signal65

Optimizing AI Workloads with Dell Ethernet Infrastructure

AUTHOR

Brian Martin
AI and Data Center Lead | Signal65

IN PARTNERSHIP WITH

DELL Technologies

JUNE 2025

Executive Summary

AI Networking Challenges

The rapid adoption of AI workloads throughout the enterprise has created unprecedented demands on networking infrastructure, requiring solutions that can efficiently handle the unique communication patterns of distributed training, real-time inference, retrieval-augmented generation (RAG), and parallel agentic operations. Organizations deploying on-premises AI clusters must carefully balance performance, scalability, and total cost of ownership while ensuring their infrastructure can adapt to evolving AI requirements. This analysis examines how advanced Ethernet-based networking delivers compelling advantages for AI workloads through superior efficiency, reduced latency, and operational simplicity.

Dell Solutions

Dell 400G Network Interface Cards, coupled with Dell PowerSwitch Z9864F-ON 800G switches deliver competitive advantages that directly address the critical requirements of modern AI infrastructure.

Superior Network Efficiency

Broadcom BCM57608 achieves 98 GB/s efficiency per core versus 71 GB/s for competing solutions, while delivering 49 GB/s in NCCL All-to-All operations compared to 39 GB/s for alternatives.

Enhanced AI Scalability

Consistent performance scaling from 4-node to 8-node configurations with maintained bandwidth and latency characteristics across distributed training workloads.

Compelling TCO Advantages

Open Ethernet ecosystem reduces vendor lock-in while higher port efficiency decreases infrastructure requirements, leading to lower power, cooling, and management costs.

These performance and economic advantages position Dell networking infrastructure as an optimal foundation for organizations seeking to maximize both computational efficiency and operational flexibility in their AI deployments.

Highlights and Value Proposition

Organizations adopting this Dell infrastructure realize significant strategic advantages through improved performance efficiency, reduced operational complexity, and compelling total cost of ownership benefits. The open Ethernet ecosystem prevents vendor lock-in while providing access to a broad range of compatible components and management tools, enabling more flexible procurement and deployment strategies. Advanced features like hardware-based congestion control, fast link failover capabilities, and comprehensive telemetry ensure reliable operation under demanding AI workloads while providing the visibility needed for optimal resource utilization. This infrastructure positions enterprises for continued AI evolution with substantial headroom for growth and compatibility with emerging protocols, making it an optimal foundation for organizations seeking to maximize both computational efficiency and operational flexibility in their AI deployments.

The Imperative for Advanced Networking

The artificial intelligence landscape has evolved dramatically with the emergence of large language models (LLMs) and generative AI, fundamentally transforming networking requirements for enterprise infrastructure. AI workloads, encompassing training, fine-tuning, and inferencing, exhibit distinct and demanding network characteristics. These workloads are characterized by the movement of massive data volumes, acute sensitivity to latency, and an absolute requirement for lossless communication to maintain computational integrity and efficiency.

Unique Demands of Modern AI Workloads

AI model training, the most computationally and network-intensive phase, involves processing colossal datasets that can span terabytes or petabytes. During distributed training, frequent synchronization of parameters (gradients) among GPUs is essential. This generates intense, often bursty, all-to-all or all-reduce communication patterns, requiring high bandwidth and low, predictable latency to prevent GPU starvation and minimize overall job completion time. The sheer scale of contemporary AI models, with parameter counts escalating from billions to trillions, directly correlates with the amount of data that must traverse the network.

Fine-tuning, the process of adapting pre-trained foundation models to specific tasks or datasets, shares many of the network demands of training from scratch, often with smaller datasets and shorter durations. Federated fine-tuning, an emerging paradigm for training on decentralized data while preserving privacy, introduces additional complexities related to communication overhead and efficiency across distributed nodes. Efficiently aggregating updates from numerous participating clients without incurring prohibitive network delays is a key challenge that high-performance networking aims to address.

