

Overview

HPE Storage Switch M-series SN3420M with NVIDIA Cumulus® Linux

HPE M-Series with NVIDIA Cumulus® Linux family of Ethernet switches are capable of addressing today's data center's complex networking requirements, growth, and expansion and are perfect for Top-Of-Rack (TOR) deployments and optimized for virtualized environments, hyperconverged infrastructure, and storage deployments. HPE M-Series Ethernet switches give you the right network bandwidth with consistent zero packet loss performance for high-performance and storage workloads.

M-series SN3420M with NVIDIA Cumulus® Linux is based on the 2nd generation of NVIDIA Spectrum™ switches, purpose-built for leaf/spine/super-spine data center applications. Allowing maximum flexibility, the 1U SN3420M provides port speeds spanning from 1GbE to 100GbE, and a port density that enables full rack connectivity to any server. SN3420M provides an on chip fully-shared 42MB packet buffer, and flexible port use in addition to advanced capabilities. The uplink ports allow a variety of blocking ratios to suit any application requirement. The SN3420M is ideal for building cloud-scale layer-2 and layer-3 networks. The SN3420M platform delivers 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.

With an increasing need to access data faster and accommodate growing workloads, rising levels of east-west traffic, and new storage arrays based on flash storage technologies, a high bandwidth, low-latency, zero packet loss network becomes paramount. The HPE SN3420M Switch with NVIDIA Cumulus® Linux offers a mix of SFP28 and QSFP28 connectivity for top-of-rack applications requiring matched downlink to uplink bandwidth. NVIDIA Cumulus® network platforms are capable of delivering unbelievable networking speed and agility to keep pace with the most intense workloads small to large scale enterprises can produce. With port speeds spanning 1Gb/s to 100 Gb/s and a switching capacity of 4.8Tb/s from 48 ports at 25GbE and 12 ports at 100GbE. This switch provides non-blocking throughput at wire-speed transfers-across all packet sizes. The SN3420M delivers a landmark 3.58 Bpps processing capacity and an uncompromising ultra-low true cut-through latency.

Delivering the highest feature set at the right price allows you to get the most out of your Ethernet infrastructure to best support a variety of use cases, including media and entertainment; streaming video, financial services industry, virtualized data centers, and next generation storage, including software-defined storage and NVMe® flash. HPE SN3420M switches are available with factory integrated NVIDIA Cumulus® Linux for immediate deployment. With HPE M-Series switches, you can:

- **Optimize Storage**— modernize your network to eliminate limitations and bottlenecks that can be caused by the addition of flash storage.
- **Enjoy efficient network performance**— avoid packet loss, provide predictable performance with line-rate packet delivery across all ports and all packet sizes
- **Realize breakthrough economics**— make better use of your data center resources with the highest port density per rack unit and the industry's lowest power consumption.
- **Accelerate business innovation**— utilize 1/10/40Gbps Ethernet connectivity for existing workloads and enhance connectivity utilizing built-in 25/100 Gbps capabilities to respond quickly to business needs and to stay on the leading edge of Ethernet switching technology.

HPE SN3420M Ethernet switches are based on the high-performance 25GbE NRZ encoding capable Spectrum-2 ASIC and are ideally suited for both top-of-rack leaf and fixed configuration spines with capabilities:

- Ethernet port speeds from 1Gbps to 100Gbps
- 12 100/40GbE QSFP28 ports and 48x1/10/25GbE SFP28 ports
 - All 12 QSFP28 ports support split x4 or x2 mode (12x4=48 breakout ports maximum)
- 42MB of fully-shared packet buffering
- 512K on-chip, dynamically shared forwarding table
- Deep visibility with 512k general purpose counters
- In-band network telemetry

Overview

- Feature-rich layer 2 and layer 3 forwarding
- Ultra-low latency with true cut through performance, Zero packet loss performance with NVMe TCP and RoCEv2 RDMA, DCBX, PFC, ECN support

Models

HPE M-Series SN3420M Ethernet Switch Model

Description

HPE 25GbE 48SFP28 12QSFP28 Power to Connector Airflow Switch SN3420M with NVIDIA Cumulus

SKU
S2T77A

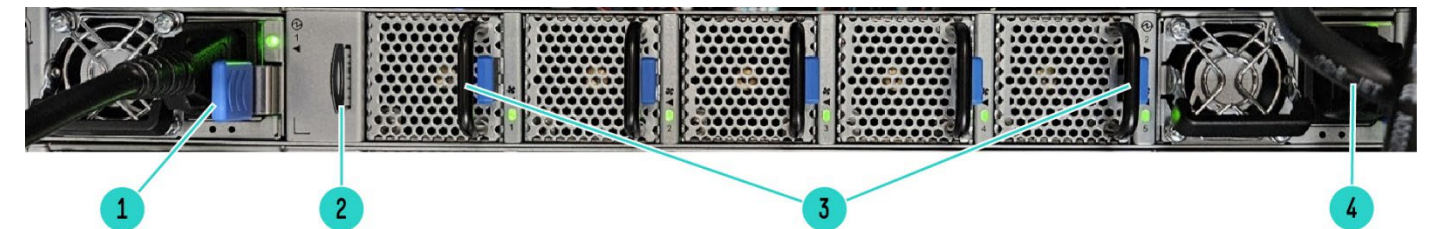
- HPE 48SFP28 12QSFP28 P2C Sw SN3420M w/NVD

Notes: HPE switch SKUs with factory installed NVIDIA Cumulus® Linux cannot be converted to HPE ONYX or ONIE switch SKUS.



HPE M-Series SN3420M - Front (Connector) View

- | | | | |
|---|--|---|--|
| 1 | Status LEDs (System; Fan; Power Supply; Unit Identifier) | 5 | Password Reset Button - refer to the Single User Mode Boot Recovery section in the Cumulus Linux User Guide |
| 2 | USB Type A 3.0 compliant (USB 1.0 not supported) | 6 | 49-56 40/100GbE QSFP28 ports |
| 3 | MGMT0 100Mb/s to 1Gb/s Port | 7 | 1-48 1/10/25GbE SFP28 ports |
| 4 | IOIO RS232 Serial Console Port: 115200 BAUD | | |



HPE M-Series SN3420M – Rear (Power) View

- | | | | |
|---|------------------------------------|---|--------------------------|
| 1 | Power Supply 1 | 3 | FANs (1 left to 5 right) |
| 2 | Inventory Information Pull-Out Tab | 4 | Power Supply 2 |

Notes: the SN3420M chassis, power supplies and fans are included in the field replaceable switch.



