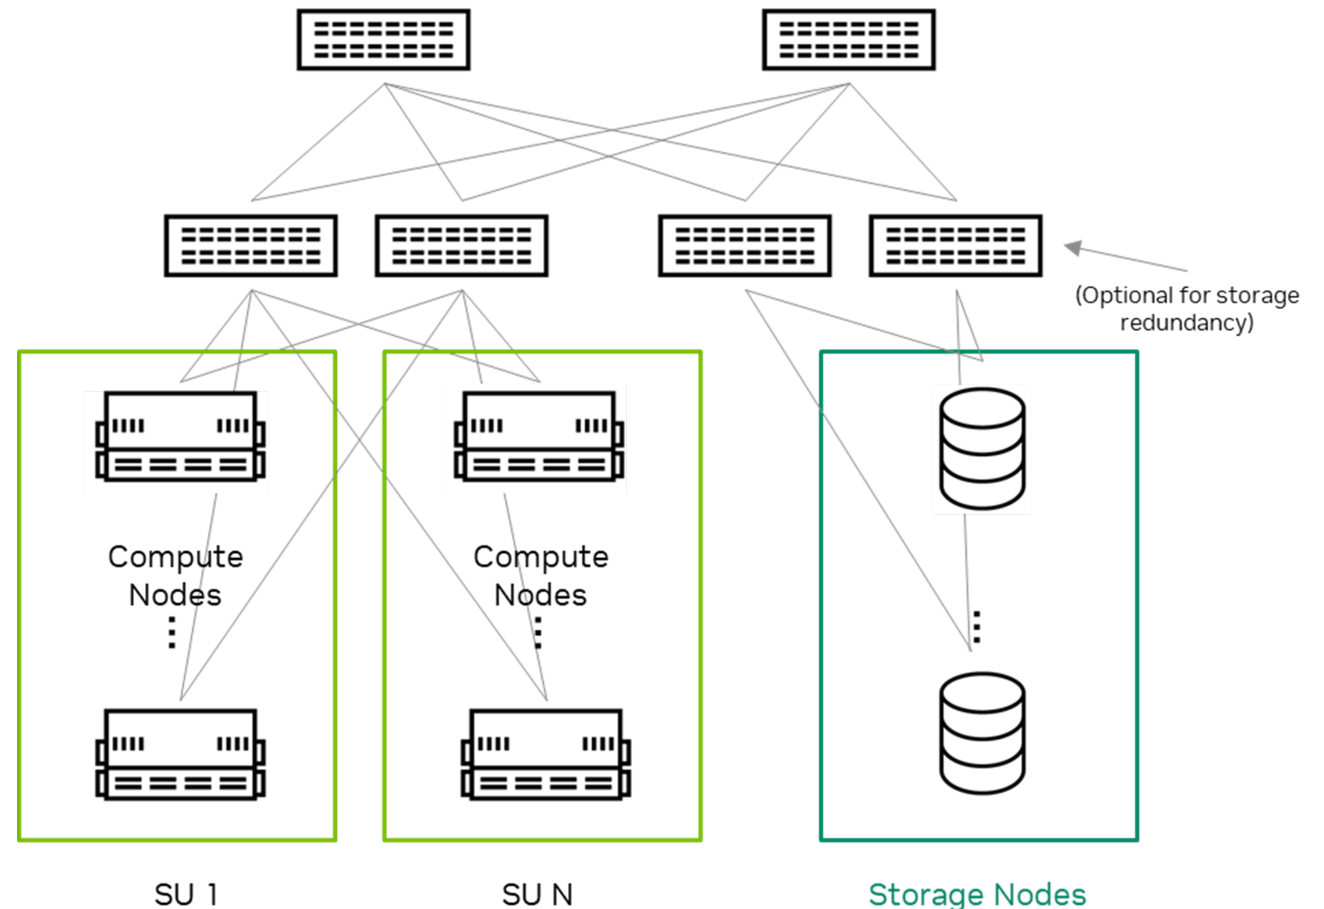
# Improved Performance for the Data Storage Fabric

