



NVIDIA Spectrum SN6000 Ethernet Switch Series

Ethernet purpose-built for AI infrastructure at any scale, across any distance.



The NVIDIA Spectrum™ SN6000 series is the sixth generation of NVIDIA Ethernet switches. Built on the Spectrum-6 ASIC, these systems deliver high throughput, low latency, and advanced data center capabilities for fabrics of any size. Dynamic, flexible shared buffers and wire-speed performance ensure consistent, predictable results for both AI and cloud-native workloads.

SN6000 switches power key NVIDIA platforms—including [Spectrum-X™ Ethernet](#), MGX™, DGX™, HGX™, and RTX PRO™ Server—and integrate seamlessly with software suites such as NVIDIA AI Enterprise and NVIDIA Mission Control™.

Together, these innovations define the next generation of the NVIDIA Spectrum-X Ethernet platform, engineered with extreme codesign for NVIDIA Vera Rubin to enable the next wave of massive-scale AI factories and pave the way for future million-GPU environments.

AI-Optimized Switching Architecture

The Spectrum-6 Ethernet switching ASIC delivers best-in-class performance for AI and cloud-native workloads across core, cloud, and edge environments. Enhanced RoCE features like adaptive routing, performance isolation, and auto-path configuration set a new benchmark for high-bandwidth AI workloads running over Ethernet.

NVIDIA Spectrum-X Ethernet Photonics

NVIDIA Spectrum-X Ethernet Photonics delivers the latest advances in co-packaged optical (CPO) technology, bringing cutting-edge photonics innovation directly into the Ethernet switching platform. Spectrum-6 CPO-based switches boast higher power efficiency, greater reliability, and longer uptimes for AI applications. Spectrum-X Ethernet Photonics switch systems integrate optical engines directly onto the same package as the switch ASIC, eliminating traditional pluggable transceivers. Internal optical routing brings fiber straight to the package, where silicon-photonics engines convert light to electrical signals adjacent to the ASIC. An external laser source (ELS) delivers continuous light to the integrated optical engines, lowering on-package heat and enhancing reliability at scale.

Key Features

Performance

- > 102.4 Tb/s aggregate bandwidth per ASIC
- > Built on 200G SerDes
- > 800 Gb/s per port
- > Optimized for Spectrum-X Ethernet platform
- > Adaptive routing
- > Congestion control
- > Performance isolation
- > Fully shared, monolithic packet buffer
- > Hardware-based multiplane support

Telemetry and Visibility

- > Hardware-accelerated high-frequency telemetry pipeline
- > Packet-granular event-level diagnostics
- > Support for NVUE, switch SDK, SONiC CLI
- > Real-time telemetry via open APIs (OTLP, gNMI, gNOI)

Achieving GigaScale AI With NVIDIA Spectrum-XGS Ethernet Technology

As AI workloads grow, data centers face hard limits in power and physical footprint, making it increasingly difficult to scale within a single site. Connecting multiple facilities with off-the-shelf Ethernet networking introduces long-tail latency, jitter, and unstable performance. NVIDIA Spectrum-XGS Ethernet technology solves this by enabling high-performance scale-across networking, allowing multiple geographically separated data centers to operate as one unified AI factory.

Built on the NVIDIA SN6000 series of Ethernet switches and Spectrum-X SuperNIC™s, Spectrum-XGS Ethernet adds a distance-aware congestion control algorithm that efficiently balances traffic across sites. With end-to-end telemetry, precise latency management, and deterministic performance, Spectrum-XGS Ethernet technology enables facilities separated by hundreds of kilometers and more to function as a single, coherent AI environment.

Spectrum-X Multiplane Topologies

Multiplane is the breakthrough that allows Ethernet to scale from thousands to hundreds of thousands of GPUs in a flat 2-tier topology. Effective scaling requires overcoming challenges due to switch radix limits, oversubscription, and component reliability. Spectrum-X multiplane solves this with hardware-coordinated SuperNIC routing. In a Spectrum-X multiplane network, each SuperNIC (which corresponds to a GPU) is split into two or more ports, which each connects into a separate network plane. Each network plane features its own distinct leaf/spine/optional superspine layer. The SuperNIC intelligently routes network traffic to different planes, load balancing packet-by-packet in hardware for optimal performance.




