

NVIDIA Spectrum-X

Purpose-built for Ethernet AI clouds.

The NVIDIA Spectrum™-X Ethernet platform is designed specifically to improve the performance and efficiency of Ethernet-based AI clouds. This breakthrough technology achieves 1.6X better AI networking performance, along with consistent, predictable performance in multi-tenant environments. Spectrum-X is built on network innovations powered by the tight coupling of the NVIDIA Spectrum-4 Ethernet switch and NVIDIA® BlueField®-3 SuperNIC. Spectrum-X network optimizations reduce runtimes of massive transformer-based generative AI models and deliver faster time to insight.

Enhancing Ethernet for AI Clouds

AI workloads require a low-latency network with high, effective bandwidth that's also compatible with collectives libraries such as the NVIDIA Collective Communications Library (NCCL). For cloud environments where multiple jobs run simultaneously, performance isolation is critical to prevent interference among tenants running on the same compute infrastructure. To meet the needs of AI clouds, which require both high-performance networking and cloud multi-tenancy, Spectrum-X leverages remote direct-memory access (RDMA) over Converged Ethernet (RoCE) extensions. Built on top of traditional Ethernet, these extensions improve effective bandwidth and ensure performance isolation between workloads.

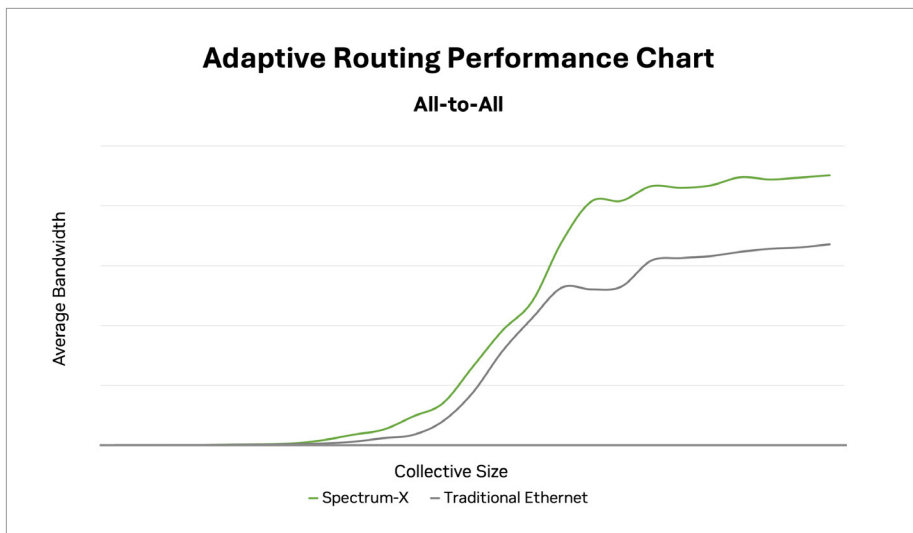
Key Features

NVIDIA Spectrum-X Networking Platform Components

- > Spectrum-4 Ethernet switches
- > BlueField-3 SuperNICs
- > LinkX® cables and transceivers

RoCE Extensions Enabling NVIDIA Spectrum-X

- > RoCE adaptive routing
- > RoCE congestion control
- > RoCE performance isolation



Spectrum-X delivers superior NCCL performance compared to traditional Ethernet

The Key Benefits of NVIDIA Spectrum-X

- > **Improved AI cloud performance:** Spectrum-X enhances AI cloud networking performance by 1.6X, accelerating processing, analysis, and execution of AI workloads and, in turn, the development and deployment of AI solutions.
- > **Standard Ethernet connectivity:** While enhancing Ethernet functionality, Spectrum-X is fully standards-based and is completely interoperable with Ethernet-based stacks.
- > **Increased power efficiency:** By improving performance, NVIDIA Spectrum-X contributes to a more power-efficient AI environment. This leads to reduced power consumption and lower operational costs for AI clusters.
- > **Enhanced multi-tenant performance:** Performance isolation in multi-tenant environments ensures that each tenant's workloads perform optimally and consistently, resulting in higher customer satisfaction and improved service quality.
- > **Better AI fabric visibility:** Visibility into the flows as they traverse multiple switches makes it possible to identify performance bottlenecks and is a key part of a modern, automated fabric-validation solution.
- > **Higher AI scalability:** Unprecedented scale to 256 200G ports in a single hop or 16K ports in a two-tier leaf/spine topology supports the expansion of the AI cloud while maintaining high levels of performance, making it an ideal solution for organizations with evolving AI infrastructure needs.
- > **Faster network setup:** The automated, end-to-end configuration of advanced networking functionality is fully tuned for AI workloads.
- > **Higher resiliency:** With higher resiliency, the cascading performance issues that occur with a lost link are eliminated, limiting the loss in bandwidth to that single link.
- > **Strong tenant isolation:** Leveraging BlueField-3 SuperNICs, Spectrum-X delivers enhanced workload isolation through secure virtual private clouds (VPCs) supporting each tenant.