Leverage Adaptive Routing for Higher Effective Bandwidth

Storage operations such as checkpointing, data ingestion, model storage, retrieval, etc.  
All generate elephant flows



Load balance with Spectrum-X  
adaptive routing for  
50% higher effective bandwidth with  
perfect collision prevention





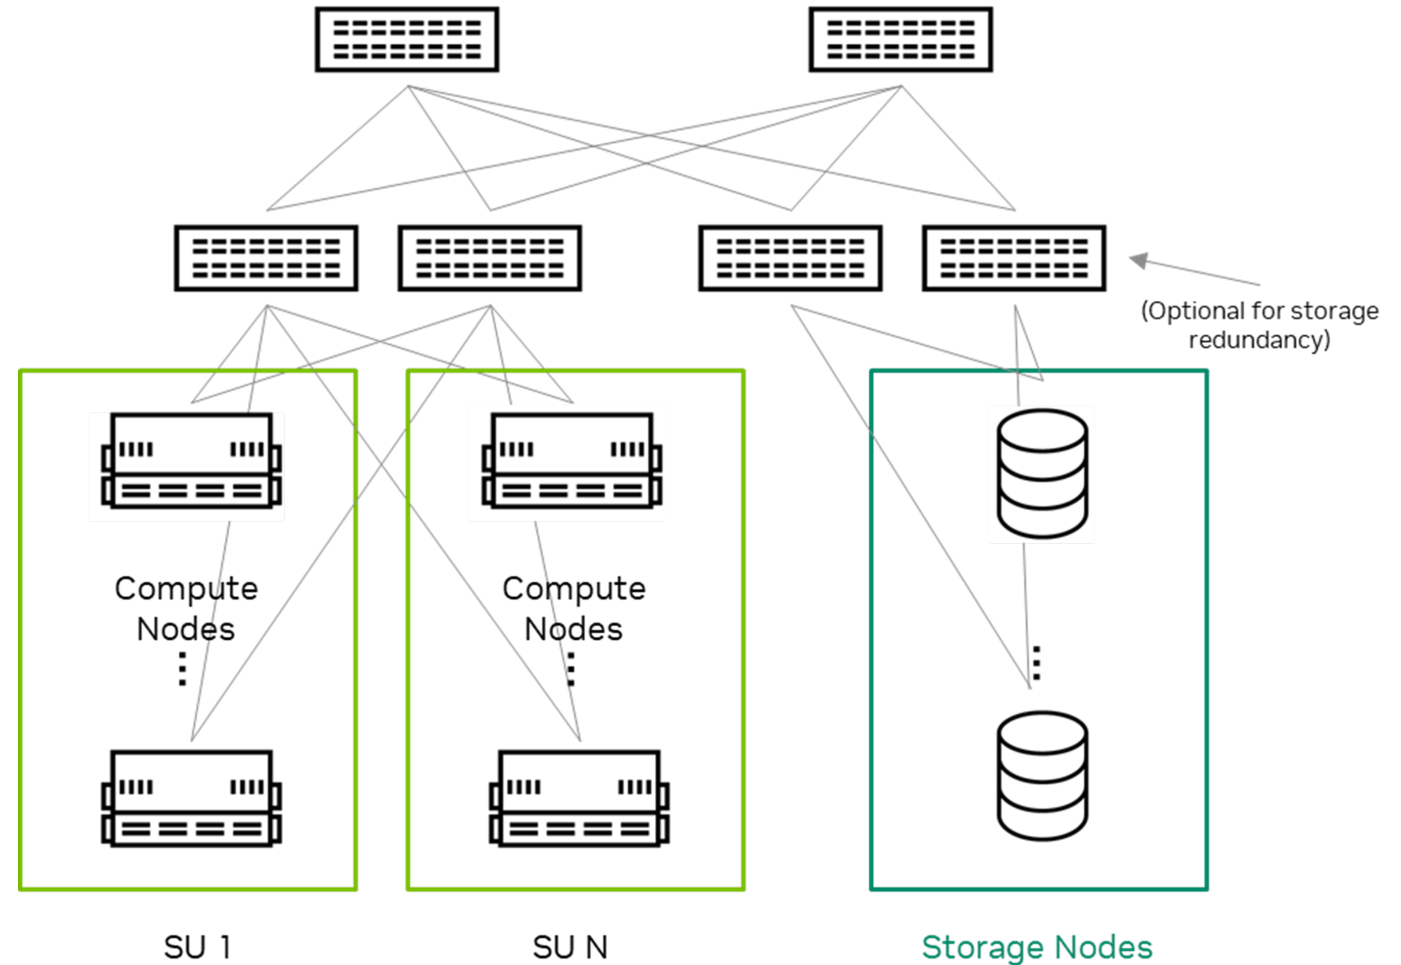
# Improved Performance for the Data Storage Fabric

