Shape the Future of QuickSpecs - Your Input Matters

NVIDIA Networking for HPE QuickSpecs

NVIDIA Networking for HPE includes NVIDIA Spectrum™-X SN5610, NVIDIA Spectrum-X SN2201, NVIDIA Quantum-2-based QM9700 switches.

The SN5600 is compatible with standard Ethernet fabric and provides accelerated Ethernet to your data center without compromising between performance and feature set. NVIDIA Networking for HPE features configurable 800GbE ports in a dense 2U form factor and can support up to 128 ports of 400GbE with bidirectional switching throughput of 51.2 Tb/s to easily address your data center networking requirements. The NVIDIA Spectrum™ SN2000 series switches are the 2nd generation of NVIDIA switches, purpose-built for leaf/spine/super-spine datacenter applications. NVIDIA Quantum-2 extends In-Network Computing acceleration technology with preconfigured, programmable engines.

Overview

NVIDIA Spectrum-X SN5610 Switch

NVIDIA SN5610 is the next generation Spectrum-X switch of SN5000 series. The NVIDIA Spectrum™-X SN5000 series switches are the fifth generation of Spectrum Ethernet switches, purpose-built to accelerate data center fabrics. With port speeds spanning from 10GbE to 800GbE, the SN5000 series switches deliver accelerated Ethernet to every data center without compromising between performance and feature set.

These switches are ideally suited for customers who deploy high-performance computing (HPC) clusters based on ProLiant DL, XL and Cray servers using Ethernet technology.

The 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 ASIC, the SN5000 series features dynamic, flexible shared buffers and predictable performance. As part of the Spectrum platform, SN5000 systems are pretested and prevalidated with NVIDIA's full portfolio of Ethernet networking technology, including NVIDIA® BlueField® data processing units (DPUs), ConnectX® smart network interface cards (SmartNICs), 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 cloud-native applications at scale.

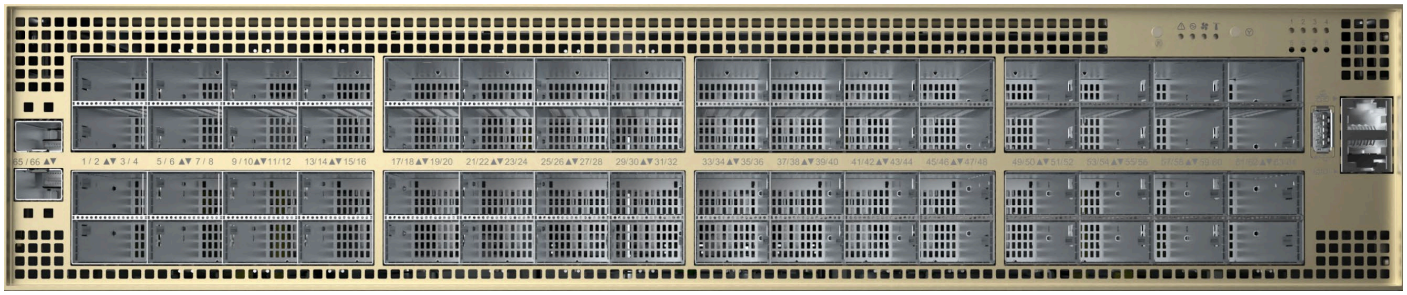

HPE Spectrum X SN5610 Switch

The SN5610 smart-leaf/spine/super-spine offers 64 ports of 800GbE in a dense 2U form factor. The SN5600 is ideal for accelerated Ethernet deployments and enables both standard leaf/spine designs with top-of-rack switches as well as end-of-row (EoR) topologies. The SN5610 offers diverse connectivity in combinations of 1 to 800GbE and boasts an industry-leading total throughput of 51.2Tb/s.

