



NVIDIA Spectrum SN2000 Series Switches

Ethernet networking switches for out-of-band management.

The NVIDIA Spectrum™ SN2000 Series of switches are ideal for out-of-band management networks in modern data centers. They offer flexibility with port speeds from 1 gigabit per second (Gb/s) to 100 Gb/s and port density, ensuring full-rack connectivity at any speed. These switches include 48 G/100M/10M Base-T host ports and non-blocking 100GbE spine uplinks. Featuring advanced hardware and software, application-specific integrated circuit (ASIC)-level telemetry, and a 16-megabyte (MB), fully shared buffer, the SN2000 Series offers differentiated switching capabilities for baseboard management controller (BMC) connectivity.

The SN2000 Series consists of the SN2201 and SN2201-M. Both SN2201 switches are one rack unit (1RU) switches that feature 48 RJ45 ports supporting 100BASE-T and 1000BASE-T speeds up to 100 meters over CAT-5E or CAT-6 cables. The four QSFP28 ports offer 1, 10, 25, 40, 50, or 100GbE speeds and can break out into 16 ports with flexible speed options. These ports support direct-attach copper (DAC) and optical transceivers ranging from 0.5 meters (m) to 80 kilometers (km). Additionally, the SN2000 Series provides full, hot-swappable redundancy for power supply units (PSUs) and fans, ensuring reliable, continuous operations.

In addition to the traditional AC-powered SN2201, the SN2201-M switch features a compact 1U form factor, compliant with NVIDIA MGX™ v1.2 rack specifications. It uses a 48 volts direct current (VDC) central busbar power architecture with connector-to-power air cooling for efficient power delivery, cooling, and easy installation.

The SN2000 Series provides ultra-rich layer-2 and layer-3 forwarding, virtual extensible LAN (VXLAN) and Ethernet VPN (EVPN) tunneling, multi-homing, telemetry, and versatile features, ideal for both data and management networking.

Data Center Simulation

NVIDIA enables simulation of the SN2000 switches through the NVIDIA Air Infrastructure Simulation Platform. The simulation includes logical instances of every switch and cable, enabling it to be used for validating security policy compliance, automation processes, monitoring tools, interoperability, and upgrade procedures. The simulation is key to transforming network operations models, enabling IT architects and infrastructure specialists to deploy and update networks up to 95 percent faster through continuous integration and continuous delivery (CI/CD).


Key Features

Visibility

- NVIDIA® What Just Happened® (WJH) telemetry: dramatically reduces mean time to issue resolution by providing answers to When, What, Who, Where, and Why
- In-Band Network Telemetry (INT)- ready hardware
- Streaming telemetry
- Up to 256,000 shared forwarding entries

Agility

- Comprehensive layer 2, layer 3, and RDMA over converged Ethernet (RoCE)
- Advanced network virtualization with high-performance single-pass VXLAN routing and IPv6 segment routing
- Programmable pipeline
- Fully shared packet buffer: provides fair, predictable, and high-bandwidth data paths



SN2000 Series Software Ecosystem

NVIDIA Cumulus Linux

NVIDIA Cumulus® Linux is an advanced, open network operating system (OS) designed for robust automation, flexibility, and scalability. Cumulus Linux stands out as the only open network OS, enabling businesses of all sizes to build affordable and efficient network operations akin to the world's largest data center operators. It also provides the tools to tailor data center and campus networks to specific business requirements.

NVIDIA NetQ

NVIDIA NetQ™ is a highly scalable network operations toolset, offering real-time network visibility, troubleshooting, and validation. It provides actionable insights into the health of data center networks, from container or host to switch and port, enabling a NetDevOps approach. NetQ leverages telemetry for deep troubleshooting, visibility, and automated workflows from a single interface, reducing mean time to innocence (MTTI) and network downtimes. With full lifecycle management, NetQ simplifies upgrading, configuring, and deploying network elements while maintaining comprehensive operations capabilities like trace and comparative look-back functionality.

NVIDIA Spectrum: Build Your Cloud Without Compromise

The NVIDIA Spectrum Ethernet switch ASIC delivers a solid balance of performance, virtualization, and telemetry capabilities.