The growth in AI model parameters influences the choice of distributed training strategies, such as data parallelism, tensor parallelism, and pipeline parallelism. Data parallelism heavily relies on efficient all-reduce operations for gradient aggregation, while pipeline parallelism involves point-to-point transfers of activations and gradients between stages. Tensor parallelism, often used for very large layers, can generate intense local and inter-node communication. The network infrastructure must be adept at handling these varied, and often concurrent communication demands without becoming a bottleneck.

Inferencing, the application of trained models to new data for generating predictions, whether direct, RAG-based, agentic, or a combination, presents a bifurcated set of network requirements. Real-time inference, crucial for applications like fraud detection or autonomous systems, mandates the lowest possible latency to ensure immediate responses. Even minor delays in the network can significantly degrade user experience or operational efficacy. Batch inference, conversely, prioritizes overall throughput and computational efficiency for processing large volumes of data where immediacy is less critical. Both mode benefit from high-speed data ingestion and result dissemination.

Evolution of Ethernet for AI

The relentless pace of AI development has catalyzed the rapid evolution of Ethernet speeds within the data center. Traditional 100G Ethernet, once considered high-performance, quickly became a bottleneck for early AI clusters. This spurred the migration to 400G Ethernet, which provided a fourfold increase in bandwidth and became a foundational technology for many AI deployments. However, the exponential growth in AI model complexity and data set sizes means that even 400G architectures are now being pushed to their limits, particularly for large-scale training operations. This has precipitated a further transition towards 800G Ethernet, expected to constitute the majority of switch ports in AI back-end networks in 2025.

Ethernet, augmented with technologies like RoCEv2, has emerged as a powerful, open, and economically attractive solution for AI fabrics. Ethernet's key advantages lie in its vast and open ecosystem, the availability of merchant silicon, a broad base of skilled professionals, and generally lower TCO. While InfiniBand may offer marginally lower latency in idealized scenarios, optimized Ethernet configurations are increasingly demonstrating performance parity characteristics for a wide array of AI workloads. Additionally, Ethernet boasts significantly faster link failover times, critical for resilience in large, long-running AI jobs.

Test Configuration

Testing was conducted on a 64 GPU NVIDIA H200 cluster with the following configuration:

- 8 Dell PowerEdge XE9680 servers with 8 H200 GPUs each
- 8 400G NICs (either Broadcom BCM57608 or CX-7) per Dell PowerEdge XE9680 server
- TH5 Ethernet ToR & Spine switches with 1:1 full bandwidth interconnect
- Ubuntu 22.04.5 LTS, NCCL 2.24.3-cuda12.6, Broadcom BCM57608 233.1 software, and CX-7 MLNX_OFED_LINUX-24.10-1.1.4.0

