



Validating High-Speed Ethernet for Next-Gen AI Networking

Data Center Networks Must Transform to Support New AI Workloads

AI model scale has continued its exponential rise, with parameter counts expanding from billions to the trillion-parameter class in just a few years. Recent analyses¹ highlight the emergence of trillion-parameter models and the accelerating growth trajectory driven by scaling laws and data-compute expansion. This rapid escalation reflects industry-wide scaling dynamics, where increasing model size, dataset size, and compute consistently yield higher capability, reinforcing why today’s systems push into the hundreds-of-billions to multi-trillion parameter range. The growing complexity and size of AI applications is driving a need for increasing numbers of AI accelerators, which in turn dictate the type and scale of the networking fabric required to interconnect the accelerators with each other and compute and memory resources.

As AI applications move from small apps with hundreds of AI accelerators to large apps that require tens of thousands of AI accelerators, the networking architectures grow more complex. Up to hundreds of accelerators leveraging CXL, NVLink and/or PCIe can effectively manage connectivity needs. A cluster with thousands of accelerators requires moving to multiple racks to host the accelerators. Clusters of resources running in these racks must be interconnected with an AI leaf layer. As we grow to 10k+ accelerators, a data center scale architecture is required with a routable fabric and AI spine.

AI XPU Size	Server IO 10-100s of XPU	Rack Scale 1000s of XPUs	DC Scale 10K+ of XPU
AI Network Options	CXL NVLink PCIe	AI Leaf Ethernet or IB	AI Spine Ethernet or IB
Type of Applications	Small AI Apps	Moderate AI Apps	Large AI Apps
	<p>Scale-up design option</p> 		<p>Scale-out design option with a new AI network that is different from traditional network</p> 
<p>Source: AI Networks for AI Workloads – Advanced Research Report, Dell’Oro Group (July 2025)</p>			

With most AI growth in moderate to large AI apps, AI backend networks leveraging AI leaf and spine approaches are critical for accelerator interconnectivity. This represents a major transformation from traditional front-end data center networks connecting general purpose servers.

AI-driven network traffic continues to accelerate at multi-fold year-over-year growth rates, driven by the shift to trillion-parameter and Mixture of Experts (MoE) based models that demand massive cross-GPU communication. Recent analyses show that communication overhead can exceed

¹Source: NVIDIA How Scaling Laws Drive Smarter, More Powerful AI (2025)

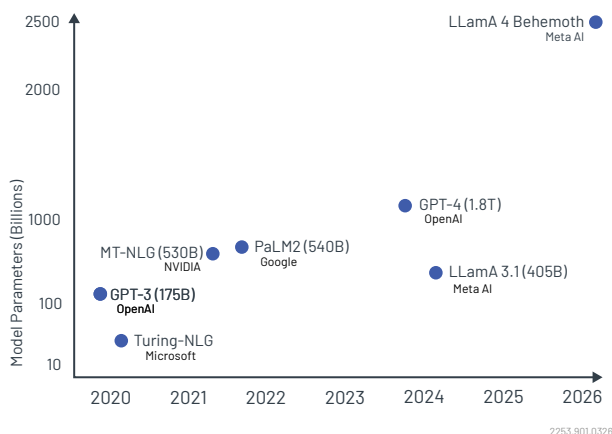
50% of total training time in large-scale clusters, making the network, not compute, the dominant bottleneck in modern AI infrastructures.

With next-generation distributed training architectures already pushing the limits of today’s 200/400/800Gbps fabrics, data centers are rapidly moving toward 800G and 1.6T Ethernet interconnects to sustain scalability and maintain accelerator utilization. Research on resilient AI fabrics highlights that emerging workloads now require high-bandwidth, low-latency fabrics (800G → 1.6T) as a baseline for efficient model parallelism and collective communication operations.

Given the long lead times for data-center construction and the pace of model growth, operators must future-proof network designs now. Without planning for 1.6T class bandwidth and congestion-resilient fabrics, new deployments risk becoming functionally obsolete at launch, unable to support the communication intensity of ever increasing AI workloads.

AI Networking Challenges

Bandwidth and Scale



High Data Volumes

AI workloads need high-speed transfers; however, traditional networks may have bottlenecks which limit transfers between compute and accelerator nodes.

Distributed Training

Growing workload complexity demands more nodes for AI model training. Traditional networks may not scale efficiently as models move to trillions of parameters and beyond.


Network Latency

Network Radix

Network topology/switch radix influences latency with higher radix providing shorter paths and lower latency for east-west traffic loads.