Groundbreaking Performance

Packet buffer architecture has a major impact on overall switch performance. The Spectrum packet buffer is monolithic and fully shared across all ports, supporting cut-through line-rate traffic from all ports without compromising scale or features. With its fast packet buffer, Spectrum provides a high-performance, fair, and bottleneck-free data path for mission-critical applications.

Centralized Visibility

Spectrum offers in-depth and contextual network visibility. It enables network operators to proactively manage issues, reducing mean times to recovery and innocence. WJH harnesses the underlying silicon and software capabilities to provide granular, event-triggered insights into infrastructure issues. In addition, Spectrum's rich telemetry information is readily available through open APIs, making it easy to integrate with third-party software tools and workflow engines.

Unprecedented Agility

For a data center to be software-defined and agile, both its compute and network building blocks must be flexible. Spectrum offers a distinctive, feature-rich, and efficient packet-processing pipeline that enables advanced data center network virtualization without sacrificing performance or scalability. It boasts a programmable pipeline and an in-depth packet parser and editor capable of handling payloads up to the first 512 bytes. Spectrum also supports single-pass VXLAN routing and bridging, advanced virtualization features like IPv6 segment routing, and network address translation (NAT).

Massive Scale

Spectrum uses intelligent algorithms and efficient resource sharing and supports unprecedented forwarding tables, counters, and policies.

- > Consistent, low cut-through latency: hardware-accelerated data movement with intelligent, congestion management and load balancing for RoCE and GPU-accelerated machine learning using NVIDIA GPUDirect®
- > Best-in-class VXLAN scale: supports 6x more tunnels and tunnel endpoints
- > Deep packet inspection: up to 512 bytes deep



SN2201



SN2201-M

For illustration only. Actual products may vary.

Technical Specifications*

	SN2201	SN2201-M
Connectors	48 RJ45 + 4 QSFP28 100GbE4	48 RJ45 + 4 QSFP28 100GbE4
Max 100GbE ports	4	4
Max 50GbE ports	8	8
Max 40GbE ports	4	4
Max 25GbE ports	16	16
Max 10GbE ports	16	16
Switching Capacity (Tb/s)	448 Gb/s	448 Gb/s
Wire-Speed Switching (Bpps)	667 Mpps	667 Mpps
Lanes per Port x Max Speed per Lane	4x 25G NRZ	4x 25G NRZ
CPU	Dual-core x86	Dual-core x86
System Memory	8 GB	8 GB
SSD Memory	20 GB	20 GB
Packet buffer	16 MB	16 MB
100/100 Management Ports	1	1
Serial Ports	1 RJ45	1 RJ45
USB Ports	1	1
Hot-Swap Power Supplies	2 (1+1 redundant)	N/A
Hot-Swappable Fans	4 (N+1 redundant)	4 (N+1 redundant)
Airflow Options	Connector to power (reverse) Power to Connector (Forward)	Connector to power (reverse)
Power Supplies	Frequency: 50–60 Hz Input range: 100–264 Vac AC input current: 4.5–2.9 A	48 VDC via busbar Input range: 40–60 VDC
Typical Power (ATIS)	100 W (estimated)	90 W (estimated)
Rack Mounting	19" EIA standard	19" MGX v1.2 NVIDIA
Size (W x H x D)	1.72" x 16.84" x 17" 43.9 mm x 428 mm x 432 mm	1.72" x 17.2" x 30.7" 43.9 mm x 438 mm x 781 mm

* 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
Operating conditions	Operating: 0–40° C
Non-operating conditions	Non-operating: -40–70° C
Relative humidity	Operating: 10–85% non-condensing
	Non-operating: 10–90% non-condensing
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 SN2000 Series Ethernet 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, review the [NVIDIA Enterprise Support and Services User Guide](#).

Product Specifications

Details of the NVIDIA Spectrum SN2000 series Ethernet switches are available in the [SN2000 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 [SN2000 Switch Systems User Manual](#) for detailed information on switch ports' power specifications.

Ready to Get Started?

To learn more about the NVIDIA Spectrum SN2000 Ethernet switches, including product specifications and ordering information, refer to the [SN2000 Switch Systems User Manual](#).