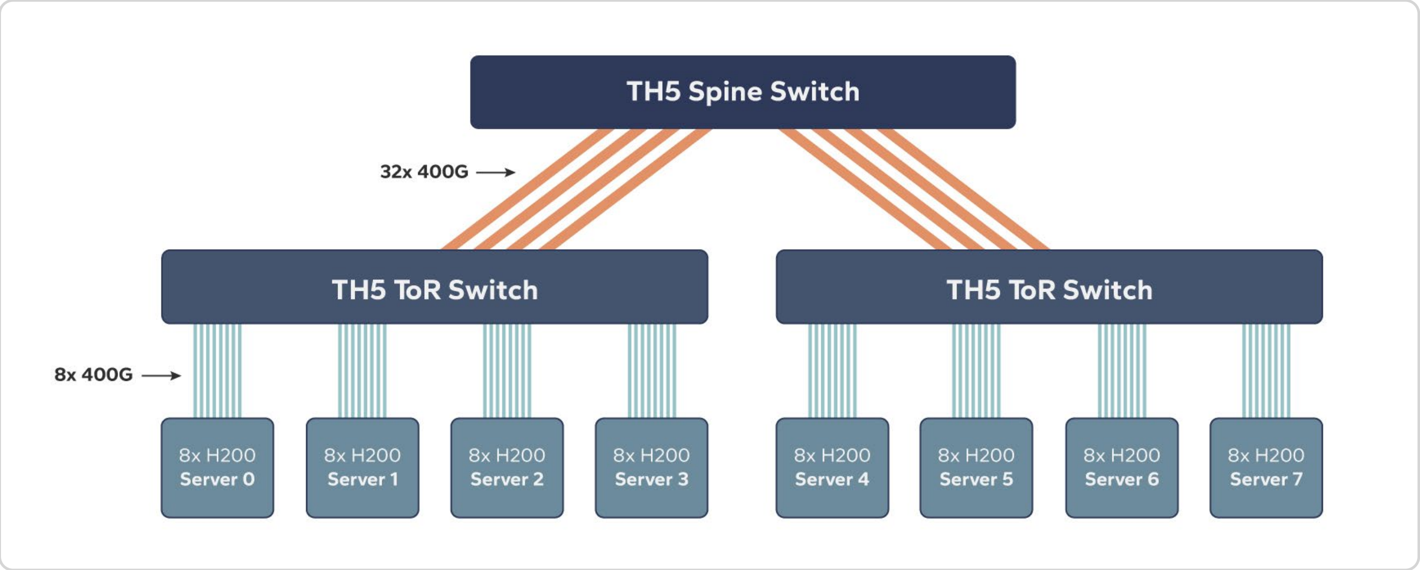

Figure 1: *Test Configuration*

Detailed Performance Analysis

NCCL Collective Operations

NCCL AllReduce Performance

AllReduce operations are fundamental to synchronous distributed training, where gradient updates must be aggregated across all nodes before model parameters can be updated. Both solutions show similar scaling behavior as collective size increases from 4MB to 4GB. At 128MB message sizes—typical for gradient tensors in large transformer models—Thor2 achieves 237 GB/s compared to CX-7's 219 GB/s. In practical terms, this 8.2% difference translates to reduced synchronization overhead during backpropagation phases of training runs that can span weeks or months. For organizations training foundation models like GPT or LLaMA variants, this efficiency gain directly impacts both time-to-convergence and compute resource utilization, particularly when training iterations number in the millions across distributed clusters.

