



NVIDIA Spectrum SN3000 Series Switches

Data center performance, scale, and rich telemetry.

Flexibility and Performance for Data Center Networks

NVIDIA® Spectrum™ SN3000 switches, based on the second-generation Spectrum switch application-specific integrated circuit (ASIC), are purpose-built for leaf, spine, and super-spine data center applications. Allowing maximum flexibility, the SN3000 series provides port speeds spanning from 1 to 100 gigabits per second (Gb/s) with a port density that enables full-rack connectivity to any server at any speed. In addition, the uplink ports allow a variety of blocking ratios to suit any application requirement.

The SN3000 series is ideal for building cloud-scale layer-2 and layer-3 networks. Delivering high performance, consistent low latency, and support for advanced software-defined networking (SDN) features, the SN3000 series is the ideal choice for web-scale IT, cloud, hyper-converged storage, and data analytics applications.

Network Disaggregation: Open Ethernet

NVIDIA Open Ethernet™ Spectrum switches break the paradigm of traditional switch systems, eliminating vendor lock-in. Instead of forcing network operators to use the specific software that's provided by the switch vendor, Open Ethernet offers the flexibility to use a choice of operating systems on top of Ethernet switches, thereby regaining control of the network and optimizing utilization, efficiency, and overall return on investment.

Encouraging an ecosystem of open-source, standard network solutions, Open Ethernet adopts the same principles as standard open solutions for servers and storage and applies them to the world of networking infrastructure. These solutions can be easily deployed into the modern data center across network equipment, easing management and ensuring full interoperability.

SN3420

In data center switching architectures built around 100GbE leaf-spine designs, the NVIDIA SN3420 provides a compact, cost-effective top-of-rack switch for mixed-speed environments. It offers 48 ports of 10/25GbE and 12 ports of up to 100GbE in a 1U form factor, using 25G NRZ per lane to support high-density server connectivity and spine uplinks. Delivering up to 2.4 Tb/s of throughput and 3.58 Bpps of processing capacity, the SN3420 enables efficient ToR aggregation and seamless QSFP28 connectivity for modern data center networks.

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.
- Hardware-accelerated histograms track and summarize queue depths at submicrosecond granularity.
- Inband network telemetry (INT)-ready hardware
- Streaming telemetry
- 512,000 on-chip flow counters

Performance

- Fully shared packet buffer provides a fair, predictable, and high-bandwidth data path.
- Consistent and low cut-through latency
- Robust remote direct-memory access (RDMA) over converged Ethernet (RoCE) transport to power non-volatile memory express (NVMe) over Fabrics and machine learning applications that leverage GPUDirect®

High Availability

The SN3420 switch is designed for high availability both from a software and hardware perspective. Key features include:

- > 1+1 hot-swappable power supplies and N+1 hot-swappable fans
- > Color-coded power supply units (PSUs) and fans
- > Up to 48x 10/25GbE or 12x 100GbE ports per link aggregation group (LAG)
- > Multi-chassis LAG for active/active L2 multipathing
- > 64-way equal-cost multi-path (ECMP) routing for load balancing and redundancy

SN3000 Series: A Rich Software Ecosystem

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 SN3000 systems.

NVIDIA Air

[NVIDIA Air](#) simplifies deployments by enabling digital twins of the entire network to design, test, and validate network provisioning, automation, and security policies. NVIDIA 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.

NVIDIA NetQ

With end-to-end network visibility powered by [NVIDIA NetQ™](#), operators gain flow-level insight across switches, GPUs, SuperNIC™s, and switch ports, including detailed per-hop behavior on RoCE queues. This integrated telemetry enables precise performance monitoring, rapid troubleshooting, and deterministic optimization of multi-tenant AI workloads running across the fabric.

ONIE

The Open Network Install Environment (ONIE) is an Open Compute Project, an open source initiative driven by a community to define an open “install environment” for bare-metal network switches, such as the NVIDIA SN3000 series. ONIE enables a bare-metal network switch ecosystem where end users have a choice of different network operating systems.

Docker Containers

NVIDIA fully supports the running of third-party containerized applications on the switch system itself. The third-party application has complete access to the bare-metal switch via its direct access to the SDK. The switch has tight controls over the amount of memory and CPU cycles each container is allowed to use, along with fine-grained monitoring of those resources.