Standard Features

Key Features and Benefits

- SN3420M switches are ideal for modern server and storage networks, supporting DAC, AOC, and optical breakout cables. The 48 SFP28 and 12 QSFP28 ports deliver predictable performance and zero packet loss at line-rate across each port and packet size.
- The 12 100GbE NRZ encoded ports in the SN3420M switches can be split to two (split-2) 50GbE ports, or to four (or less) 1/10/25GbE NRZ encoded ports, using a splitter cable.



All 12 QSFP28 ports 49-60 support x4 or x2 breakout

- SN3420M can be deployed to support 1GbE and 10GbE ports, including 10Gbase-T RJ45 transceivers, and is designed to be able to evolve over time to support 25, 40, 50, 100GbE speeds. This helps future-proof your network architecture and allows for implementing significant speed upgrades to the architecture over time.
- SN3420M provide ultra-low cut-through latency port-to-port. This is advantageous for flash storage which moved the latency bottleneck from storage media to the network, as well as for the bursty nature of today's software-defined and cloud-driven data center traffic flows.
- M-series SN3420M provide high port density in a single rack unit, allowing for higher capacity and efficiency, simplifying scale-out environments and saving on total cost of ownership. Unique breakout cables fan out individual switch ports to multiple device ports.
- Provides wire-rate performance with zero packet loss across all frame sizes, avoiding any negative impact on applications that could occur with frame loss as unexpected packet loss is unacceptable in modern data centers, especially within a storage network.
- Capable of forwarding 100% capacity wire rate performance with zero packet loss across all ports concurrently at 25GbE and 100GbE speeds while transferring data across both Layer 2 and Layer 3 networks.
- Designed to use less electric power than competing switches, providing one of the industry's lowest power draws, producing less heat than competing products, providing reduced OPEX cost.
- Provides enough switching bandwidth to transport all ports at 25GbE and 100GbE bandwidth concurrently. This allows the switches to avoid head-of-line blocking which can reduce a switches overall performance and throughput.

NVIDIA® Cumulus Linux

Cumulus Linux is a powerful open network operating system enabling advanced automation, customization and scalability using web-scale principles like the world's largest data centers. It accelerates networking functions and provides choice from an extensive list of supported switch models including NVIDIA Spectrum™ based switches. Cumulus Linux was built for automation, scalability and flexibility, allowing you to build data center and campus networks that ideally suits your business needs. Cumulus Linux is the only open network OS that allows you to build affordable and efficient network operations like the world's largest data center operators, unlocking webscale networking for businesses of all sizes.

The Ideal Solution for Your Network Challenges, Cumulus Linux enables modern data center architectures, while providing a transition path for traditional data center architectures, with support for layer 2, layer 3, and overlay networks. This open approach enables a wide range of solutions. Centralized, remote management of AI deployments enables over-the-air software updates, remote debugging, and system monitoring, as well as other features like self-healing systems. These remote management features make maintenance and upkeep easier, and AI more accessible and practical for locations that are difficult to access or far from headquarters. This results in faster, more comprehensive insights that can drive real-time decisions.



Standard Features

Deployment Models:

L3 network, L2 network, Clos, Out-of-band management, Overlay network,

Use cases:

Containers, Big data, Private cloud, Network virtualization, DevOps / automation, Monitoring and analytics, Hyperconverged infrastructure

Key Features

Unnumbered Interfaces: Automation gets even easier with this simplified Internet Protocol (IP) approach for Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF). All you need is one IP template for leaf nodes and one for spine nodes. BGP Unnumbered, the only difference between a BGP unnumbered configuration and the BGP numbered configuration is that the BGP neighbor is as an interface (instead of an IP address). There is no need to configure an IP address on the interface between the two peers on each side.

Redistribute Neighbor (RDNR): Get virtual machine (VM) and host mobility by plugging your server into any RDNR switch and making it layer-3 discoverable on the fabric. If you need to move the server, there's no need to reconfigure.

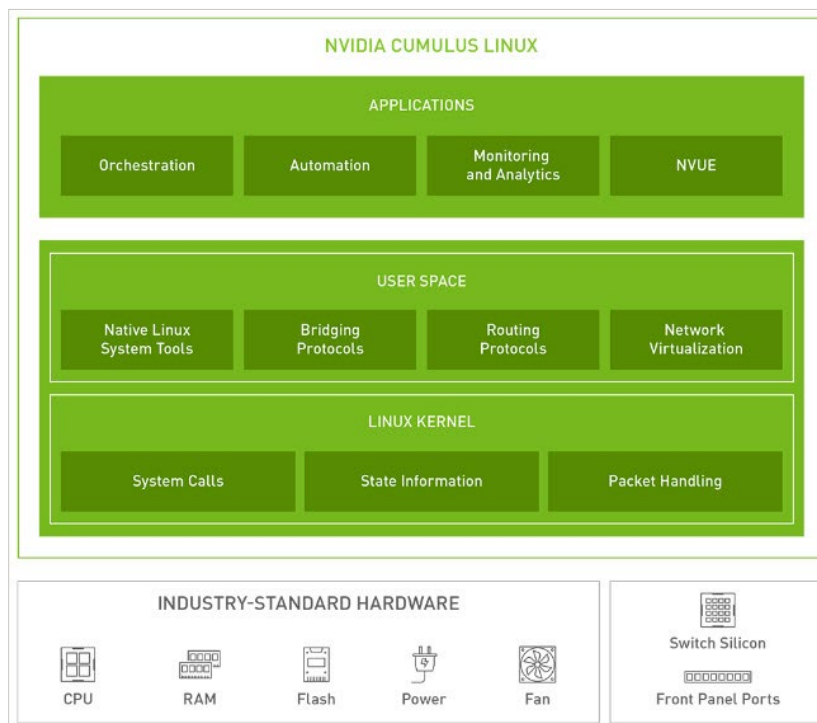
Prescriptive Topology Manager (PTM): Efficiently go from whiteboard to physical cable. With PTM, you can program your data center to verify connections and resolve issues faster.

Virtual Routing and Forwarding: Run multiple network paths without the need for multiple switches, giving you traffic isolation and network segmentation for multiple devices.

Ethernet Virtual Private Networks (EVPN): The most advanced capabilities available for EVPN allow legacy layer-2 applications to operate over next-generation layer-3 networks.