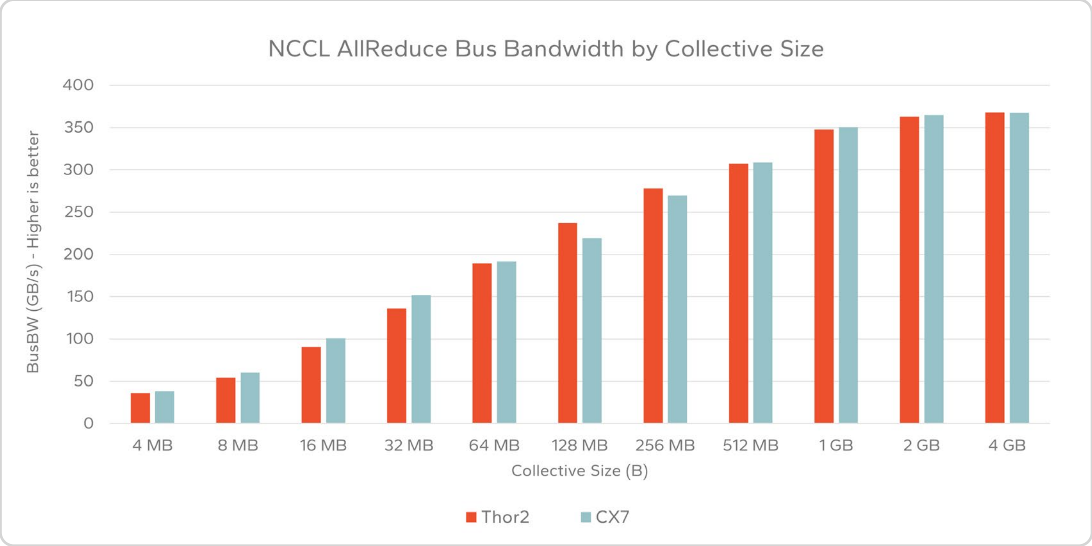

Figure 2: *NCCL AllReduce Bus Bandwidth by Collective Size*

NCCL All-to-All Performance

All-to-All communication patterns are critical for advanced distributed training strategies including model parallelism, where different layers of large neural networks are distributed across multiple GPUs. Thor2's 49 GB/s performance at 128MB versus CX-7's 39 GB/s represents a 25.6% advantage that directly affects training throughput for models that exceed single-GPU memory capacity. This communication pattern is particularly relevant for training and tuning large language models with billions of parameters, where model sharding becomes necessary. The performance difference translates to measurable reductions in training wall-clock time and improved GPU utilization rates, significantly impacting the economics of large-scale AI development.

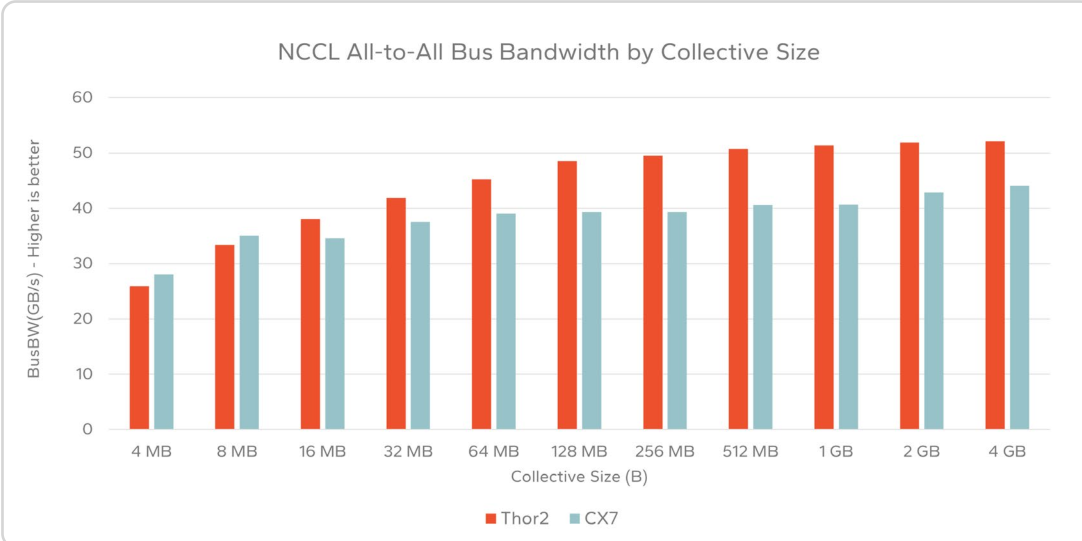

Figure 3: *NCCL All-to-All Bus Bandwidth by Collective Size*

NCCL AllGather Performance

AllGather operations facilitate the distribution of model parameters and intermediate results across distributed training processes, commonly used in data parallel training scenarios. Thor2 and CX-7 deliver equivalent performance at 128MB message sizes. This parity means that training pipelines relying heavily on parameter broadcasting—such as federated learning implementations or distributed inference serving—will experience similar network-bound performance characteristics regardless of NIC choice.

