



NVIDIA Spectrum SN5000 Series Switches

Bringing accelerated Ethernet to every data center for AI and cloud.

The NVIDIA Spectrum™ SN5000 series of Ethernet switches is the fifth generation of NVIDIA Ethernet switches, purpose-built to accelerate data center fabrics. With port speeds spanning from 10 to 800 gigabits per second (Gb/s), the SN5000 series switches deliver accelerated Ethernet to every data center without compromising performance or feature set.

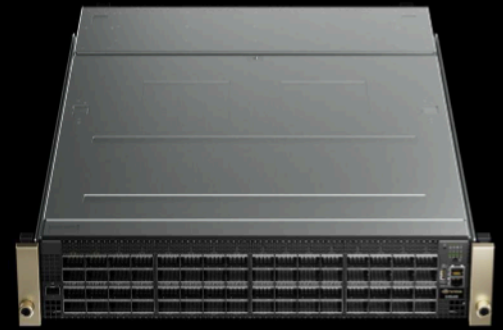
Accelerating Data Center Fabrics

The SN5000 series is ideal for enabling cloud-scale infrastructures for data centers of any size. SN5000 switch systems provide high performance, consistent low latency, and support for advanced data center networking features, making them ideal for cloud networks and end-to-end data center fabrics.

Powered by the NVIDIA Spectrum-4 application-specific integrated circuit (ASIC), the SN5000 series features dynamic, flexible shared buffers and predictable performance. SN5000 switches are built to accelerate clouds and power NVIDIA platforms, including the NVIDIA Spectrum-X™ networking platform, NVIDIA MGX™, NVIDIA DGX™, NVIDIA HGX™, NVIDIA RTX™ Pro Server, and AI software tools such as NVIDIA AI Enterprise and NVIDIA Mission Control™.

As part of the Spectrum platform, SN5000 systems are pretested and prevalidated with NVIDIA's full portfolio of Ethernet networking technology, including NVIDIA® BlueField® SuperNIC™ cards and data processing units (DPUs), ConnectX® SuperNICs and network interface cards (NICs), and LinkX® interconnects. While each element is fully compatible with standard Ethernet fabrics, the end-to-end switch-to-host solution is critical to powering accelerated workloads, delivering the high performance and innovative features needed to supercharge AI and cloud-native applications at scale.

SN5000 switches are also a key component of the NVIDIA Spectrum-X networking platform. This platform integrates Spectrum 5000 series Ethernet switches (excluding SN5400) with NVIDIA SuperNICs, standing out as the world's first Ethernet fabric built specifically for AI. This unique design accelerates generative AI network performance, boasting a 1.6x improvement over traditional Ethernet fabrics. Additionally, Spectrum-X ensures consistent and predictable outcomes for thousands of simultaneous AI jobs at every scale by maximizing resource utilization and performance isolation. It not only enables advanced cloud multi-tenancy but also offers robust zero-trust security. For cloud service providers, this translates into faster AI solution development, quicker deployment, and accelerated time to market, ultimately enhancing their return on investment.



Key Features

Performance

- > Up to 800 Gb/s per port for 51.2 Tb/s aggregate switch bandwidth
- > Fully shared packet buffer provides a fair, predictable, and high-bandwidth data path
- > Consistent and low cut-through latency
- > Accelerated remote direct-memory access (RDMA) over converged Ethernet (RoCE) with extensions for AI cloud servers built on Ethernet
- > Best-in-class Virtual Extensible LAN (VXLAN) scale
- > 512,000 forwarding entries flexibly shared across access control list (ACL), longest prefix match (LPM) routes, host routes, media access control (MAC) list, equal-cost multi-path (ECMP), and tunnel applications

Flexible, High-Performance Configuration

With a bidirectional switching capacity of 51.2 terabits per second (Tb/s), SN5000 platforms are available in a range of configurations. Each delivers high performance with feature-rich layer-2 and layer-3 forwarding, ideally suited for both end-of-row and spine and suitable for replacing modular chassis switches. SN5000 switches provide full wire-speed performance with ultra-low cut-through latency, leveraging the fully shared 160 megabyte (MB) packet buffers for fair and predictable performance. By combining the low latency of the Spectrum-4 ASIC with the large switch radix of SN5000 systems, thousands of hosts can be connected in a two-tier network architecture while maintaining minimal port-to-port latencies.

Adding a wide range of innovations in the areas of programmability, telemetry, and tunneling, the NVIDIA SN5000 series can address the complex networking requirements of today's data centers. With a range of system form factors and a rich software ecosystem, NVIDIA SN5000 series allows you to pick and choose the right components for your data center.

SN5610/SN5600/SN5600D

The SN5610 and SN5600 spine and super-spine switches offer 64 ports of 800 gigabit Ethernet (GbE) in a dense 2U form factor. These switches are optimized for NVIDIA Spectrum-X deployments, supporting standard spine architectures with middle-of-row (MoR) and end-of-row (EoR) designs. The SN5600 and SN5610 offer diverse connectivity options, including combinations of 100 GbE, 200 GbE, 400 GbE, and 800 GbE, and boast a total throughput of 51.2 Tb/s. The SN5600D is a variation of the SN5600 that features a direct current (DC) bus bar for power intake, making it ideal for NVIDIA DGX SuperPOD™ offerings.