NVIDIA User Experience (NVUE): A full command-line interface (CLI) object model of Cumulus Linux enables advanced programmability, extensibility, and usability.

Digital twins: NVIDIA Air makes physical deployments seamless by validating and simplifying deployments and upgrades in a virtual network.



Standard Features

User Interface: Command Line

With the NVIDIA User Experience (NVUE) object-oriented management tool, Cumulus Linux customers can go beyond the CLI and unify their network management with the rest of the data center management. NVUE enables any APIs to tie into NOS management, including REST, gRPC, RestConf, NetConf, and OpenConfig. Additionally, NVUE is Git-based, enabling Diff, revert, apply, branch, and stash behaviors. NVUE's configuration is simple: one YAML file ties all Linux configurations together, making it easy to copy configurations from switch to switch.

Operating System Install and Upgrade

- Server-style upgrade/patching across minor releases, server-style process restart/termination
- Support for zero touch OS installation using ONIE loaded on industry-standard switches
- Standard mechanism for authentication, authorization, and accounting with TACACS+

Extensibility

- Cumulus Linux works with any language supported in Linux today, including scripting with Bash. Perl. Python, and Ruby

Hardware Management

- The switch hardware abstraction layer accelerates Linux kernel networking constructs in hardware, including the routing table, ARP table, bridge FIB, IP/FIB tables, bonds, VLANs, and VXLAN bridges
- Hardware management also includes jumbo frames support and environmental management
- Forwarding table profiles on the ASIC provide flexible partitioning of resources

Layer 3 Features

- IPv4/v6 routing suite including OSPFv2, OSPFv3, and BGPv4/v6
- RDMA over Converged Ethernet (RoCEv2) support for Layer 2 and Layer 3
- Virtual routing and forwarding (VRF) and VRF route leaking
- Equal-cost multi-path (ECMP) and ECMP resilient hashing for IPv4 and IPv6 traffic
- Bidirectional forwarding detection (BFD) across all platform and interface types, IPv4 and IPv6, BGP and OSPF, VXLAN, BGP conditional route advertisement
- Protocol-independent multicast (PIM, PIM-SM, PIM-SSM)
- Policy-based routing
- Generic routing encapsulation (GRE) tunneling
- Precision time protocol (PTP) Boundary Clock
- VNI scaling: supports 6 bridges with up to 1,000 VNIs
- GTP Hashing
- Adaptive Routing with RoCEv2

Layer 2 features

- Bridge management with STP (IEEE 802.1d). RSTP (IEEE 802.1w), PVST, PVST+, bridge assurance, BPDU guard, and BPDU filter
- VLANs. VLAN trunks (IEEE 802.1q). LACP (IEEE 802.3ad). LACP bypass, unicast/broadcast storm control, LLDP, CDP, IPv6 neighbor discovery, and IPv6 route advertisement
- MLAG (cladg daemon)
- IGMPv2/v3 snooping, MLDv1/v2 snooping, Optimized Multicast Flooding (OMF)
- Virtual router redundancy (VRR - active-active first hop redundancy protocol)
- LLDP DCB IEEE TLVs



Standard Features

Network Virtualization

- VXLAN support
- VXLAN Routing - symmetric and asymmetric
- L2 gateway services integration with VMware NSX
- VXLAN head end replication
- VXLAN active-active bridging with MLAG
- Controller-less network virtualization with EVPN

Management

- Object-oriented API-compatible switch management with NVIDIA User Experience (NVUE)
- ISSU: in-service software upgrades
- Warm boot on bonds
- Native Linux management tools, such as OpenSSH, SCP, and FTPS
- Automated install and provisioning: zero touch install and zero touch provisioning
- Management VRF
- DHCP and v4/v6 DHCP relays
- Authentication with LDAP and authorization with sudo NTP
- Interface configuration management (ifupdown2)
- Advanced management/orchestration through third-party add-on packages
- Snapshot and rollback of the entire system to eliminate risk from system updates

Monitoring and Troubleshooting

- Monitor traffic patterns and preemptive capacity planning with buffer monitoring
- Traditional monitoring with SNMPv2 and SNMPv3 and network-specific MIBs, hardware monitoring via watchdog, analytics with SPAN, ERSPAN, ACL-based counters, DOM optics data, thermal sensors, real time queue-depth, and buffer utilization reporting
 - HPE recommends use of SNMPv3 username and password instead of the read-only community; SNMPv3 does not expose the password in the GetRequest and GetResponse packets and can also encrypt packet contents. You can configure multiple usernames for different user roles with different levels of access to various MIBs.
 - The `/usr/share/snmp/mibs/Cumulus-Snmp-MIB.txt` file defines the overall Cumulus Linux MIB.
- Troubleshooting with `dnsutils`, `syslog`, reachability tools, hardware inventory, log files, server-style filesystem, and Spectrum ASIC commands
- The `cl-support` script generates a compressed archive file (.txz) of useful information for troubleshooting. The system either creates the archive file automatically or you can create the archive file manually. The system creates the `cl-support` archive file automatically for the following reasons:
 - When there is a core dump file for any application (not specific to Cumulus Linux, but something all Linux distributions support), located in `/var/support/core`.
 - When one of the monitored services fails for the first time after you reboot or power cycle the switch
- sFlow monitoring for system statistics and network traffic.
- The NVIDIA Firmware Tool (MFT) `mlxlink` tool is integrated into Cumulus and is used to check and debug link status and related issues. The tool can be used on different links and cables (passive, active, transceiver and backplane).
- What Just Happened (WJH) provides real time visibility into network problems and has two components:
 - The WJH agent enables you to stream detailed and contextual telemetry for off-switch analysis with tools such as NVIDIA NetQ.
 - The WJH service (what-just-happened) enables you to diagnose network problems by looking at dropped packets. WJH can monitor layer 1, layer 2, layer 3, tunnel, buffer and ACL related issues. Cumulus Linux enables and runs the WJH service by default.



Standard Features

- You can choose which packet drops you want to monitor and show by creating channels and setting the packet drop categories (layer 1, layer 2, layer 3, tunnel, buffer and ACL) you want to monitor.
- NVUE does not provide commands to set the buffer and ACL packet drop categories. You must edit the /etc/what-just-happened/what-just-happened.json file. NVUE supports WJH show commands.

