

January 16, 2025


This Former Spirent Business is Now Part of VIAVI

AI data centers are brimming with traffic, relying heavily on high-speed Ethernet and back-end fabrics that must perform with low latency and close to zero packet loss to avoid the expensive impact of idle GPUs. A new approach to testing can significantly improve performance and resiliency by stress-testing these fabrics with complex AI traffic patterns and congestion scenarios to identify issues.

High-Speed Ethernet (HSE) is Racing to Support AI

AI/ML applications and next-gen workloads are driving network traffic to unprecedented levels. Hyperscalers are pushing the boundaries of Ethernet to keep pace with these increasing bandwidth demands, rapidly transitioning to 800 Gbps for backend AI clusters and preparing for 1.6T in the near-future.

Migration to High-Speed Networks in AI Clusters (Back-End)



Source: Dell'Oro Group "AI Networks for AI Workloads" Report

Behind hyperscalers, Ethernet's high-speed adopters are service providers, large enterprises, and eventually smaller enterprises—sometimes years later. Lower speeds like 1G, 2.5G, and 5G are also in demand to support emerging applications in areas like access networks, automotive, industrial automation, and IoT.

800G and Higher Speeds are Fast-Tracking Complexity

Migrating to 800G to meet market needs is easier said than done. Upgrade paths have become more challenging compared to the adoption of previous generations of Ethernet.


So, what's changed?

- **More vendors.** Earlier Ethernet generations were dominated by one or two switch ecosystem vendors, so there was little need for interoperability testing. That changed with 400G and 800G as many vendors want to demonstrate they conform to standards, are interoperable, and can meet performance requirements. The challenge for a hyperscaler or provider is to ensure selected optics, cables/transceivers, and switches function together before being deployed in production networks.
- **Fast-paced cycles.** Adoption cycles of Ethernet speeds are getting shorter as demand and technology innovations accelerate. Chipset vendors and testing can't wait for standards to be ratified, but must evolve as the industry evolves, supporting early specifications and then IEEE standards.
- **Rapidly changing tech.** Fundamental—and complex—technology innovations, such as new forward error corrections, increased capacity per electrical lane, and core signaling are being introduced to support bandwidth and latency needs. All of these must be tested.
- **Demanding real-world conditions.** Systems may work properly under normal conditions but fail with real-world traffic. Before being placed in the production network, a system's ability to perform reliably under scale and stress needs to be validated.

To ensure the successful deployment of these rapidly evolving technologies, it's essential to thoroughly validate and test them. That includes rigorous testing solutions from chip, transceiver, and cable vendors, as well as network equipment manufacturers. Once these components are proven to function seamlessly together, deploying service providers and hyperscalers must prioritize vendor compliance, system interoperability, and performance under real-world traffic conditions. Issues must be identified and resolved in the lab, rather than within the live production network, where they can significantly increase operating costs and impact service delivery.

AI's Impact on Data Centers

AI data centers are crucial to the AI ecosystem and are experiencing explosive growth: massive GPU deployments, clusters quadrupling every two years, and traffic growing tenfold every two years. The demand for scale and capacity required to support AI is unprecedented. HSE is essential for AI data center back-end fabrics to achieve the required performance.



Adding to the complexity caused by growth, the massive volume of data communication between GPUs in an AI data center cannot be handled by a single GPU. Instead, processing is spread across several GPUs in parallel using high-performance computing Collective Communication Library (CCL) traffic patterns. Parallel processing is essential to meet AI performance expectations, but it is extremely sensitive to latency and packet loss. If all the distributed packets don't come back together in a timely manner or are out of sequence, retransmissions are required, which slow application response time.

As a result, AI data centers face significant challenges:

- **Scale.** AI data centers must support hundreds of servers and thousands of GPUs, which can cost billions of dollars to deploy.
- **Resiliency.** Complex traffic workloads put pressure on network resiliency.
- **Robustness.** Sensitivity to latency, congestion, and packet loss can cause expensive GPUs to sit idle waiting for retransmissions, as much as one-third of the time.

Embracing Essential New Testing Approaches for AI Data Centers

Test methods must evolve and innovate in alignment with AI and HSE's advancements.

Traditional tools provide basic AI testing but aren't suitable for validating back-end AI fabrics. Industry-standard RFC 2544 tests on an Ethernet fabric with a few GPU clusters may indicate that throughput, latency, jitter, and frame loss, are within acceptable limits. However, when the fabric is deployed in the production network and subjected to realistic, complex workloads and congestion at scale, numerous issues typically arise. These problems must be discovered in the lab, not in the production network, where they can impact service performance.

We have taken a new approach to AI testing by creating an at-scale production network scenario in the lab. Once RFC 2544 AI baseline tests are completed, it's essential to stress-test the fabric with CCL traffic patterns and congestion scenarios to identify and remove traffic bottlenecks. This advanced AI testing rigorously evaluates the fabric's readiness for deployment and ensures optimal GPU utilization.

With advanced AI testing, GPUs are emulated on test hardware to create realistic, complex AI workloads to test the fabric. This solution stress-tests the fabric in the lab with emulated CCL traffic patterns and congestion scenarios to ensure the AI data center infrastructure operates at peak performance and avoids idle time. Tests measure service performance, such as job completion time and packet performance. The inclusion of negative and impaired traffic allows for testing of recovery capabilities. These tests are reproducible and consistent, enabling benchmarking of different vendor fabrics.

Another benefit of this new era of AI testing is the reduction in test cost, complexity, and maintenance. Now, vendors and hyperscalers no longer need to build expensive labs with real servers and hard-to-obtain, pricey GPUs to test their Ethernet switch performance. These physical labs incur significant management and maintenance costs, and few people have the skills to create test methodologies for real GPUs.

Step into the New Era of AI and High-Speed Ethernet Testing

Advanced AI test solution emulates realistic AI workloads, going beyond basic tests. This reveals how well the switch fabric performs under complex traffic and scale, identifying potential bottlenecks that cause resources underutilization.

Test configuration is simplified with an intuitive wizard that easily sets up and generates complex traffic patterns. The solution has a small footprint, is easy to manage, and is backed by world-class test expertise.