Overview

HPE Spectrum X SN5610 Switch	
Part Number	920-9N42F-00RI-3C1
ASIC	Spectrum-4, ASIC, 51.2T
Power Supply	4 x AC, 2kW
	PSU Inlet Connector: Type C-165
Airflow	Connector-to-Power, port side intake
Thermal	5 Fans. N+1 Redundant 60mm x 60mm
System's Compute	AMD EPYC 3251, 8-Cores, Secured boot
SSD	80GB NVMe
Front Access (Service) Ports	2 x SFP28 [1G/10G/25GbE]
Management Interfaces	Serial Console port RS-232 (1 x RJ45 port) USB port (1 x USB 2.0) OOB management Interface 1GbT, RJ-45

Models

NVIDIA Spectrum SN5610 64-port 800GbE OSFP 4x AC PSU Connector to Power Airflow Switch for HPE — P85922-B21

- NVIDIA Spectrum-4-based 800 Gb/s Ethernet switch with Cumulus Linux Authentication
- 64 OSFP ports in 2U form factor 4 x AC, 2kW
- 5 Fans. N+1 Redundant 60mm x 60mm
- 2 SFP28 port
- 128 400GbE
- Bidirectional switching capacity of 51.2Tb/s
- Two power supplies (AC),
- Standard depth
- Managed switch
- C2P airflow
- Tool-less Rail kit
- Quick Start Guide

Overview

Key Features

- Up to 800GbE per port for 51.2Tb/s aggregate switch bandwidth
- 64 OSFP ports in 2U form factor
- 51.2 Tb/s non-blocking aggregate bandwidth
- It supports both standard leaf/spine designs with top-of-rack switches and end-of-row topologies
- Quick and easy setup and management
- Fully shared packet buffer provides a fair, predictable, and high-bandwidth data path
- Ultra-low latency for efficient communication
- Network operating system NVIDIA Cumulus Linux
- Accelerated RDMA over Converged Ethernet (RoCE) for AI, cloud, storage, and high-performance computing (HPC)
- Serial Console port RS-232 (1 x RJ45 port) USB port (1 x USB 2.0)
- OOB management Interface 1GbT, RJ-45

Switch Specifications

- NVIDIA Spectrum-4-based SN5610 switch system
- 64 800 Gb/s non-blocking ports
- 128 400 Gb/s
- 64 octal small form-factor pluggable (OSFP) connectors
- Ethernet (800 Gbps) per port bidirectional
- Switching capacity 51.2 Tbps
- SSD Memory 80GB NVMe
- Thermal: 5 Fans. N+1 Redundant 60mm x 60mm
- Supports passive copper or active copper or fiber cable with optical module
- 4 AC hot-swappable power supplies
- AMD EPYC 3251, 8-Cores, Secured boot

Management Ports (Managed Switch)

- 1x RJ45 (1 Gb/s Ethernet)
- 1x RJ45 (UART) 1x USB 3.0

Connectors and cabling

- OSFP ports
- Passive or active copper or optical module with fiber cables

Overview

Indicators

- Per OSFP port status LED for link and activity
- System statues LEDs for system, fans, power supplies
- Unit ID (beacon) LED

Power Supply

- N+N Redundant, 5 FANS
- Input range: 200-240Vac; US min: 2 phases of 100-110v – in total, at least 208v
- 80 Gold+ and ENERGY STAR certified

Cooling

- Connector to power airflow (reverse)
- Hot-swappable fan units

Overview

NVIDIA Spectrum-X SN2201 Switch

The NVIDIA® Spectrum™ SN2000 series switches are the 2nd generation of NVIDIA switches, purpose-built for leaf/spine/super-spine datacenter applications. Allowing maximum flexibility, SN2000 series provides port speeds spanning from 1GbE to 100GbE, with a port density that enables full rack connectivity to any server at 1/20/25/40/50/100GbE speeds. In addition, the uplink ports allow a variety of blocking ratios to suit any application requirement.

The SN2000 series is ideal for building wire-speed and cloud-scale layer-2 and layer-3 networks. The SN2000 platforms deliver high performance, consistent low latency along with support for advanced software defined networking features, making it the ideal choice for web scale IT, cloud, hyperconverged storage and data analytics applications.

The SN2201 is ideal as an out-of-band (OOB) management switch, or as a top of rack (ToR) switch connecting up to 48 x 1G Base-T host ports with non-blocking 100GbE spine uplinks. Featuring highly advanced hardware and software along with ASIC-level telemetry and a 16MB fully shared buffer, the SN2201 delivers unique and innovative features to 1G switching.