Figure 1: Spectrum-X Multiplane Topology

In addition, Spectrum-X multiplane improves resiliency by offering multiple independent paths between GPUs. The SuperNIC leverages high-frequency telemetry from the switch and integrates with the NVIDIA Collective Communications Library (NCCL) to gain a complete view of the plane topologies and route traffic away from planes with problematic link outages or congestion issues.

SN6000 Series Switches

The SN6000 series, based on the 102.4 Tb/s Spectrum-6 ASIC, is available in multiple configurations for high-performance L2/L3 forwarding across leaf, spine, and superspine deployments. The platforms deliver full wire-speed throughput, ultra-low latency, and 160 MB per ASIC of fully shared packet buffer for fair and predictable performance.

Virtualization, Agility, and Security


- > Programmable packet-processing pipeline
- > Deep packet parser/editor supporting up to 512 bytes
- > Single-pass VXLAN routing and bridging
- > Centralized VXLAN routing
- > Hardware acceleration for routing, NAT, and tunnel termination
- > Hardware root of trust for secured switching
- > Zero-trust security foundation for multi-tenant AI clouds
- > Open network operating system (NOS) ecosystem support

Cloud-Tier Scale

- > Programmable Network Control and Policy support for 750,000 entries (shared across routes, media access control (MAC), access control list (ACL))
- > Precision Time Protocol (PTP) time synchronization



SN6600-LD*



SN6600*



SN6800-LD*



SN6810-LD*

* For illustration only. Actual products may vary.

SN6600

The SN6600 system is an AC-powered, air-cooled, connector-to-power (port-side intake) Ethernet switch, featuring 64 integrated heatsink (IHS) type OSFP cages supporting 1.6 Tb/s optics connecting 128 ports at speeds of up to 800 Gb/s in a 3RU form factor. It delivers 102.4 Tb/s of total throughput for scalable GPU-cluster connectivity. Leveraging the ecosystem of fully retimed and transit-retimed pluggable transceivers, it is optimized for the NVIDIA Spectrum-X Ethernet reference architecture for AI-optimized deployments.

SN6600-LD

The SN6600-LD is a fully liquid-cooled switch with UQD8 v2 connectors, featuring 64 OSFP cages supporting 1.6 Tb/s optics connecting 128 ports of up to 800 Gb/s in a compact 2RU form factor. It provides 102.4 Tb/s of total throughput for scalable GPU-cluster connectivity. Leveraging the existing pluggable transceiver ecosystem, it is optimized for NVIDIA Spectrum-X Ethernet and supports AI-optimized leaf/spine architectures in middle-of-row (MoR) and end-of-row (EoR) designs. Its 48–60 VDC busbar power design reduces cabling, improves power efficiency, and simplifies maintenance in high-density data center environments.

SN6800-LD

SN6800-LD is the industry's first 4-ASIC co-packaged Ethernet switch for AI scale-out, delivering unmatched radix, optical density, and power efficiency for next-generation AI fabrics. Optimized for Spectrum-X Ethernet, the SN6800-LD is a 5RU switch featuring 512 MMC-16 4-lane optical connectors, each providing four 200 Gb/s lanes, with one lane connected to each ASIC. It delivers 409.6 Tb/s of switching capacity and supports up to 2,048 breakout connections for massive GPU-cluster scale-out. The SN6800-LD features an integrated shuffle-box to natively support multiplane topologies.

The SN6800-LD features a 5RU liquid-cooled design with UQD8 v2 connectors and a 48–60 VDC busbar power architecture for efficient rack integration into high-density AI data centers.

SN6810-LD

The SN6810-LD is a Spectrum-6 switch designed to deliver high-performance Ethernet connectivity for AI clusters. Driven by a single Spectrum-6 ASIC, it delivers 102.4 Tb/s of switching capacity leveraging silicon-photonics co-packaged optics (CPO) for efficient, low-latency single-mode fiber connectivity.

The SN6810-LD features a 2RU liquid-cooled design with UQD8 v2 connectors and a 48–60 VDC busbar power architecture for efficient rack integration into high-density AI data centers.

Rack-Scale Ethernet Switch Systems – SN6200-LD

The NVIDIA SN6200-LD is a rack-integrated Spectrum-X Ethernet switch tray purpose-built for NVIDIA MGX ETL256 rack architectures. Implemented as a compact 1U system, the SN6200-LD is designed as an integral component of the ETL256 rack, leveraging the rack's shared backplane, 48–60 VDC busbar power, and liquid-cooling infrastructure. Powered by the Spectrum-6 Ethernet ASIC, the SN6200-LD delivers 102.4 Tb/s of switching throughput. It provides the integrated Spectrum-X Ethernet fabric for high-bandwidth, low-latency connectivity for up to 32 compute nodes and 256 GPUs within the rack, with AI-optimized capabilities for RoCE-based training and inference.