Security

- Access control lists (ACLs) L2-L4 classification through IP/EP tables and CPU protection through hardware enforced ACL-based rate limiting DoS control
- Authenticate and authorize attached devices with 802.1x
- Kernel Address Space Randomization

QoS

- Link PAUSE
- Classification based on Class of Service (CoS) (IEEE 802.1p) or DSCP (queuing, scheduling-DWRR and Strict Priority - and buffer allocation)
- Ingress ACL-based classification/policing
- Priority flow control and explicit congestion notification (ECN)
- Dynamic buffer configuration as default

Cumulus VX and NVIDIA AIR

- Getting started with Cumulus Linux is easy. Customers can explore, test, and prototype the technology.
 - To build a data center digital twin or mock-up a deployment with Cumulus, visit [NVIDIA Air](#)
- Download the free NVIDIA Cumulus VX appliance, supported virtual appliance to test and stage production rollouts

Third-party Packages

- Orchestration: Ansible, CFEngine, Chef, and Puppet
- Monitoring: Collectd, Ganglia, Graphite, Nagios/Icinga, and NetSNMP
- You can use Cumulus Linux to run the Docker container platform. You can install Docker Engine directly on a Cumulus Linux switch and run Docker containers natively on the switch.



Service and Support

Warranty

(3-3-3) Hardware Warranty; 3-year parts; 3-year on-site (standard business hours, next business day response) and 3-year labor.

Notes: The hardware warranty covers firmware. For extended hardware support and installation information, please see the “Services and Support” Section.

HPE Services

No matter where you are in your 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>

Recommended Services

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 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/completecure>



Service and Support

Other related services from HPE Services

HPE Installation and Start-up Service

Provides for the hardware installation and startup of HPE branded M-Series switches with operating system, according to the product specifications. The HPE service delivery technician will assist you in bringing your new hardware into operation in a timely and professional manner.

<https://www.hpe.com/h20195/v2/Getdocument.aspx?docname=a00025816enw>

HPE Hardware Installation

Provides for the basic hardware installation of HPE branded M-Series ONIE switches to assist you in bringing your new hardware into operation in a timely and professional manner.

<https://www.hpe.com/h20195/v2/Getdocument.aspx?docname=5981-9356enw>

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.

<https://www.hpe.com/services/lifecycle>

- For a list of the most frequently purchased services using service credits, see the [HPE Service Credits Menu](#)

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.



Service and Support

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 QuickSpecs, 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/>

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>



Configuration Information

Model Description

HPE M-series SN3420M (P2C) 48SFP28 12QSFP28 Switch

- 2 x 550w Power Supply with Intake Fan
- 5 x Intake Fan Tray
- 1 x Fixed Rack Mount Kit MTEF-KIT-SP Standard: 23.6"-31.5" (60 to 80 cm)
- 2 x Power cord, 1.83m, C14-C13
- 1 x Serial cable (DB9 to RJ45)
- 2 x PSU CABEL RETAINER KIT
- 1 x HPE Warranty and Installation instructions

Notes: SN3420M with NVIDIA Cumulus® Linux requires transceivers listed below

Step 1 - Base Configuration. Select one Model:

Description

HPE 25GbE 48SFP28 12QSFP28 Power to Connector Airflow Switch SN3420M with NVIDIA Cumulus

SKU

S2T77A

- HPE 48SFP28 12QSFP28 P2C Sw SN3420M NVD, AC power

Step 2 – Options

Refer to HPE M-Series Switches **SPOCK** Connectivity Stream for latest SN3420M with NVIDIA Cumulus® Linux interconnect support matrix

Transceivers – for M-series SN3420M NVIDIA Cumulus® Linux switches

Notes #	Descriptions	SKU
	HPE 100GbE QSFP28 SR4 100m Transceiver	Q2F19A
	HPE 100GbE QSFP28 500m 1310mm PSM4 Transceiver	Q8J73A
5, 6	HPE 100Gb QSFP28 LC SWDM4 Multi-mode 100m Transceiver	R0R40A
2	HPE 10GbE SFP+ SR Multi-mode 300m Transceiver	Q6M30A
2	HPE 10Gb SFP+ Short Wave 1-pack Pull Tab Optical Transceiver	Q2P65A
2,7,8	HPE 10GBASE-T SFP+ RJ45 30m 1-pack Transceiver	R0R41B
2	HPE 25Gb SFP28 SR 30m Transceiver	R0R42B
2	HPE BladeSystem c-Class 10Gb SFP+ SR Transceiver	455883-B21
2	HPE Aruba Networking 1G SFP RJ45 T 100m Cat5e TAA Transceiver	R9Q45A
2	HPE Aruba Networking 10G SFP+ LC LR 10km SMF Transceiver	J9151E
4	HPE 7.6m/25ft CAT5 RJ45 M/M Ethernet C/O Cable	C7539A
2	HPE Networking X130 10G SFP+ LC SR Transceiver	JD092B
2	HPE Networking X130 10G SFP+ LC LR Transceiver	JD094B
2	HPE Networking X130 10G SFP+ LC ER 40km Transceiver	JG234A
2	HPE Aruba Networking SFP-10GE-ZR 10GBASE-ZR SFP+ 1310nm LC Connector Pluggable 10GbE Transceiver	JW148A

Configuration Information

Notes #	Descriptions	SKU
1, 2	HPE Networking X130 10G SFP+ LC LH 80km Transceiver	JG915A
	HPE Networking X140 40G QSFP+ CSR4 300m Transceiver	JG709A
	HPE Networking X140 40G QSFP+ LC LR4L 2km SM Transceiver	JL286A
	HPE Networking X140 40G QSFP+ LC LR4 SM 10km 1310nm Transceiver	JG661A
	HPE BladeSystem c-Class 40Gb QSFP+ MPO SR4 100m Transceiver	720187-B21
	HPE Networking X140 40G QSFP+ LC BiDi 100m MM Transceiver	JL251A
	HPE 40Gb QSFP+ Bidirectional Transceiver	841716-B21
2	HPE Networking X190 25G SFP28 LC SR 100m MM Transceiver	JL293A
2	HPE 25Gb SFP28 Short Wave Extended Temperature 1-pack Pull Tab Optical Transceiver	Q2P64B
	HPE Networking X140 40G QSFP+ MPO SR4 Transceiver	JG325B
	HPE Aruba Networking 40G QSFP+ MPO SR4 Transceiver	R9F97A
2	HPE 25Gb SFP28 SR 100m Transceiver	845398-B21
2	HPE Aruba Networking 25G SFP28 LC LR 10km SMF Transceiver	JL486A
2	HPE QSFP28 to SFP28 Adapter	845970-B21
	HPE Networking X150 100G QSFP28 LC LR4 10km SM Transceiver	JL275A
	HPE Networking X150 100G QSFP28 CWDM4 2km SM Transceiver	JH673A
5, 6	HPE Networking X150 100G QSFP28 LC SWDM4 100m MM Transceiver	JH419A
5	HPE 100Gb QSFP28 Bidirectional Transceiver	845972-B21
12	HPE 100GbE QSFP28 LC DR1 500m 1-pack Transceiver	R8M61A
	HPE Alletra 6000 2x100Gb QSFP28 MPO SR4 100m FIO Transceiver	R7D08A
	HPE Alletra 6000 2x100Gb QSFP28 MPO SR4 100m Transceiver	R7D12A
2	HPE Alletra 6000 2x10Gb SFP+ SR FIO Transceiver	R7D05A
2	HPE Alletra 6000 2x10Gb SFP+ SR Transceiver	R7D09A
2	HPE Alletra 6000 2x25Gb SFP28 SR 100m FIO Transceiver	R7D07A
2	HPE Alletra 6000 2x25Gb SFP28 SR 100m Transceiver	R7D11A