NVIDIA Spectrum-X SN2201 Switch

Models

NVIDIA 48-port 1GbE RJ45 Connector to Power Airflow Switch for HPE P81269-B21

- 48 RJ45 + 4 QSFP28 100GbE
- Cumulus Linux Authentication
- Switching capacity 448 Gb/s
- Standard depth
- 1G switching
- C2P airflow
- Tool-less Rail kit
- Quick Start Guide
- RoHS compliant
- Operating temperature: 0–40ºC; Non-operating: -40–70ºC

Overview

Key Features

- Hardware-accelerated histograms track and summarize queue depths at sub-microsecond granularity
- Inband Network Telemetry (INT)–ready hardware
- Streaming telemetry
- Up to 256,000 shared forwarding entries

Overview

NVIDIA NDR IB Quantum-2 QM9700 Switch

The NVIDIA® Quantum-2 InfiniBand platform provides AI developers and scientific researchers with the highest networking performance available to take on the world's most challenging problems. Innovative NVIDIA In-Network Computing acceleration engines provide ultra-low latency and double the data throughput, while delivering the scalability and feature-rich capabilities required for supercomputers, artificial intelligence, and hyperscale cloud data centers.

NVIDIA Quantum-2 enhances and extends In-Network Computing acceleration technology with preconfigured and programmable engines such as NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP) ™, Message Passing Interface (MPI) tag matching, MPI_Alltoall, and programmable cores, as well as full transport offload with RDMA, GPUDirect RDMA, and GPUDirect Storage to deliver the best cost per node and return on investment (ROI).

The NVIDIA Quantum-2-based QM9700 and QM9790 switch systems deliver 64 ports of 400Gb/s InfiniBand per port, packed into a 1U standard chassis. Each switch carries an impressive 51.2 terabits per second (Tb/s) of aggregated bidirectional throughput and a landmark capacity exceeding 66.5 billion packets per second (BPPS). Supporting NVIDIA's advanced 400Gb/s interconnect technology, NVIDIA Quantum-2 is a high-speed, extremely low-latency, and scalable networking solution. Key features incorporated include state-of-the-art technologies such as remote direct-memory access (RDMA), adaptive routing, and NVIDIA® Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™. Unlike other networking solutions, NVIDIA InfiniBand incorporates self-healing network capabilities, quality of service, enhanced virtual lane mapping, and advanced congestion control, maximizing overall application throughput. As ideal rack-mounted InfiniBand solutions, the QM9700 and QM9790 400Gb/s InfiniBand fixed-configuration switches offer exceptional flexibility, with support for various topologies like Fat Tree, SlimFly, DragonFly+, multi-dimensional Torus, and others.

They maintain backward compatibility with previous generations and support a broad software ecosystem. The QM9700 400Gb/s InfiniBand switches, enhancing NVIDIA In-Network Computing technologies, incorporate the third generation of NVIDIA SHARP technology, SHARPv3.

NVIDIA NDR IB Quantum-2 QM9700 Switch

Models

NVIDIA InfiniBand NDR 64-port OSFP Managed Connector to Power Airflow Switch P81263-B21

- NVIDIA Quantum-2-based 400Gb/s InfiniBand switch
- 64 400Gb/s ports
- 32 OSFP ports
- Non-blocking switching capacity of 51.2Tb/s
- Two power supplies (AC)
- Standard depth
- Unmanaged switch
- Power-to-connector (P2C) airflow
- Rail kit
- Quick Start Guide
- RoHS compliant
- Operating temperature: Reverse Airflow: 0–40ºC; Non-operating: -40–70ºC

Key Features

- 64 ports, each providing NDR (400Gbps) performance
- 51.2 Tb/s non-blocking aggregate bandwidth
- NDR200 Supports up to 128 ports of HDR (200 Gbps)
- Quick and easy setup and management
- Improved performance by removing fabric congestion
- Backward compatibility to Nvidia HDR InfiniBand technology for ease of migration

NVIDIA NetQ Software for Spectrum X Switches (On Prem)

Network Operations with NetQ

NVIDIA NetQ™ is a highly scalable, modern network operations toolset that provides visibility, troubleshooting, and validation of your Cumulus fabrics in real time. NetQ utilizes telemetry and delivers actionable insights about the health of your data center network, ensuring your AI network fabric is operating smoothly

How NetQ works?

