



NVIDIA Spectrum-X

Purpose-built for Ethernet AI clouds.

The NVIDIA Spectrum™-X networking platform is the first Ethernet platform designed specifically to improve the performance and efficiency of Ethernet-based AI clouds. This breakthrough technology achieves 1.7X better overall AI performance and energy efficiency, 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 plus NVIDIA® BlueField®-3 data processing unit (DPU). The delivery of end-to-end AI capabilities reduces runtimes of massive transformer-based generative AI models and allows network engineers, AI data scientists, and cloud service providers to obtain accurate results and make informed decisions faster.

The Challenge of Traditional Ethernet With AI

AI clouds that use traditional Ethernet for their compute fabric can only achieve a fraction of the MLPerf performance that they would achieve with an optimized network. In multi-tenant environments where multiple AI jobs run simultaneously, performance isolation is critical to prevent further degradation of performance. And if there's a link fault, the traditional Ethernet fabric can cause the cluster's AI performance to drop by half. This is because traditional Ethernet has primarily been optimized for everyday enterprise workflows and isn't designed to meet the demands of high-performance AI applications that rely on the NVIDIA Collective Communications Library (NCCL).

These performance issues are due to factors inherent to traditional Ethernet:

- > Higher switch latencies, common across commodity ASICs
- > Split buffer switch architecture, which can lead to bandwidth unfairness
- > Load balancing that's sub-optimized for the large flows generated by AI workloads
- > Performance isolation and noisy neighbor issues

The Spectrum-X networking platform, based on Spectrum-4 and BlueField-3, solves these issues with traditional Ethernet.

Key Features

NVIDIA Spectrum-X Networking Platform Components

- > Spectrum-4 Ethernet switches (SN5000 Series)
- > BlueField-3 DPUs
- > NVIDIA LinkX® cables and transceivers
- > Spectrum-X license

RoCE Extensions Enabling Spectrum-X

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

The Key Benefits of NVIDIA Spectrum-X

- > **Improved AI cloud performance:** Spectrum-X enhances AI cloud performance by 1.7X and higher, accelerating processing, analysis, and execution of AI workloads and, in turn, the development and deployment of AI solutions.
- > **Standard Ethernet connectivity:** Spectrum-X is powered by NVIDIA innovation, is fully standards-based Ethernet, and is completely interoperable with Ethernet-based stacks.
- > **Increased power efficiency:** By improving performance, 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 running across the AI cloud 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 256x 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.
- > **Secure infrastructure:** Leveraging BlueField-3, Spectrum-X enhances encryption functionality and features such as deep packet inspection to ensure security and isolation of the control plane.

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 outperforms all other Ethernet solutions, delivering 1.7X superior power efficiency (performance per watt).

NCCL-Optimized Networking

Spectrum-4 switches and BlueField-3 DPUs work in tight coordination to combine with the Spectrum-X license. Together, they 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 DPUs 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 DPUs 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 DPU delivers high-performance compute operating at 400 gigabits per second (Gb/s) with line-rate processing of software-defined networking, storage, and cybersecurity. BlueField-3 combines powerful computing capabilities, high-speed networking, and extensive programmability to power Spectrum-X with accelerated RoCE, along with multi-tenancy and zero-trust security, crucial for the control plane.

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.

The Spectrum-X license unlocks the full power of Spectrum-4 and BlueField-3, enabling the RoCE extension features and optimized network parameter settings that deliver enhanced AI performance.

Spectrum-X Networking Platform Ordering Specifications

For more information about ordering Spectrum-4 switches, BlueField-3 DPUs, and LinkX cables and transceivers, please see the [SN5000 switch datasheet](#) and the [BlueField-3 DPU datasheet](#).

The Spectrum-X license is a license that enables the software-defined, hardware-accelerated feature set required for generative AI on Ethernet. Without the license, these features, including the RoCE extensions, aren't available.

For deployments that use standard, out-of-the-box RoCE, the Spectrum-X license isn't needed. The Spectrum-X license is licensed on a per-switch basis but unlocks end-to-end value from switch to DPU to GPU.

Spectrum-X license part numbers are bundled to include support.

Support is available at business-standard and business-critical service-level agreements (SLAs) for three- and five-year periods. It's recommended that support levels for the Spectrum-X license align with the base switch hardware support levels. That is, business-critical Spectrum-X support should be paired with business-critical hardware support, and vice versa. For dynamic renewal part numbers, please contact your NVIDIA representative.

Spectrum-X licenses are applicable to both NVIDIA Cumulus® Linux and Open Network Install Environment (ONIE) Spectrum-4 deployments. Licenses may be bought at time of switch purchase or may be added later as a standalone purchase when you need Spectrum-X functionality.

NVIDIA Spectrum-X License Ordering Information

OPN	Description
798-XSPEXZ-P3CMI36	Spectrum-X license subscription with business-standard support for three years
798-XSPEXZ-P3CMI60	Spectrum-X license subscription with business-standard support for five years
798-YSPEXZ-P3CMI36	Spectrum-X license subscription with business-critical support for three years
798-YSPEXZ-P3CMI60	Spectrum-X license subscription with business-critical support for five years
798-SPEXEVAL	Spectrum-X license evaluation (90 days)

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. 2761336. MAY23