Direct Attach Copper Cables (DAC)

Notes #	Descriptions	SKU
	HPE Networking Comware X240 40G QSFP+ QSFP+ 1m Direct Attach Copper Cable	JG326A
	HPE Networking Comware X240 40G QSFP+ QSFP+ 3m Direct Attach Copper Cable	JG327A
	HPE Networking Comware X240 40G QSFP+ QSFP+ 5m Direct Attach Copper Cable	JG328A
3	HPE 100GbE QSFP28 to 4x25GbE SFP28 1m Direct Attach Copper Cable	Q9S72A
3	HPE Networking Comware X240 40G QSFP+ to 4x10G SFP+ 1m Direct Attach Copper Splitter Cable	JG329A
3	HPE Networking Comware X240 40G QSFP+ to 4x10G SFP+ 3m Direct Attach Copper Splitter Cable	JG330A
3	HPE Networking Comware X240 40G QSFP+ to 4x10G SFP+ 5m Direct Attach Copper Splitter Cable	JG331A
	HPE BladeSystem c-Class 40G QSFP+ to QSFP+ 3m Direct Attach Copper Cable	720199-B21
	HPE BladeSystem c-Class 40G QSFP+ to QSFP+ 5m Direct Attach Copper Cable	720202-B21
	HPE Networking X242 40G QSFP+ to QSFP+ 1m Direct Attach Copper Cable	JH234A
	HPE Networking X242 40G QSFP+ to QSFP+ 3m Direct Attach Copper Cable	JH235A
	HPE Networking X242 40G QSFP+ to QSFP+ 5m Direct Attach Copper Cable	JH236A
3	HPE Networking X240 QSFP28 4xSFP28 1m Direct Attach Copper Cable	JL282A

Configuration Information

Notes #	Descriptions	SKU
3	HPE Networking X240 QSFP28 4xSFP28 3m Direct Attach Copper Cable	JL283A
3	HPE Networking X240 QSFP28 4xSFP28 5m Direct Attach Copper Cable	JL284A
2	HPE Aruba Networking 25G SFP28 to SFP28 0.65m Direct Attach Cable	JL487A
2	HPE Aruba Networking 25G SFP28 to SFP28 3m Direct Attach Copper Cable	JL488A
2	HPE Aruba Networking 25G SFP28 to SFP28 5m Direct Attach Copper Cable	JL489A
2	HPE Aruba Networking 25G SFP28 to SFP28 3m Direct Attach Copper Cable	R9F92A
2	HPE Aruba Networking 25G SFP28 to SFP28 5m Direct Attach Copper Cable	R9F93A
	HPE Aruba Networking 100G QSFP28 to QSFP28 1m Direct Attach Copper Cable	R0Z25A
	HPE Aruba Networking 100G QSFP28 to QSFP28 5m Direct Attach Copper Cable	R0Z26A
2	HPE 25Gb SFP28 to SFP28 0.5m Direct Attach Copper Cable	R4G18A
2	HPE 25Gb SFP28 to SFP28 1m Direct Attach Copper Cable	R4G19A
2	HPE 25Gb SFP28 to SFP28 3m Direct Attach Copper Cable	844477-B21
2	HPE 25Gb SFP28 to SFP28 5m Direct Attach Copper Cable	844480-B21
2	HPE Aruba Networking 100G QSFP28 to QSFP28 3m Direct Attach Copper Cable	JL307A
	HPE Aruba Networking 100G QSFP28 to QSFP28 1m Direct Attach Copper Cable	R9F77A
	HPE Aruba Networking 100G QSFP28-QSFP28 3m Direct Attach Copper Cable	R9F74A
	HPE Aruba Networking 100G QSFP28 to QSFP28 5m Direct Attach Copper Cable	R9F78A
	HPE 100Gb QSFP28 to 4x25Gb SFP28 3m Direct Attach Copper Cable	845416-B21
	HPE 100Gb QSFP28 to QSFP28 5m Direct Attach Copper Cable	845408-B21
	HPE Networking X240 100G QSFP28 1m DAC Cable	JL271A
	HPE Networking X240 100G QSFP28 3m DAC Cable	JL272A
	HPE Networking X240 100G QSFP28 5m DAC Cable	JL273A
	HPE Aruba Networking 10G SFP+ to SFP+ 7m Direct Attach Copper Cable	J9285D
	HPE 100Gb QSFP28 to QSFP28 3m Direct Attach Copper Cable	845406-B21
	HPE 100Gb QSFP28 to QSFP28 0.5m Direct Attach Copper Cable	R8M59A
2	HPE Alletra 6000 2x10Gb SFP+ to SFP+ 3m Direct Attach Copper Cable	R7D16A
2	HPE Alletra 6000 2x25Gb SFP28 to SFP28 3m Direct Attach Copper Cable	R7D17A
	HPE Alletra 6000 2x100Gb QSFP28 to QSFP28 3m Direct Attach Copper Cable	R7D18A
	HPE Aruba Networking 10G SFP+ to SFP+ 3m Direct Attach Copper Cable	R9F84A

Active Optical Cable (AOC)