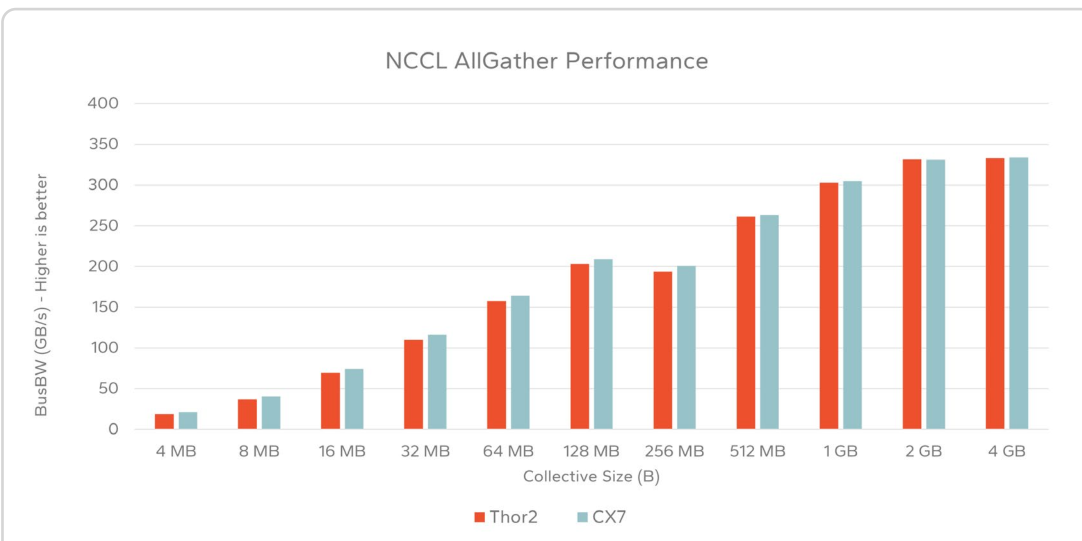

Figure 4: *AllGather Performance by Collective Size*

NCCL ReduceScatter Performance

ReduceScatter operations combine reduction and scattering phases, commonly used in optimized distributed training algorithms like ZeRO (Zero Redundancy Optimizer) that partition optimizer states across nodes. CX-7's 233 GB/s versus Thor2's 229 GB/s at 128MB represents a 1.7% difference, not likely to impact training job completion times. For memory-efficient training of very large models where optimizer state partitioning is essential, such as training 175B+ parameter models, both solutions provide adequate bandwidth to prevent communication from becoming the limiting factor in training throughput.