Factory-integrated 256 × 200GbE copper backplane connections eliminate intra-rack optics and loose cabling, reducing power, cost, complexity, and potential points of failure. Pre-integrated into the ETL256 rack with factory-validated wiring and topology, the Spectrum SN6200-LD simplifies deployment and lifecycle operations. Up to 32 front-panel OSFP ports supporting 800GbE enable seamless rack-to-rack and spine connectivity, while the resilient multiplane architecture enhances scale and maintains connectivity across all 256 GPUs during switch tray servicing or failure. For more details, see [Technical Specifications: Rack-Scale Switch Systems](#).

Platform Security

SN6000 series switches offer complete security across all switch layers. Hardware, firmware, and software are authenticated by a built-in root of trust, from the basic input and output system (BIOS) to the network operating system (NOS). Any attempt to run an altered component or image is automatically blocked, ensuring secure operation across data center environments.

Baseboard Management Controller

The platform uses a baseboard management controller (BMC), an integral part of the security root of trust, to provide an interface to the baseboard management system (BMS). The BMC can be managed via either a serial port or a BMC Ethernet port. It supports remote system power management and reports liquid leak detection events.

Network Operating Systems

- > **NVIDIA Cumulus Linux:** [NVIDIA Cumulus™ Linux](#) is a powerful network operating system (NOS) that enables advanced automation, customization, and scalability using API-first principles hardened in the world's largest data centers. Cumulus Linux is the only NOS that brings operational efficiency to every AI factory.
- > **SONiC:** SONiC is a fully open source, hardware-agnostic NOS designed for hyperscalers, service providers, and enterprises. NVIDIA's [Pure SONiC](#) distribution adds NVIDIA expertise and support. Pure SONiC is fully supported across all SN6000 systems.

Platform Software Options

- > **NVIDIA DSX Air:** NVIDIA DSX Air simplifies deployments by enabling digital twins of the entire network to design, test, and validate network provisioning, automation, and security policies. NVIDIA DSX Air allows day-zero operations before hardware is deployed by simulating and automating changes ahead of production, and accelerates deployment of networking infrastructure, reducing time to first token while ensuring cloud-scale efficiency and reliability.

Technical Specifications* (SN6800-LD, SN6810-LD, SN6600-LD, and SN6600)

Switch Model	SN6800-LD	SN6810-LD	SN6600-LD	SN6600
Optical Form Factor	512 MMC-16 800 Gb/s co-packaged optics (CPO)	128 MMC-16 800 Gb/s co-packaged optics (CPO)	64 OSFP cages supporting up to 1.6 Tb/s optical transceivers 2x 800 Gb/s) liquid-DC	64 OSFP cages supporting up to 1.6 Tb/s optical transceivers (2x 800 Gb/s) air-DC
800 Gb/s Ports	-	128	128	128
400 Gb/s Ports	-	256	256	256
200 Gb/s Ports	2,048	512	512	512
100 Gb/s Ports	2,048+8	512+2	512+2	512+2
25 Gb/s Ports	-	-	-	2
External Laser Source (ELS)	64	16	-	-
Switching Capacity (Tb/s)	409.6 Tb/s [4x 102.4 Tb/s]	102.4 Tb/s	102.4 Tb/s	102.4 Tb/s
Lanes per OSFP/ MMC-16	4x 200G PAM4	4x 200G PAM4	8x 200G PAM4	8x 200G PAM4
CPU	16-core x86, AMD	8-core x86, AMD	8-core x86, AMD	8-core x86, AMD
System Memory	64 GB DDR5	64 GB DDR5	32 GB DDR5	32 GB DDR5
SSD Memory	160 GB NVMe	160 GB NVMe	80 GB NVMe	80 GB NVMe
Packet Buffer	4x 256 MB	256 MB	256 MB	256 MB
OOB Management Ports, RJ45	2x 1GbT	2x 1GbT	1x 1GbT	1x 1GbT
Serial Console Ports	1 RJ45	1 RJ45	1 RJ45	1 RJ45
USB Ports	USB-A v2	USB-A v2	USB-A v2	USB-A v2
BMC Ports RJ45	1x 1GbT	1x 1GbT	1x 1GbT	1x 1GbT
Cooling Specifications	Liquid connector: 10x UQD8 v2 LFR: 13.4 LPM Coolant temp: 17–45°C	Liquid connector: 4x UQD8 v2 LFR: 3.4 LPM Coolant temp: 17–45°C	Liquid connector: 2x UQD8 v2 LFR: 5.7 LPM Coolant temp: 17–45°C	Air-cooled
Operating Conditions	Operational ambient: 5–40°C, ASHRAE A3	Operational ambient: 5–40°C, ASHRAE A3	Operational ambient: 5–40°C, ASHRAE A3	Operational ambient: 0°–30°C, ASHRAE A3
System Power	48 VDC/54 VDC [busbar] Permissible range: 40–60 VDC	48 VDC/54 VDC [busbar] Permissible range: 40–60 VDC	48 VDC/54 VDC [busbar] Permissible range: 40–60 VDC	High Line, Rated VAC 220V–240 V 50/60 Hz
Size (H x W x D)	H: 8.75" (220 mm), 5RU W: 17.2" (438 mm) D: 35.4" (900 mm)**	H: 3.43" (87 mm), 2RU W: 17.2" (438 mm) D: 35.4" (900 mm)**	H: 3.43" (87 mm), 2RU W: 17.2" (438 mm) D: 35.4" (900 mm)**	H: 5.15" (133.35 mm), 3RU W: 17.2" (438 mm) D: 30.1" (767 mm)