Leverage Congestion Control for Performance Isolation

Multiple clients read or write to a single storage server creates incast congestion





Spectrum-X congestion control avoids hotspots and provide performance isolation



# Resilient Adaptive Routing Performance

Link Failures on Traditional Ethernet Lead to Disproportionate Drop in AI Performance



Spectrum-X utilizes global adaptive routing to rebalance flows and avoid failed paths

# Enhanced Visibility for AI Storage Fabrics

## NetQ Accelerated Telemetry



Topology Validations – Reduce **Time To AI**



AI Performance – Maximize Storage Utilization



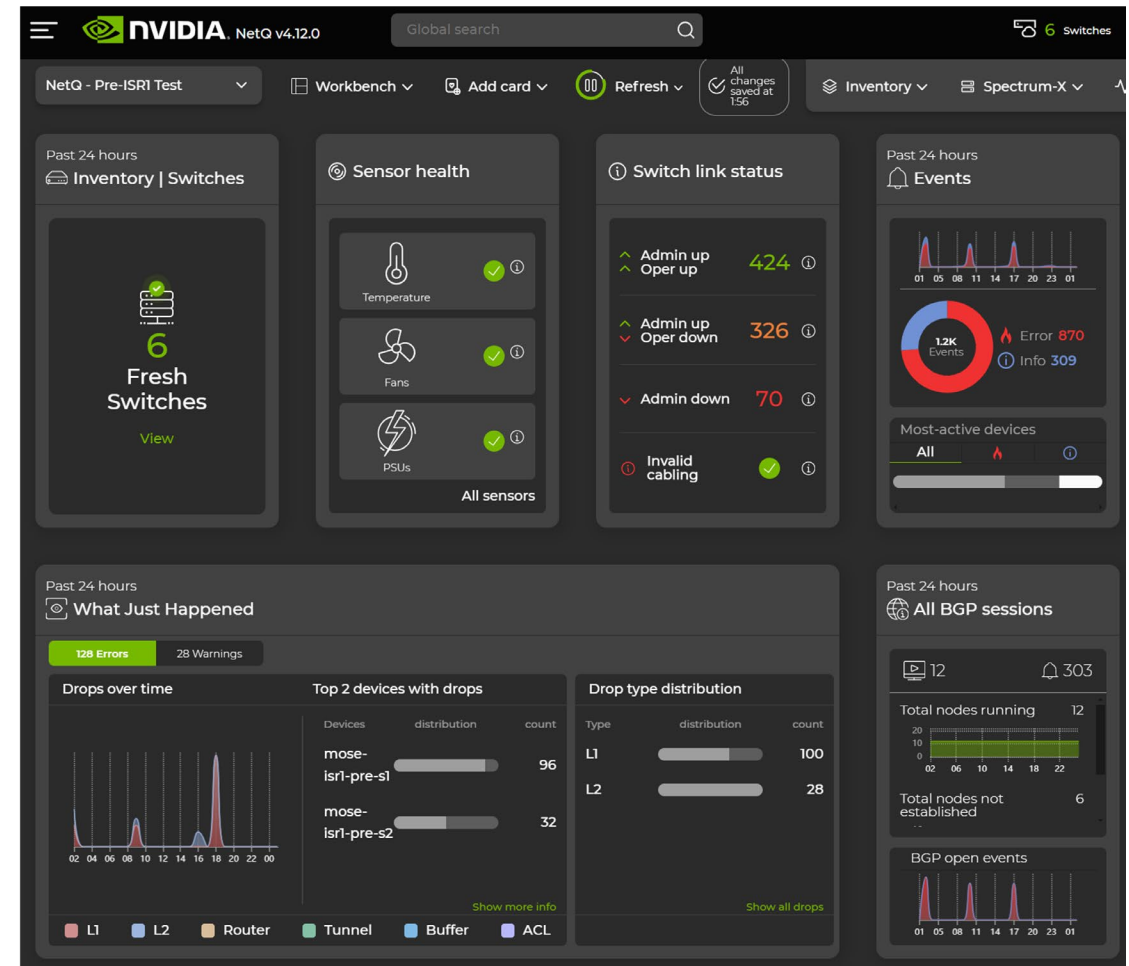
Actionable Visibility – Reduce Time To Identify Issues