Notes #	Descriptions	SKU
3	HPE 40GbE QSFP+ to 4x10GbE SFP+ 5m Active Optical Cable	Q9S66A
2	HPE 25GbE SFP28 to SFP28 3m Smart Active Optical Cable	Q9S67A
2	HPE 25GbE SFP28 to SFP28 5m Smart Active Optical Cable	Q9S68A
2	HPE 25GbE SFP28 to SFP28 10m Smart Active Optical Cable	Q9S69A
2	HPE 25GbE SFP28 to SFP28 15m Smart Active Optical Cable	Q9S70A
	HPE 100GbE QSFP28 to QSFP28 5m Active Optical Cable	Q9S71A
	HPE Storage 100GbE QSFP28 to QSFP28 10m Extended Temperature Active Optical Cable	S2T39A
	HPE Storage 100GbE QSFP28 to QSFP28 15m Extended Temperature Active Optical Cable	S2T40A
	HPE Storage 100GbE QSFP28 to QSFP28 25m Extended Temperature Active Optical Cable	S2T41A

Configuration Information

Notes #	Descriptions	SKU
	HPE Aruba Networking 100G QSFP28 to QSFP28 2m Active Optical Cable	R9F76A
	HPE Aruba Networking 100G QSFP28 to QSFP28 7m Active Optical Cable	R9F79A
	HPE Aruba Networking 100G QSFP28 to QSFP28 15m Active Optical Cable	R9F80A
	HPE Aruba Networking 100G QSFP28 to QSFP28 30m Active Optical Cable	R9F81A
	HPE BladeSystem c-Class 40G QSFP+ to QSFP+ 15m Active Optical Cable	720211-B21
	HPE BladeSystem c-Class QSFP+ to 4x10G SFP+ 15m Active Optical Cable	721076-B21
	HPE 100Gb QSFP28 to QSFP28 7m Active Optical Cable	845410-B21
	HPE 100Gb QSFP28 to QSFP28 15m Active Optical Cable	845414-B21
	HPE QSFP28 to 4x25Gb SFP28 7m Active Optical Cable	845420-B21
	HPE QSFP28 to 4x25Gb SFP28 15m Active Optical Cable	845424-B21

M-series 200GbE pluggables – DAC/AOC/Optical transceivers

Notes #	Descriptions	SKU
9	HPE 200Gb QSFP56 to QSFP56 0.5m Direct Attach Copper Cable	R5Z76A
9	HPE 200Gb QSFP56 to QSFP56 1m Direct Attach Copper Cable	R5Z77A
9	HPE 200Gb QSFP56 to QSFP56 2m Direct Attach Copper Cable	R5Z78A
9	HPE 200Gb QSFP56 to QSFP56 2.5m Direct Attach Copper Cable	R5Z79A
9	HPE 200Gb QSFP56 to QSFP56 5m Active Optical Cable	R5Z80A
9	HPE 200Gb QSFP56 to QSFP56 10m Active Optical Cable	R5Z81A
9	HPE 200Gb QSFP56 to QSFP56 15m Active Optical Cable	R5Z82A
9,10	HPE 200Gb QSFP56 MPO SR4 100m Transceiver	R5Z83A
9,11	HPE 200Gb QSFP56 LC CWDM4 FR4 Transceiver	R5Z84A
9,14	HPE 200Gb QSFP56 to 2x100Gb QSFP56 3m Active Optical Cable	R6F24A
9,14	HPE 200Gb QSFP56 to 2x100Gb QSFP56 5m Active Optical Cable	R6F25A
9,14	HPE 200Gb QSFP56 to 2x100Gb QSFP56 15m Active Optical Cable	R6F26A
9,13	HPE 200Gb QSFP56 to 4x50/25Gb SFP56 2.5m Direct Attach Copper Cable	R6F27A
14	HPE 200GbE QSFP56 to 2xQSFP56 2m Direct Attach Copper Cable	R8M57A
14	HPE 200GbE QSFP56 to 2xQSFP56 2.5m Direct Attach Copper Cable	R8M58A
13	HPE 400GbE QSFP-DD to 4xQSFP56 1m Direct Attach Copper Cable	R8M55A
13	HPE 400GbE QSFP-DD to 4xQSFP56 2m Direct Attach Copper Cable	R8M56A

Notes:

- 1. JG915A - Storage connectivity support for this transceiver is limited to 40Km
- 2. HPE QSA28 (QSFP28 to SFP28) adapter (845970-B21) is compatible with all M-Series switches and is required with this transceiver to convert a QSFP+/QSFP28 slot to a single SFP+/SFP28 slot for 25G or 10G operation with this switch model.
- 3. The SN2010M 25G DAC connectivity to another M-Series switch or 3rd party switch is limited to a 0.5m DAC cable. End device connectivity may use 1m or a 3m (26 gauge) DAC.
- 4. This RJ45 crossover cable is compatible and supported for use when directly connecting the two M-Series switch MGMT ports. When configuring MLAG and also utilizing in-band management, the MGMT0 ports of the two switches shall be connected (with MLAG the MGMT0 ports shall be reachable from each switch).
- 5. The 845972-B21 HPE 100Gb QSFP28 Bidirectional XCVR does not interoperate with the JH419A and R0R40A HPE 100Gb QSFP28 LC SWDM4 MM 100m transceivers.
- 6. 100GbE SWDM4 LC transceivers JH419A and R0R40A are interoperable.

Configuration Information

- 7. 10Gbase-T SFP+ RJ45 transceiver supports maximum length 30M CAT6a cable. This 10Gbase-T transceiver is not qualified for use at 1GbE and shall be operated only at 10GbE.
- 8. Cumulus 5.9 and later revisions are required for support of R0R41B in the SN3420M and with a QSA28 in SN3420M, SN3700CM, SN3700M, and SN4600cM switches. Prior to Cumulus 5.9, SN3420M limits R0R41B support to Ports 1-6, Cumulus 5.9 expands support to ports 1-48. R0R41B transceiver supports maximum length 30M CAT6a cable. R0R41B 10Gbase-T transceiver is not qualified for use at 1GbE and shall be operated only at 10GbE.
- 9. Cumulus 5.6 and later revisions are required for support of the QSA28 and QSFP56 200G pluggables with the SN3700M, SN3700cM, SN3420M, SN2100M, SN2010M and SN4600cM switches
- 10. R5Z83A operates at 200GbE, 100GbE, and 4x50GbE breakout. Operation at 4x25/100GbE breakout is not supported. Use a 100GbE transceiver for this breakout application.
- 11. R5Z84A requires a single mode duplex LC cable
- 12. The R8M61A class 6 power requirements limit use in the SN3700cM to ports 1, 2, 31, 32; SN3700M has no power transceiver class 6 port use limit.
- 13. All cable ports must be set as NRZ or PAM4. A mix between the two technologies is not supported.
- 14. All cable ports must be set as PAM4, NRZ is not supported.