Data Collection

Using agents on the switches and hosts, NetQ is able to collect telemetry data across the entire network.

Data Processing

As a central control point, NetQ stores and processes information to provide actionable insights and complete visibility.

Data Visualization

The rich graphical user interface (GUI) simplifies operations and increases efficiency by quickly highlighting issues and alerts.

Further information available on https://www.nvidia.com/en-us/networking/ethernet-switching/netq/

Data Sheet: https://nvdam.widen.net/s/cqlbxvtw98/ethernet-switching-netq-datasheet

Overview

Key Features

- Validations, Tracing, Flow telemetry analysis
- RDMA over Converged Ethernet (RoCE) monitoring
- Software upgrade management
- Snapshot and compare
- Topology
- High-availability clustering
- Precision Time Protocol (PTP) monitoring
- Application-specific integrated circuit (ASIC) monitoring

Proof Points

- Simplifies scaling of NVIDIA Cumulus® Linux
- Speeds mean time to innocence
- Reduces opex, downtime
- Increases productivity
- Maximizes value of network infrastructure

Service and Support

HPE Services

No matter where you are in your digital transformation journey, you can count on HPE Services to deliver the expertise you need when, where and how you need it. From planning to deployment, ongoing operations and beyond, our experts can help you realize your digital ambitions.

https://www.hpe.com/services

Consulting Services

No matter where you are in your journey to hybrid cloud, experts can help you map out your next steps. From determining what workloads should live where, to handling governance and compliance, to managing costs, our experts can help you optimize your operations.

https://www.hpe.com/services/consulting

HPE Managed Services

HPE runs your IT operations, providing services that monitor, operate, and optimize your infrastructure and applications, delivered consistently and globally to give you unified control and let you focus on innovation.

HPE Managed Services | HPE

Operational services

Optimize your entire IT environment and drive innovation. Manage day-to-day IT operational tasks while freeing up valuable time and resources. Meet service-level targets and business objectives with features designed to drive better business outcomes.

https://www.hpe.com/services/operational

HPE Complete Care Service

HPE Complete Care Service is a modular, edge-to-cloud IT environment service designed to help optimize your entire IT environment and achieve agreed upon IT outcomes and business goals through a personalized experience. All delivered by an assigned team of HPE Services experts. HPE Complete Care Service provides:

- A complete coverage approach -- edge to cloud
- An assigned HPE team
- Modular and fully personalized engagement
- Enhanced Incident Management experience with priority access
- Digitally enabled and AI driven customer experience

https://www.hpe.com/services/completecare

HPE Tech Care Service

HPE Tech Care Service is the operational support service experience for HPE products. The service goes beyond traditional support by providing access to product specific experts, an AI driven digital experience, and general technical guidance to not only reduce risk but constantly search for ways to do things better. HPE Tech Care Service delivers a customer-centric, AI driven, and digitally enabled customer experience to move your business forward. HPE Tech Care Service is available in three response levels. Basic, which provides 9x5 business hour availability and a 2-hour response time. Essential which provides a 15-minute response time 24x7 for most enterprise level customers, and Critical which includes a 6-hour repair commitment where available and outage management response for severity 1 incidents.

https://www.hpe.com/services/techcare

HPE Lifecycle Services

HPE Lifecycle Services provide a variety of options to help maintain your HPE systems and solutions at all stages of the product lifecycle. A few popular examples include:

- Lifecycle Install and Startup Services: Various levels for physical installation and power on, remote access setup, installation and startup, and enhanced installation services with the operating system.
- HPE Firmware Update Analysis Service: Recommendations for firmware revision levels for selected HPE products, taking into account the relevant revision dependencies within your IT environment.
- HPE Firmware Update Implementation Service: Implementation of firmware updates for selected HPE server, storage, and solution products, taking into account the relevant revision dependencies within your IT environment.
- Implementation assistance services: Highly trained technical service specialists to assist you with a variety of activities, ranging from design, implementation, and platform deployment to consolidation, migration, project management, and onsite technical forums.