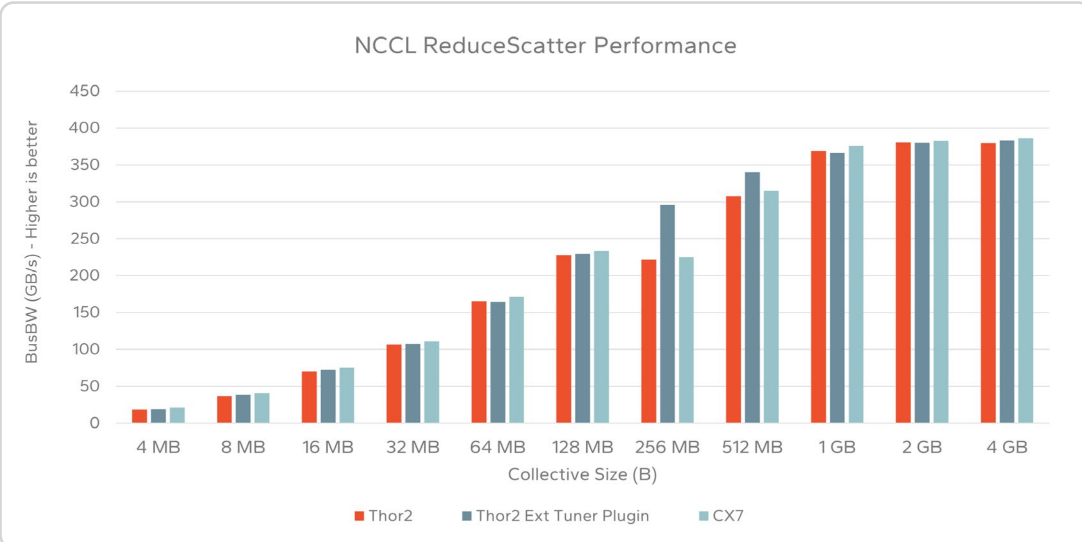

Figure 5: *NCCL ReduceScatter Performance by Collective Size*

Large Language Model Training

Llama2 70B Fine-Tuning Scaling

Fine-tuning large language models represents a common production workload where pre-trained models are adapted for specific domains or tasks. Both Thor2 and CX-7 achieve identical scaling to approximately 12 samples per second across 64 H200 GPUs, demonstrating that neither solution introduces network bottlenecks in realistic training scenarios. This performance level translates to processing thousands of training examples per hour, enabling practical fine-tuning timelines measured in days rather than weeks for large datasets.

- Both solutions scale similarly from 1 to 8 nodes (8 to 64 GPUs)
- Performance tracks closely with NVIDIA's published results across node counts
- At full scale (8 nodes, 64 GPUs), both solutions achieve approximately 12 samples per second

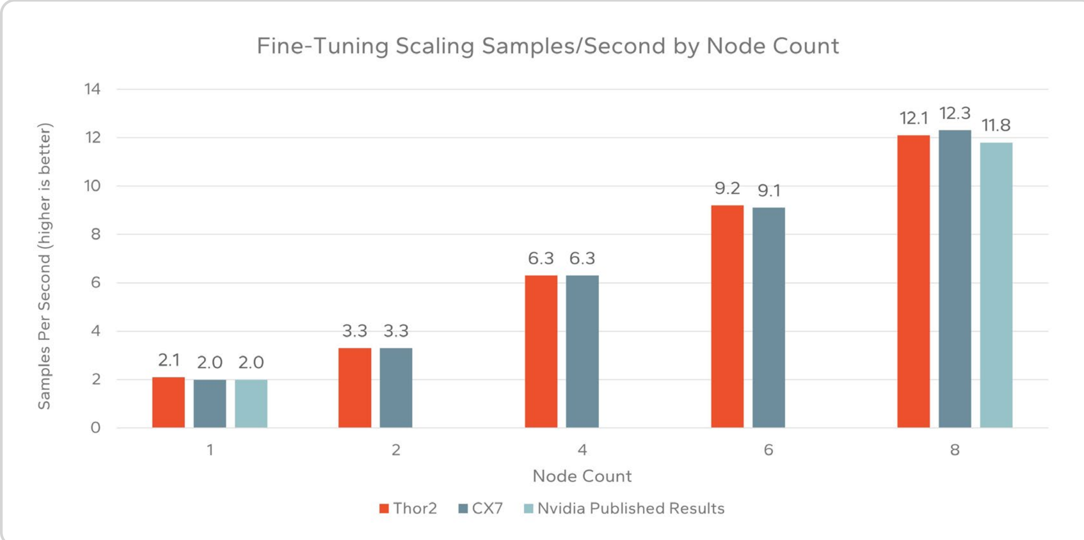

Figure 6: *Fine-Tuning Scaling Samples/Second by Node Count*

Additional Benchmarks

Broadcom BCM57608 also shows competitive performance in additional benchmarks. In traditional machine learning tasks, Thor2 demonstrates marginal advantages in BERT-large training (304s vs 318s) and ResNet-50 training (129s vs 128s), indicating consistent performance for computer vision and natural language processing workflows commonly deployed in production environments. These results are particularly relevant for organizations running mixed workloads that combine legacy ML models with modern transformer architectures, where network infrastructure must support both established and emerging training patterns without performance degradation.

Thor2's performance across workloads demonstrates its effectiveness in supporting the full spectrum of distributed training scenarios commonly encountered in production AI environments. The slight advantage in AllReduce operations (237 GB/s vs 219 GB/s) and stronger performance in All-to-All communication (49 GB/s vs 39 GB/s) translate to measurable improvements in training throughput for complex model architectures that rely heavily on inter-node communication. For organizations operating large-scale training clusters where marginal efficiency gains compound over extended training runs, these performance differences represent tangible reductions in compute costs and time-to-market for AI model development initiatives.