Link Health Monitoring – Detect Links Causing Issues



Hardware Monitoring – Identify Faulty Hardware





# Performance Benchmark Results and Partner Solutions

# Spectrum-X Accelerates Israel-1 Storage


Storage IO performance at scale

- FIO Write Benchmark
  - Sends data from GPU nodes to Storage
  - Like saving Checkpoints during training
- FIO Read Benchmark
  - Sends data from Storage to GPU nodes
  - Like Inference or restoring checkpoints during training



**NVIDIA Israel-1 AI Supercomputer**

Validated / Optimized / Quantified



**Israel-1 Storage Test Results**

300 GPUs across 4 SUs

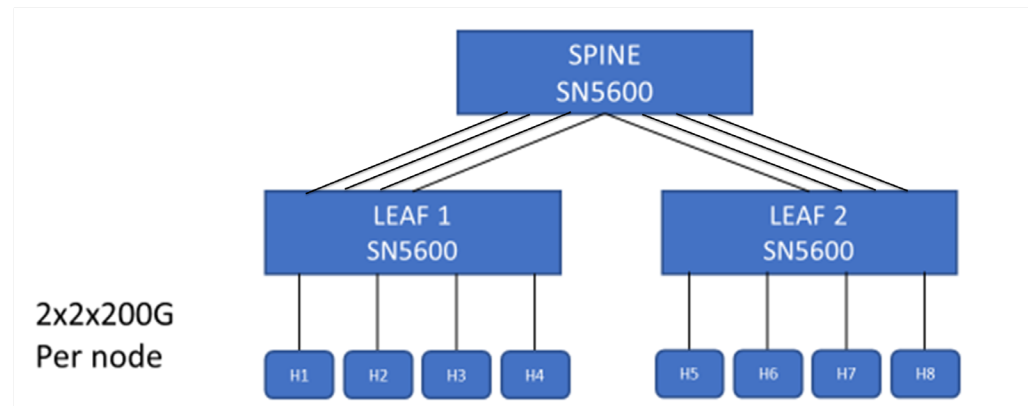
# Spectrum-X Accelerates GPFS

Up to 55% increased performance


- GPFS results:
  - Spectrum-X improved storage performance up to 55%
  - Bandwidth per port varied significantly with ROCEv2
  - Bandwidth per port was very consistent with Spectrum-X

- GPFS test setup:
  - 3 switches in a simplified leaf/spine topology
  - 8 servers in various storage client/target modes
  - 200GB and 1000GB files

- GPFS test topology:



**GPFS Write Performance**



# Benchmarks from NVIDIA Storage Partners

Summary of results from VAST, DDN, and Weka

- Significantly higher throughput
- Lower Latency
- Noisy Neighbor protection in the face of multiple workloads running simultaneously
- Higher Scalability

# Resources

# Spectrum-X for Storage Resources

August 2025

## NVIDIA Collateral:

- [White paper](#)
- [Technical blog](#)
- [FAQ/positioning doc](#)
- [Spectrum-X Product Webpage](#)
- [NVShare](#)

## Partner Collateral:

- [DDN white paper](#)
- [DDN blog](#)
- [VAST Data white paper](#)
- [VAST Data blog](#)
- [Weka blog](#)



[Enhancing AI Storage Fabrics with NVIDIA Spectrum-X](#)



# Thank You!