* This document describes hardware features and capabilities. For information on feature availability, refer to the switch software release notes.

**Including UQD and front levers (as applicable).

Technical Specifications: Rack-Scale Switch Systems*

Switch Model	SN2600-LD
Form Factor	32 OSFP pluggable cages supporting up to 1.6 Tb/s (2x 800 Gb/s) + 256x 200GbE backplane ETL256 Rack
800 Gb/s Ports	32
400 Gb/s Ports	64
200 Gb/s Ports	256
100 Gb/s Ports	256
25 Gb/s Ports	-
External Laser Source (ELS)	-
Switching Capacity (Tb/s)	102.4 Tb/s
Lanes per OSFP/MMC-16	8x 200G PAM4
CPU	8-core x86, AMD
System Memory	32 GB DDR5
SSD Memory	90 GB NVMe
Packet Buffer	256 MB
OOB Management and BMC	1GbT RJ45
Serial Ports	1 USB-C
USB Ports	-
BMC Ports	-
Cooling Specifications	Liquid-cooled
Operating Conditions	Operational ambient: 0–30°C, ASHRAE A3
System Power	48 VDC/54 VDC [busbar] Permissible range: 40–60 VDC
Height	1U

* This document describes hardware features and capabilities. For information on feature availability, refer to the switch software release notes.

Compliance

Standards Compliance

Safety	CB, CE, cTUVus, CU
EMC	CE, ICES, FCC, RCM, VCCI
Non-Operating Conditions	-40–70°C
Relative Humidity	5%–85%
Operating Altitude	0–3,050 m
RoHS	RoHS-compliant

Enterprise Support and Services

A minimum of one year of [Enterprise Business-Standard Support](#) is required when purchasing NVIDIA Spectrum SN6000 switches:

- NVIDIA Enterprise Support provides access to NVIDIA experts, the NVIDIA Enterprise Support Portal, advanced return material authorization (RMA), and more.
- Add-on services—including installation, configuration, technical account manager, four-hour on-site engineer, expedited RMA, media retention, and more—are available.

For details, visit the [NVIDIA Enterprise Support and Services User Guide](#)

Product Specifications

Details of the NVIDIA Spectrum Ethernet SN6000 series switches are available in the [SN6000 Switch Systems User Manual](#).

Transceivers and Cables

- For details on NVIDIA cables and transceivers, visit the [Interconnect documentation hub](#).
- Some transceivers may require higher than typical power delivery. Please refer to the [SN6000 Switch Systems User Manual - Interconnects](#) for detailed information on switch ports' power specifications.

Ready to Get Started?

To learn more about NVIDIA Spectrum Ethernet SN6000 switches, including product specifications and ordering information, refer to the [SN6000 Switch Systems User Manual](#).