While sharing identical scale and connectivity capabilities with the SN5600, the SN5610 introduces a modified hardware specification that optimizes system power consumption and enhances efficiency.

SN5400

The SN5400 smart-leaf, spine, and super-spine switch offers 64 ports of 400 GbE in a 2U form factor. The SN5400 offers diverse quad small form factor double density (QSFP-DD) connectivity options ranging from 1 to 400 GbE and boasts a total throughput of 25.6 Tb/s, making it ideal for 400 G data center fabrics.

High Availability

The NVIDIA Spectrum SN5000 series of Ethernet switches is designed with the following features for high availability:

- 1+1 or 2+2 hot-swappable power supplies
- N+1 hot-swappable fans
- Ethernet VPN (EVPN) multi-homing
- Multi-chassis link aggregation group (MLAG) for active/active L2 multipathing
- Hardware-assisted in-service software upgrades (ISSU)
- 256-way equal-cost multi-path (ECMP) routing for load balancing and redundancy

Agility and Security

- Comprehensive layer 2 and layer 3
- Advanced network virtualization with high-performance, single-pass VXLAN routing
- Cloud-scale network address translation (NAT)
- Programmable pipeline that can programmatically parse, process, and edit packets
- Secured switch offerings via hardware root of trust

Visibility

- NVIDIA® What Just Happened® (WJH) telemetry dramatically reduces mean time to issue resolution
- Hardware-accelerated histograms powered by NetQ track and summarize queue depths at submicrosecond granularity
- Inband network telemetry (INT)
- Streaming telemetry
- 512,000 on-chip flow counters

Platform Security

SN5000 series switches offer a complete security suite at 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 that wasn't originally signed by NVIDIA is automatically blocked, ensuring the safest network installation within data centers.

NVIDIA Cumulus Linux

NVIDIA Cumulus® Linux is a powerful open network operating system that enables advanced automation, customization, and scalability using web-scale principles hardened in the world's largest data centers. Cumulus Linux was built for building data center networks ideally suited to diverse business needs. And it's the only open NOS that enables the building of affordable and efficient network operations like the world's largest data center operators, unlocking web-scale networking for businesses of all sizes.

SONiC

Designed for hyperscalers, service providers, and enterprises, SONiC is a fully open-sourced, hardware-agnostic NOS, perfect for preventing vendor lock-in and ideal for next-generation data centers. SONiC's containerized design makes it flexible and customizable, allowing customers to combine and manage SONiC and non-SONiC switches within the same networking fabric. NVIDIA's pure SONiC offering removes distribution limitations and enables enterprises to take full advantage of the benefits of open networking while adding the NVIDIA expertise and support that best guarantee success. NVIDIA is proud to guide the direction of SONiC as a member of the SONiC Governing Board and Technical Steering Committee. Pure SONiC is fully supported on all SN5000 switch systems.

Platform Software Options

NVIDIA Air

The NVIDIA Air infrastructure simulation platform creates digital twins of SN5000 switch systems (as well as the rest of the Spectrum portfolio). The digital twin includes logical instances of every switch and cable, so it can be used to validate security policy compliance, automation processes, monitoring tools, interoperability, and upgrade procedures. The digital twin is key to transforming network operations models, allowing IT architects and infrastructure specialists to deploy and update networks up to 95% faster through continuous integration and continuous delivery (CI/CD) integration.

NVIDIA NetQ

NVIDIA NetQ™ is a highly scalable network operations toolset that provides visibility, troubleshooting, and validation of networks in real time. NetQ delivers actionable insights and operational intelligence about the health of data center networks—from the container or host all the way to the switch and port—enabling a NetDevOps approach. NVIDIA NetQ is the leading network operations tool that utilizes telemetry for deep troubleshooting, visibility, and automated workflows from a single interface, reducing mean time to innocence (MTTI) and network downtimes.



SN5610 switch image



SN5600/SN5600D switch image



SN5400 switch image

NVIDIA Spectrum-4: An ASIC for Accelerated Data Centers

NVIDIA's Spectrum-4 is the fifth-generation Ethernet switch ASIC, delivering a solid balance of performance, virtualization, telemetry, and scalability capabilities.

Groundbreaking Performance and Features

Packet buffer architecture has a major impact on overall switch performance. The Spectrum-4 ASIC features a fully shared and monolithic packet buffer across all ports, supporting cut-through traffic from all ports without compromising scale or features. With its fast packet buffer, Spectrum-4 provides a high-performance, fair, and bottleneck-free data path for mission-critical applications. In addition, Spectrum-4 provides the most advanced and innovative feature set yet. This includes unique features like RoCE extensions for NVIDIA Spectrum-X and end-to-end nanosecond-level timing precision from switch to host.