Increased Power Efficiency

Power capping has become a common practice in data centers due to the growing demand for computing resources and the need to control energy costs. Overall, improving AI performance per watt is essential to achieve greater computational efficiency and deliver insights faster, while staying within power budgets. This is particularly important in applications such as deep learning, where training models can be computationally intensive and require a large amount of power. In addition to its raw performance benefits, Spectrum-X improves power efficiency, delivering superior performance per watt over traditional Ethernet.

Spectrum-X Technology Innovations

Spectrum-4 switches and BlueField-3 SuperNICs work in tight coordination to form a **NCCL-optimized network fabric** built to optimize AI cluster performance using a suite of end-to-end innovations:

- > **RoCE adaptive routing** avoids congestion by dynamically routing large AI flows away from congestion points. This approach improves network resource utilization, leaf/spine efficiency, and performance. The Spectrum-4 switch employs fine-grained load balancing, re-routing active flows to eliminate congestion. Additionally, the BlueField-3 SuperNICs work in tandem to handle out-of-order packets, placing packets in the correct order in the destination memory. RoCE adaptive routing supports profiles for efficient provisioning and automation.
- > **RoCE congestion control** collects network performance data with in-band network telemetry. The BlueField-3 SuperNICs use the collected switch telemetry data to optimize network data rates. BlueField algorithms use deep learning models for data metering, optimizing settings for multi-job, multi-tenant systems.
- > **End-to-end visibility with NVIDIA NetQ™** traces flow-level performance from the GPU to the DPU and maps the path and per-hop behavior across switch ports and RoCE queues.
- > **NVIDIA full-stack integration** includes NVIDIA NetQ, NCCL, Nsight™, H100 Tensor Core GPU, BlueField-3, and Spectrum-4— all configurable with NVIDIA Bright Cluster Manager for simple, production-ready deployment and faster time to AI.
- > **Very low latency** is critical for AI and machine learning workloads that require real-time processing. The Spectrum-4 switch delivers the industry's lowest-latency 200/400GbE switching, ensuring ultra-low latency and jitter for 256 ports.
- > **NCCL-optimized** switch behavior thresholds enable the optimization of buffer and congestion thresholds for AI workloads. This ensures synchronized collective operations, reducing the likelihood of congestion and packet loss.

Spectrum-X offers a groundbreaking solution for organizations building Ethernet-based, multi-tenant AI clouds. Spectrum-X enhances performance and energy efficiency of AI clouds across various applications, resulting in higher predictability and consistency. This leads to faster time to market and a stronger competitive advantage.

NVIDIA Spectrum-X Networking Platform

The Spectrum-X networking platform consists of the following components.

The Spectrum-4 Ethernet switch for smart-leaf, spine, and super-spine designs offers 64 ports of 800GbE in a dense 2U form factor, playing a vital role in NVIDIA Spectrum-X deployments. The SN5600 switch supports both standard leaf and spine designs with top-of-rack switches and rail-optimized end-of-row topologies. The SN5600 offers diverse connectivity in combinations of 10-800GbE and boasts an industry-leading total throughput of 51.2 terabits per second (Tb/s).

The BlueField-3 SuperNIC delivers up to 400GbE RoCE network connectivity between GPU servers and enables NVIDIA GPUDirect® RoCE for optimizing peak AI workload efficiency. The BlueField-3 SuperNIC accelerates AI networking and performs low-compute tasks. It is best suited for the B3140H model, providing an HHHH form-factor, and a sub-75 watt power envelope, making it an ideal choice for NVIDIA Spectrum-X deployments. NVIDIA integrates BlueField-3 SuperNICs across its data center computing systems, including the HGX™ H100 and OVX™ L40S-based systems.

LinkX transceivers and cables provide a robust selection of direct attach copper cables (DACs), active copper cables (ACCs), active optical cables (AOCs), and optics needed to connect octal small form-factor pluggable (OSFP)- and quad small form-factor pluggable (QSFP)-based fabrics. Spectrum-4 switches, combined with the LinkX portfolio, enable the most power-efficient AI fabric on the market.

Ready to Get Started?

Learn more about the NVIDIA Spectrum-X Networking Platform at: nvidia.com/en-us/networking/spectrumx

© 2023 NVIDIA Corporation and affiliates. All rights reserved. NVIDIA, the NVIDIA logo, BlueField, Cumulus, NetQ, Nsight, and Spectrum are trademarks and/or registered trademarks of NVIDIA Corporation and affiliates in the U.S. and other countries. Other company and product names may be trademarks of the respective owners with which they are associated. 2954800. OCT23