- > Best-in-class Virtual Extensible LAN (VXLAN) scale—10x more tunnels and tunnel endpoints than others
- > 512,000 forwarding entries flexibly shared across access control list (ACL), longest prefix match (LPM) routes, host routes, media access control list (MAC), equal-cost multi-path (ECMP), and tunnel applications

Agility

- > Comprehensive layer-2, layer-3, and RoCE
- > Advanced network virtualization with high-performance single-pass VXLAN routing and IPv6 segment routing
- > Cloud-scale network address translation (NAT)—100,000+ sessions
- > Programmable pipeline that can programmatically parse, process, and edit packets
- > Deep packet inspection—512 bytes deep



SN3420



Docker Containers Support

NVIDIA Spectrum-2 ASIC: 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-2 packet buffer is 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-2 provides a high-performance, fair, and bottleneck-free data path for mission-critical applications.

Pervasive Visibility

Spectrum offers in-depth and contextual network visibility, enabling network operators to proactively manage issues, thereby reducing mean time to recovery or establishing innocence. WJH harnesses the underlying silicon and software capability 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 modern data center infrastructure to be software-defined and agile, both its compute and network building blocks must be agile. Spectrum-2 ASIC features a distinctive, feature-rich, and efficient packet processing pipeline that delivers advanced data center network virtualization without sacrificing performance or scalability. Not only does it have a programmable pipeline and in-depth packet parser and editor that can process payloads up to the first 512 bytes, Spectrum-2 also supports single-pass VXLAN routing and bridging, advanced virtualization features like IPv6 segment routing, and network address translation (NAT).

Massive Scale

The number of endpoints in the data center is experiencing exponential growth. This growth is further amplified by the ongoing shift from virtual machine-based architectures to container-based architectures, necessitating high-scale forwarding tables that modern data centers and mega clouds require—sometimes increasing by an order of magnitude or more. In response to the need for greater scalability and flexibility, Spectrum uses intelligent algorithms and efficient resource sharing and supports unprecedented forwarding tables, counters, and policy scale.

Building on this commitment to adaptability and precision, fine-grained resource allocation is tailored to fit all specific needs, allowing up to 512,000 entries to be dynamically shared across MAC, ARP, IPv4/IPv6 routes, ACLs, ECMP, and tunnels. An innovative algorithmic TCAM is further optimized for data centers and cloud environments, which can scale the number of rules to up to half a million.

End-to-End Solution

The SN3000 series is an integral component of NVIDIA's comprehensive end-to-end solutions, providing 10–100Gb/s of interconnectivity within the data center. In addition to the SN3000 series, NVIDIA's solutions also include NVIDIA ConnectX® smart network interface cards (SmartNICs), BlueField® data processing units (DPUs), and LinkX® copper and fiber cabling.

Technical Specifications

Switch Model	SN3420
Connectors	48 SFP28 25GbE + 12 QSFP28 100GbE
Max. 100GbE ports	12
Max. 50GbE ports	24
Max. 40GbE ports	12
Max. 25GbE ports	48+48
Max. 10GbE ports	48+48
Max. 1GbE ports	48+48
Switching capacity (Tb/s)	2.4 Tb/s
Wire-speed switching (Bpps)	3.58 Bpps
Lanes per port x max speed per lane	4x 25G NRZ
CPU	Dual-core x86
System memory	8 GB
SSD memory	32 GB
Packet buffer	42 MB
100/1,000Mb/s management ports	1
Serial ports	1 RJ45
USB ports	1
Hot-swap power supplies	2 (1+1 redundant)
Hot-swappable fans	5 (N+1 redundant)
Airflow options	Forward or reverse
Power supplies	Frequency: 50–60 Hz Input range: 100–264 Vac AC input current: 2.9–4.5 A
Size	1.72" (H) x 17.24" (W) x 18.29" (D) 44 mm (H) x 438 mm (W) x 464.6 mm (D)
Weight	8.5 kg (18.73 lb)

*This section describes hardware features and capabilities. For 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: -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 SN3000 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 more details, refer to the [NVIDIA Enterprise Support and Services User Guide](#).

Product Specifications

Details of the NVIDIA Spectrum SN3000 series of Ethernet switches are available in the [SN3000 Switch Systems User Manual](#).

Transceivers and Cables

- For details on NVIDIA cables and transceivers, visit the [Networking Interconnect documentation hub](#).
- Some transceivers may require higher than typical power delivery. Please refer to the [SN3000 Switch Systems User Manual](#) for detailed information on switch ports' power specifications.

Ready to Get Started?

To learn more about the NVIDIA Spectrum SN3000 series of Ethernet switches, including product specifications and ordering information, refer to the [SN3000 Switch Systems User Manual](#)