- HPE Service Credits: Access to prepaid services for flexibility to choose from a variety of specialized service activities, including assessments, performance maintenance reviews, firmware management, professional services, and operational best practices.

Notes: To review the list of Lifecycle Services available for your product go to:
https://www.hpe.com/services/lifecycle

For a list of the most frequently purchased services using service credits, see the HPE Service Credits Menu

---

Other Related Services from HPE Services:

HPE Education Services

Training and certification designed for IT and business professionals across all industries. Broad catalogue of course offerings to expand skills and proficiencies in topics ranging from cloud and cybersecurity to AI and DevOps. Create learning paths to expand proficiency in a specific subject. Schedule training in a way that works best for your business with flexible continuous learning options.

https://www.hpe.com/services/training

Defective Media Retention

An option available with HPE Complete Care Service and HPE Tech Care Service and applies only to Disk or eligible SSD/Flash Drives replaced by HPE due to malfunction.

Consult your HPE Sales Representative or Authorized Channel Partner of choice for any additional questions and services options.

Parts and Materials

HPE will provide HPE-supported replacement parts and materials necessary to maintain the covered hardware product in operating condition, including parts and materials for available and recommended engineering improvements.

Parts and components that have reached their maximum supported lifetime and/or the maximum usage limitations as set forth in the manufacturer's operating manual, product quick-specs, or the technical product data sheet will not be provided, repaired, or replaced as part of these services.

How to Purchase Services

Services are sold by Hewlett Packard Enterprise and Hewlett Packard Enterprise Authorized Service Partners:

- Services for customers purchasing from HPE or an enterprise reseller are quoted using HPE order configuration tools.
- Customers purchasing from a commercial reseller can find services at https://ssc.hpe.com/portal/site/ssc/

Service and Support

AI Powered and Digitally Enabled Support Experience
Achieve faster time to resolution with access to product-specific resources and expertise through a digital and data driven customer experience

Sign into the HPE Support Center experience, featuring streamlined self-serve case creation and management capabilities with inline knowledge recommendations. You will also find personalized task alerts and powerful troubleshooting support through an intelligent virtual agent with seamless transition when needed to a live support agent.

https://support.hpe.com/hpesc/public/home/signin

Consume IT On Your Terms

HPE GreenLake edge-to-cloud platform brings the cloud experience directly to your apps and data wherever they are—the edge, colocations, or your data center. It delivers cloud services for on-premises IT infrastructure specifically tailored to your most demanding workloads. With a pay-per-use, scalable, point-and-click self-service experience that is managed for you, HPE GreenLake edge-to-cloud platform accelerates digital transformation in a distributed, edge-to-cloud world.

- Get faster time to market
- Save on TCO, align costs to business
- Scale quickly, meet unpredictable demand
- Simplify IT operations across your data centers and clouds

To learn more about HPE Services, please contact your Hewlett Packard Enterprise sales representative or Hewlett Packard Enterprise Authorized Channel Partner. Contact information for a representative in your area can be found at "Contact HPE" https://www.hpe.com/us/en/contact-hpe.html

For more information
http://www.hpe.com/services

Summary of Changes

Date	Version History	Action	Description of Change
06-Oct-2025	Version 4	Changed	Information for NVIDIA Spectrum-X SN5610 Switch was added in Overview section.
28-Jul-2025	Version 3	Changed	Update survey link.
02-Jun-2025	Version 2	Changed	Overview section was updated.
04-Nov-2024	Version 1	Created	New QuickSpecs

Copyright

Shape the Future of QuickSpecs - Your Input Matters

Chat now

To learn more, visit http://www.hpe.com/

a00038966enw - 16132 - Worldwide - V4 - 06-October-2025
HEWLETT PACKARD ENTERPRISE
Hpe.com