Supported Optical Cables for all M-Series switch models

Description	SKU
HPE Premier Flex MPO/MPO Multi-mode OM4 12 Fiber 10m Cable	QK729A
HPE Premier Flex MPO/MPO Multi-mode OM4 8 Fiber 50m Cable	QK731A
HPE Premier Flex LC/LC Multi-mode OM4 2 Fiber 1m Cable	QK732A
HPE Premier Flex LC/LC Multi-mode OM4 2 Fiber 2m Cable	QK733A
HPE Premier Flex LC/LC Multi-mode OM4 2 Fiber 5m Cable	QK734A
HPE Premier Flex LC/LC Multi-mode OM4 2 Fiber 15m Cable	QK735A
HPE Premier Flex LC/LC Multi-mode OM4 2 Fiber 30m Cable	QK736A
HPE Premier Flex LC/LC Multi-mode OM4 2 Fiber 50m Cable	QK737A
HPE LC to LC Multi-mode OM3 2-Fiber 0.5m 1-Pack Fiber Optic Cable	AJ833A
HPE LC to LC Multi-mode OM3 2-Fiber 1.0m 1-Pack Fiber Optic Cable	AJ834A
HPE LC to LC Multi-mode OM3 2-Fiber 2.0m 1-Pack Fiber Optic Cable	AJ835A
HPE LC to LC Multi-mode OM3 2-Fiber 5.0m 1-Pack Fiber Optic Cable	AJ836A
HPE LC to LC Multi-mode OM3 2-Fiber 15.0m 1-Pack Fiber Optic Cable	AJ837A
HPE LC to LC Multi-mode OM3 2-Fiber 30.0m 1-Pack Fiber Optic Cable	AJ838A
HPE LC to LC Multi-mode OM3 2-Fiber 50.0m 1-Pack Fiber Optic Cable	AJ839A
HPE Premier Flex MPO/MPO OM4 100m Cable	H6Z30A
HPE Multi Fiber Push On to 4 x Lucent Connector 5m Cable	K2Q46A
HPE Multi Fiber Push On to 4 x Lucent Connector 15m Cable	K2Q47A
HPE Premier Flex MPO/MPO Multi-mode OM4 12 Fiber 1m Cable	Q1H63A
HPE Premier Flex MPO/MPO Multi-mode OM4 12 Fiber 2m Cable	Q1H64A
HPE Premier Flex MPO/MPO Multi-mode OM4 12 Fiber 5m Cable	Q1H65A
HPE Premier Flex MPO/MPO Multi-mode OM4 12 Fiber 15m Cable	Q1H66A
HPE Premier Flex MPO/MPO Multi-mode OM4 12 Fiber 30m Cable	Q1H67A
HPE Premier Flex MPO to 4xLC 30m Cbl	Q1H68A
HPE Premier Flex MPO to 4 x Lucent Connector 50m Cable	Q1H69A

Configuration Information

Description	SKU
HPE 5m Single-Mode LC/LC Fibre Channel Cable	AK346A
HPE Premier Flex MPO12 to MPO12 APC Single-mode 50m 1-pack Cable	R6F28A
HPE Premier Flex MPO12 to MPO12 APC Single-mode 100m 1-pack Cable	R6F29A
HPE Premier Flex MPO12 to MPO12 APC Single-mode 300m 1-pack Cable	R6F30A
HPE Premier Flex MPO12 to MPO12 APC Single-mode 500m 1-pack Cable	R6F31A
HPE Premier Flex Wide Band OM5 LC 2 Fibers 100m Cable	R9M64A
HPE Premier Flex Wide Band OM5 LC 2 Fibers 150m Cable	R9M65A
HPE Premier Flex Wide Band MPO OM5 8 Fibers 100m Cable	R9M62A
HPE Premier Flex MPO16 to 2xMPO8 OM4 5m Cable	S1H57A
HPE Premier Flex MPO16 to 2xMPO8 OM4 10m Cable	S1H58A
HPE Storage Premier Flex MPO16 to 4xMPO8 OM4 5m Cable	S2T36A
HPE Storage Premier Flex MPO16 to 4xMPO8 OM4 10m Cable	S2T37A
HPE Premier Flex MPO16 to MPO16 OM4 2m Cable	R4D51A
HPE Premier Flex MPO16 to MPO16 OM4 5m Cable	R4D52A
HPE Premier Flex MPO16 to MPO16 OM4 10m Cable	R4D53A
HPE Premier Flex MPO16 to MPO16 OM4 15m Cable	R4D54A
HPE Premier Flex MPO16 to MPO16 OM4 30m Cable	R4D55A
HPE Premier Flex MPO16 to 8xLC OM4 5m Cable	R4D56A
HPE Premier Flex MPO16 to 8xLC OM4 10m Cable	R4D57A
HPE Premier Flex MPO16 to 8xLC OM4 15m Cable	R4D58A
HPE Premier Flex MPO16 to 8xLC OM4 30m Cable	R4D59A
HPE MPO12 to 4xLC Single-mode 15m Fibre Cable	R4D60A
HPE MPO12 to 4xLC Single-mode 30m Fibre Cable	R4D61A
HPE MPO12 to 4xLC Single-mode 2m Fibre Cable	R4D62A
HPE MPO12 to 4xLC Single-mode 5m Fibre Cable	R4D63A
HPE MPO12 to 4xLC Single-mode 10m Fibre Cable	R4D64A