Congestion Hotspots

Network congestion slows data transfers between accelerators nodes or into the AI cluster itself, causing lags in training and inference.



Source: Mellanox/NVIDIA (2025), AI Training Cluster Network Bottlenecks: High-Performance Interconnect Solutions

Packet Latency



Packet Delivery

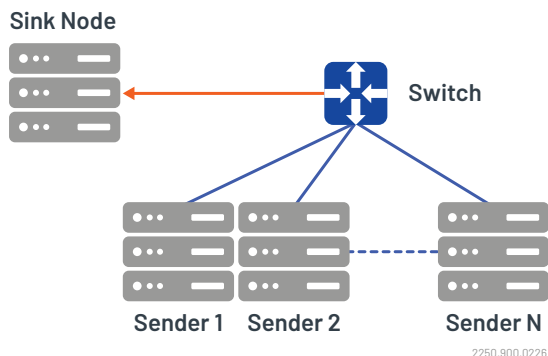
Out-of-order packets can increase latency significantly due to buffering or retransmission.

Long-tail Latency

Many AI operations such as model training cannot complete until all required packets are received. The highest 1% or 0.1% of packet latencies can cause extreme delays that disrupt critical AI operations.

Challenges of Using Traditional Ethernet for AI Networking


Running Ethernet as it exists today to serve AI/ML workloads is not well-suited to the unique requirements of such traffic.



Bandwidth Limitation

Incast Use Case

- Traffic spike in the network overwhelms the switch from Sender to Sink Node due to bandwidth bottleneck
- Traffic rate from senders gets reduced to avoid congestion




Multi-Path Packet Spraying

Head-of-Line (HOL) blocking

- Packets get blocked because the first packet in a queue is held up

Go back N routing algorithm

- Out of order packets will be dropped and all the subsequent frames will need to be retransmitted



Congestion Control

Priority Flow Control (PFC) storm and deadlock

- Packets paused in a loop across multiple switches block all packets from being transmitted

New Ethernet Technologies for Better AI Networking

The following technologies will help in alleviating network issues when serving AI workloads.

- Higher bandwidth networking
 - 400G, 800G, and 1.6T Ethernet on AI data center switch
 - 200G, 400G, and 800G interfaces on SmartNIC/GPU/DPU/TPU
 - 400G and 800G low-latency and low-power optics transceivers
- Advanced protocols for achieving a low latency, congestion-free network
 - Optimized RoCEv2 – RoCE selective repeat
 - Multi-pathing packet spraying
 - AI optimized congestion control
 - High scale – High-radix fabric enabling the building of massive flat networks
- Emerging Ultra Ethernet Transport (UET) standards
 - Re-engineering Ethernet to meet the stringent demands of large-scale AI training, introducing transport, congestion-control, and reliability mechanisms

Test and Validation Solutions

As AI networks adopt these advanced technologies, from 1.6T link speeds and multi-path packet spraying to new congestion control schemes and emerging UET protocols, the risk of subtle interoperability failures, performance regressions, and configuration errors grows significantly. Unlike traditional data center traffic, AI workloads are uniquely unforgiving: a single misconfigured congestion threshold or an out-of-order packet handling defect can stall an entire distributed training job across thousands of accelerators, wasting millions of dollars in compute time. Validating that each layer of the stack, optics, link integrity, RoCEv2 transport, congestion response, and collective communication patterns, performs correctly and at scale is no longer optional. It is a prerequisite for ensuring that the substantial capital invested in next-generation AI infrastructure delivers its intended return.

Ethernet Validation

- Up to 1.6T, multi-speed appliances
- Modular chassis based solutions
- Easy to use interface and test configuration
- License just the speeds and port blocks you need
- Full support, automation, per port reservation
- Breakouts supported
- L1 to L3 statistics and performance
- RoCEv2 hardware-based congestion control
- Bidirectional RoCEv2 traffic
- Comprehensive Collective Communications Library (CCL) AI workloads supported.



D2 1.6T 2U Appliance



B3 800G 2U Appliances



A1 400G 2U Appliance



A2 400G 2U Appliance



MX2 / MX3 Series
Multiprotocol Performance



FX2 / FX3 Series
Flexible Functionality



Contact Us: +1 844 GO VIAVI | (+1 844 468 4284). To reach the VIAVI office nearest you, visit viasolutions.com/contact

© 2026 VIAVI Solutions Inc. Product specifications and descriptions in this document are subject to change without notice. Patented as described at viasolutions.com/patents

highspeedethernet-nextgenai-wp-hse-nse-ae
30194943 901 0326

viasolutions.com