Category	Benchmark	Broadcom Thor2	NVIDIA CX-7
TCP	Bidirectional Throughput	800Gb/s	800Gb/s
	Efficiency (Gb/s per core)	98Gb/s	71Gb/s
DPDK	Packet Rate	230Mpps	268Mpps
	Throughput @400B packets	800Gb/s	800Gb/s
NVIDIA GPU	**NCCL All Reduce** busBW, 8 nodes, 64 H200 GPUs, 128MB	237GB/s	219GB/s
	NCCL All to All busBW, 8 nodes, 64 H200 GPUs, 128MB	49GB/s	39GB/s
	NCCL All Gather busBW, 8 nodes, 64 H200 GPUs, 128MB	203GB/s	209GB/s
	NCCL Reduce Scatter busBW, 8 nodes, 64 H200 GPUs, 128MB	229GB/s	233GB/s
	BERT-large - MLPerf v3.1 (Training time - lower is better)	304s	318s
	ResNet - 50 v1.5 - MLPerf v3.1 (Training time - lower is better)	129s	128s
CPU/HPC	**OSU All to All** latency, 64 nodes, 16 PPN, 64KB message (blocking)	28ms	68ms
	OSU All Reduce latency, 64 nodes, 16 PPN, 64KB message (blocking)	179us	159us
	LAMMPS - Rhodo Scaled Benchmark (CPU/atom/steps - lower is better)	6.92E-09	7.16E-09
	LAMMPS - Chain Scaled Benchmark (CPU/atom/steps - lower is better)	3.51E-10	3.66E-10
	HPCG - (GFLOP/s - higher is better)	2,571	2,492

Figure 7: *Additional Performance Metrics*

Conclusion

Dell networking solutions based on Broadcom BCM57608 and Dell PowerSwitch Z9864F-ON, provide enterprise organizations a compelling foundation for scaling AI workloads without compromising performance or flexibility. The demonstrated performance advantages in critical communication patterns, particularly a 25% improvement in All-to-All operations and superior per-core efficiency, directly address the intensive networking demands of modern AI distributed training, fine-tuning, and inference workloads. These performance gains become increasingly significant as organizations deploy larger models and longer training runs, where network efficiency improvements translate to substantial reductions in both time-to-convergence and operational costs.

Beyond raw performance metrics, Ethernet-based approaches deliver strategic advantages that position organizations for long-term AI infrastructure success. The open ecosystem prevents vendor lock-in while providing access to a broad range of compatible components and evolving protocols, ensuring infrastructure investments remain viable as AI technologies continue to advance. Combined with faster link failover capabilities, comprehensive telemetry, and the operational simplicity inherent to Ethernet architectures, Dell's networking solutions enable organizations to focus on AI innovation rather than infrastructure complexity. For enterprises seeking to maximize both computational efficiency and operational flexibility in their AI deployments, Dell's proven networking infrastructure provides the performance headroom and adaptability required to support the next generation of AI workloads.

Important Information About this Report

CONTRIBUTORS

Brian Martin
AI and Data Center Lead | Signal65

Russ Fellows
VP, Labs | Signal65

PUBLISHER

Ryan Shrout
President and GM | Signal65

INQUIRIES
Contact us if you would like to discuss this report and Signal65 will respond promptly.

CITATIONS
This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "Signal65." Non-press and non-analysts must receive prior written permission by Signal65 for any citations.

LICENSING

DISCLOSURES
Signal65 provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

IN PARTNERSHIP WITH

DELL Technologies

ABOUT SIGNAL65
Signal65 is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

signal65

signal65

CONTACT INFORMATION
Signal65 I signal65.com