



NVIDIA Spectrum-X800 Ethernet Platform

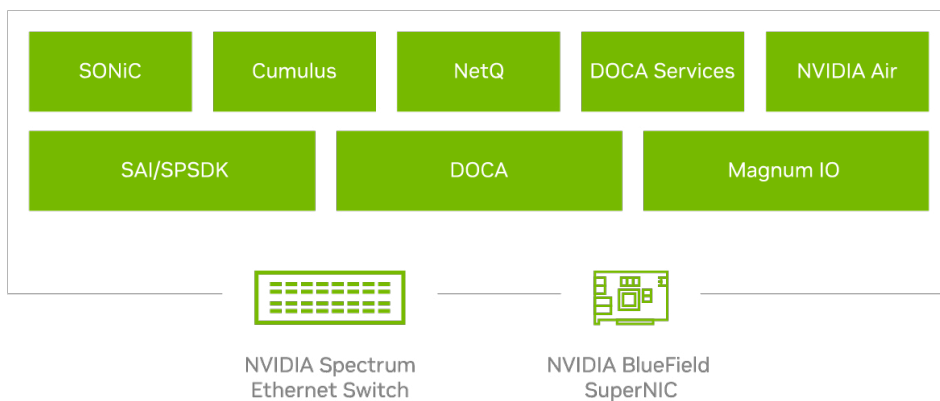
AI-optimized networking for every data center.



NVIDIA Spectrum™-X800 is the next generation of Spectrum-X technology and the world's first end-to-end, AI-optimized Ethernet platform for generative AI clouds and enterprises. Featuring the NVIDIA Spectrum-X800 SN5600 800 gigabits per second (Gb/s) Ethernet switch and NVIDIA® BlueField SuperNIC™, the Spectrum-X800 platform provides the optimal solution for building multi-tenant, hyperscale AI clouds using standards-based Ethernet. By leveraging adaptive routing and congestion control, Spectrum-X800 provides the optimum effective bandwidth and noise isolation, ensuring the highest-performance Ethernet networking for AI.

Purpose-built for multi-tenant environments, Spectrum-X800 enables the highest-performance Ethernet for generative AI. It offers performance isolation to ensure that the AI workloads of each tenant perform optimally and consistently, resulting in higher customer satisfaction and improved service quality. In addition, it includes visibility features centered around streaming telemetry for identifying performance bottlenecks and complete automated fabric validation.

Spectrum-X800 delivers consistent, predictable outcomes for thousands of simultaneous AI jobs at every scale. This empowers cloud service providers (CSPs) and large enterprises to accelerate the development, deployment, and time to market of AI solutions, improving return on investment.



A block diagram showcasing the software stack powering Spectrum-X800.

Key Features

Spectrum-X800 Ethernet further boosts innovative features introduced with Spectrum-X to provide the highest performance for Ethernet networks.

- > **Remote direct-memory access (RDMA) over converged Ethernet (RoCE) adaptive routing:** Spectrum-X800 features adaptive routing for lossless networks, closely integrating the switch and SuperNIC to boost bandwidth and resilience in AI fabrics.
- > **Programmable congestion control:** Spectrum-X800 uses advanced congestion control techniques to enhance noise isolation in multi-tenant AI environments.
- > **Low latency:** The Spectrum-4 switch offers the lowest-latency Ethernet and large radix, enabling Spectrum-X800 deployments to scale to hundreds of thousands of GPUs without latency-related performance loss.

NVIDIA Spectrum-X800 Technical Overview

The Spectrum-X800 SN5600, with 64 ports of 800G octal small form-factor pluggable (OSFP) and 51.2 terabits per second (Tb/s) of switching capacity, is the fastest single application-specific integrated circuit (ASIC) Ethernet switch (by aggregate bandwidth) in production. Based on Spectrum-4, it has 4X switching capacity and 2X port count compared to the previous-generation switch. The SN5600 switch, which also powers the 400G Spectrum-X platform, has already been widely adopted by OEMs and customers.

The BlueField-3 SuperNIC is an advanced network accelerator, purpose-built for supercharging hyperscale AI workloads. Boasting a power-efficient, lean design, the BlueField-3 SuperNIC is optimized for network-intensive, massively parallel computing, accelerating RoCE network connectivity between GPU servers, and delivering peak AI workload efficiency. Ushering in a new era of AI cloud computing, the BlueField-3 SuperNIC enables secure, multi-tenant data center environments while ensuring deterministic and isolated performance between jobs and tenants.

NVIDIA Spectrum-X800 platform connectivity options with the NVIDIA LinkX® interconnect portfolio of products provide the maximum flexibility to build a preferred network topology. This is achieved by using optical transceivers with passive fiber cables, direct attached copper (DAC) cables, and linear active copper cables (LACCs). The 800Gb/s links can reach up to 2 kilometers.

- > **NVIDIA Collective Communications Library (NCCL) optimization:** Spectrum-X800 optimizes buffer and congestion thresholds for NCCL, enhancing operations and mitigating congestion and packet loss.
- > **NVIDIA NetQ™ visibility for AI:** NetQ tracks GPU-to-SuperNIC performance, mapping paths and per-hop behavior across switch ports and RoCE queues.
- > **NVIDIA full-stack integration:** Includes NVIDIA NetQ, Air, NCCL, Nsight™, DOCA™, BlueField-3 SuperNIC, and Spectrum-4, all configurable with simple production-ready deployment for faster time to AI.

“NVIDIA Spectrum-X800 is the next generation of Spectrum-X technology and the world’s first end-to-end, AI-optimized 800G Ethernet platform for generative AI clouds and enterprises.”

Ready to Get Started?

Learn more by contacting an NVIDIA sales representative:

www.nvidia.com/en-us/contact/sales