Centralized Visibility

Spectrum-4 provides deep and contextual network visibility, which allows network operators to proactively manage issues and reduce mean time to recovery (MTTR) and innocence (MTTI). Available through the switch SDK, NVIDIA User Experience (NVUE), or SONiC command-line interface (CLI), What Just Happened leverages the underlying silicon and software capability to provide granular and event-triggered information about infrastructure issues. In addition, Spectrum-4's rich telemetry information is readily available via open APIs and tools such as gNMI that can be integrated with NetQ, as well as third-party software tools and workflow engines.

Unprecedented Agility

For modern data center infrastructure to be software-defined and agile, both its compute and network building blocks need to be agile. Spectrum-4 includes a unique, feature-rich packet processing pipeline that offers data center network virtualization without compromising on performance or scale. Spectrum-4 has a programmable pipeline and a deep packet parser and editor that can process payloads up to the first 512 bytes. Spectrum-4 supports single-pass VXLAN routing and bridging. In addition, Spectrum-4 supports advanced virtualization features such as routing and NAT.

Massive Scale

Spectrum-4 uses intelligent algorithms and efficient resource sharing and supports unprecedented scale for forwarding tables, counters, and policies. This fine-grained resource allocation enables Spectrum-4 to fit all specific needs, allowing up to 512,000 entries to be dynamically shared across media access control (MAC), Address Resolution Protocol (ARP), Internet Protocol version 4 (IPv4) routes, ACLs, ECMP, and VXLAN tunnels.

Technical Specifications

	SN5610	SN5600	SN5600D	SN5400
Connectors	64 OSFP 800 GbE + 2 SFP28 25 GbE	64 OSFP 800 GbE + 1 SFP28 25 GbE	64 OSFP 800 GbE + 1 SFP28 25 GbE	64 QSFP-DD 400 GbE + 2 SFP28 25 GbE
Max. 800 GbE ports	64	64	64	-
Max. 400 GbE ports	128	128	128	64
Max. 200 GbE ports	256	256	256	128
Max. 100 GbE ports	256	256	256	256
Max. 50 GbE ports	256	256	256	256
Max. 40 GbE ports	128	128	128	128
Max. 25 GbE ports	256+2	256+1	256+1	256+2
Max. 10 GbE ports	256+2	256+1	256+1	256+2
Max. 1 GbE ports	2	1	1	2
Switching capacity (Tb/s)	51.2 Tb/s	51.2 Tb/s	51.2 Tb/s	25.6 Tb/s
Wire-speed switching (Bpps)	33.3 Bpps	33.3 Bpps	33.3 Bpps	33.3 Bpps
Lanes per port x max. speed per lane	8x 100 G PAM4	8x 100 G PAM4	8x 100 G PAM4	8x 50 G PAM4
CPU	Octa-core AMD	Hexa-core x86	Hexa-core x86	Hexa-core x86
System memory	32 GB	32 GB	32 GB	32 GB
SSD memory	80 GB NVMe	160 GB SATA-3	160 GB SATA-3	160 GB SATA-3
Packet buffer	160 MB	160 MB	160 MB	160 MB
100/1000 Mb/s management ports	RJ45	RJ45	RJ45	RJ45
Serial ports	RJ45	RJ45	RJ45	RJ45
USB ports	USB Type 3	USB Type 3	USB Type 3	USB Type 3
Hot-swap power supplies	4 (2+2 redundant)	2 (1+1 redundant)	2 (1+1 redundant)	2 (1+1 redundant)
Hot-swappable fans	5 (N+1 redundant)	4 (N+1 redundant)	4 (N+1 redundant)	4 (N+1 redundant)
Airflow	Reverse	Reverse	Reverse	Reverse, forward
Power supplies	Frequency: 50–60 Hz Input range: 200–240 VAC	Frequency: 50–60 Hz Input range: 200–240 VAC	DC bus bar Input range: 40–60 VDC	Frequency: 50–60 Hz Input range: 200–240 VAC
Operating conditions	0–40°C	0–35°C	0–35°C	0–35°C
Size (H x W x D)	3.39" (H) x 17.2" (W) x 31" (D) 86.2 mm (H) x 438 mm (W) x 788 mm (D)	3.39" (H) x 17.2" (W) x 31" (D) 86.2 mm (H) x 438 mm (W) x 788 mm (D)	3.39" (H) x 17.2" (W) x 31" (D) 86.2 mm (H) x 438 mm (W) x 788 mm (D)	3.39" (H) x 17.2" (W) x 31" (D) 86.2 mm (H) x 438 mm (W) x 788 mm (D)

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 SN5000 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 SN5000 series switches are available in the [SN5000 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 [SN5000 Switch Systems User Manual](#) for detailed information on switch ports' power specifications.

Ready to Get Started?

To learn more about NVIDIA Spectrum SN5000 switches, including product specifications and ordering information, refer to the [SN5000 Switch Systems User Manual](#).