Technical Specifications

Family Information (M-series with NVIDIA Cumulus® Linux)				
	HPE SN2010M 100GbE x4QSFP28 + x18 1/10/25GbE switch S2T75A	HPE SN2100M 100GbE 16QSFP28 switch S2T76A	HPE SN3420M 100GbE x12QSFP28 + x48 1/10/25GbE switch S2T77A	HPE SN3700cM 100GbE 32QSFP28 switch S2T78A
Description	Ideal ½ width ToR 1/10/25/40/100GbE	Ideal ½ width ToR 1/10/25/40/100GbE	1/10/25GbE ToR 40/100GbE	40/100GbE ToR
Ports Speeds	18 x 1/10/25GbE + 4x40/100GbE Breakout - 16x1/10/25GbE	16x40/100GbE Breakout - 64x1/10/25GbE	48x1/10/25GbE + 12x40/100GbE Breakout - 48x1/10/25GbE	32x40/100GbE Breakout - 128x1/10/25GbE
Minimum Configuration	18 + 4 Ports	16 Ports	48 + 12 Ports	32 ports
Size	1U (½ 19" wide)	1U (½ 19" wide)	1U	1U
Switching Capacity	1.7Tb/s	3.2Tb/s	4.8Tb/s	6.4Tb/s
Processing Capacity	1.26Bpps	2.38Bpps	3.58Bpps	4.76Bpps
Forwarding Technology	Cut Through	Cut Through	Cut Through	Cut Through
Latency	300ns	300ns	425ns	425ns
Typical Power Consumption (ATIS)	66W	74W	202W	242W
Energy Efficiency	full load: 91.3% @ 115Vac/60Hz, 92.6% @ 230Vac/50Hz	full load: 91.3% @ 115Vac/60Hz , 92.6% @ 230Vac/50Hz	80 Plus Gold	80 Plus Gold
Supported NOS**	NVIDIA Cumulus® Linux	NVIDIA Cumulus® Linux	NVIDIA Cumulus® Linux	NVIDIA Cumulus® Linux
System Memory	8GB	8GB	8GB	8GB
SSD Memory	16GB	16GB	32GB	32GB
Packet Buffer	16MB	16MB	42MB	42MB
1GbE Mgmt Port	1 RJ45	1 RJ45	1 RJ45	1 RJ45
Serial Port	1 RJ45	1 RJ45	1 RJ45	1 RJ45
USB Port	1 Mini USB 2.0	1 Mini USB 2.0	1 Type A USB 3.0	1 Micro USB 2.0
Airflow	Power-to-Connector (P2C) airflow	Power-to-Connector (P2C) airflow	Power-to-Connector (P2C) airflow	Power-to-Connector (P2C) airflow
Power Supplies	2 (1+1 redundant) not replaceable	2 (1+1 redundant) not replaceable	2 (1+1 redundant)	2 (1+1 redundant)
Fans	2 fans not replaceable	2 fans not replaceable	5 (N+1 redundant)	4 (N+1 redundant)
Power Supplies with integrated fans	Frequency: 50-60Hz Input range: 100-264Vac Input current 4.5A-2.9A IEC C14 To C13	Frequency: 50-60Hz Input range: 100-264Vac Input current 4.5A-2.9A IEC C14 To C13	Frequency: 50-60Hz Input range: 100-264Vac Input current 7.1A-2.8A IEC C14 To C13	Frequency: 50-60Hz Input range: 100-264Vac Input current 13A-7A IEC C14 To C13
Size	1.72" x 7.87" x 20" (43.9mm x 200mm x 508mm)	1.72" x 7.87" x 20" (43.9mm x 200mm x 508mm)	1.72" x 17.24" x 18.29" (44mm x 438mm x 464.6mm)	1.72" x 16.84" x 22" (44mm x 428mm x 559mm)
Weight	4.5kg (9.9lbs)	4.54kg (10lbs)	8.5kg (18.8lbs)	14kg (30.9lbs)

Technical Specifications

Model	HPE SN4600cM 100GbE 64QSFP28 Switch S2T80A	HPE SN3700M 200GbE 32QSFP56 Switch S2T79A	HPE SN4700M 400GbE 32QSFP-DD Switch S2T81A
Description	40/100GbE Aggregation/ToR/EoR/ Super Spine	40/50/100/200GbE Aggregation/ToR/ Super Spine	40/50/100/200/400GbE Aggregation/ToR/EoR/ Super Spine
Ports Speeds	64 x40/100GbE Breakout – 128 x 1/10/25GbE	32x40/100/200GbE Breakout – 128x1/10/25/50GbE	32x40/100/200/400GbE Breakout – 128x1/10/25/50GbE
Minimum Configuration	64 QSFP28 Ports	32 QSFP56 Ports	32 QSFP-DD Ports
Size	2U	1U	1U
Switching Capacity	12.8Tb/s	12.8Tb/s	25.6Tb/s
Processing Capacity	8.4Bpps	8.33Bpps	8.4Bpps
Forwarding Technology	Cut Through	Cut Through	Cut Through
Latency	500ns	425ns	500ns
Typical Power Consumption (ATIS)	466 watts	250 watts	630 watts
Supported NOS**	NVIDIA Cumulus® Linux	NVIDIA Cumulus® Linux	NVIDIA Cumulus® Linux
System Memory	8GB	8GB	16GB
SSD Memory	30GB	32GB	60GB
Packet Buffer	64MB	42MB	64MB
1GbE Mgmt Port	1 RJ45	1 RJ45	1 RJ45
Serial Port	1 RJ45	1 RJ45	1 RJ45
USB Port	1 Type A USB 2.0	1 Micro USB 2.0	1 Type A USB 2.0
Airflow	Power-to-Connector airflow (P2C) airflow	Power-to-Connector airflow (P2C) airflow	Power-to-Connector airflow (P2C) airflow
Power Supplies	2 (1+1 redundant)	2 (1+1 redundant)	2 (1+1 redundant)
Energy Efficiency	80 Plus Gold	80 Plus Gold	80 Plus Gold
Fans	3 (N+1 redundant)	6 (N+1 redundant)	6 (N+1 redundant)
Power Supplies with integrated fans	Frequency: 50-60Hz Input range: 100-264Vac Input current 10A-6A @200Vac-240Vac C14 TO C15 EUR + CCC C14 TO C15 UL	Frequency: 50-60Hz Input range: 100-264Vac Input current 13A-7A @200Vac- 240Vac IEC C14 To C13	Frequency: 50-60Hz Input range: 100-264Vac Input current 15A-10A @200Vac-240Vac C14 TO C15 EUR + CCC C14 TO C15 UL
Size	3.46" x 16.85" x 22.3" (88mm x 428mm x 568.5mm)	1.72" x 16.84" x 22" (44mm x 428mm x 559mm)	1.72" x 16.85" x22.3" (44mm x 428mm x 568.5mm)
Weight	14.64kg (32.3lbs)	14kg (30.9lbs)	11.6kg (25.6lbs)

Notes: **There are separate M-Series SKUs for ONIE, ONYX™ and NVIDIA Cumulus® switch models.

Technical Specifications

Environment

Operating temperature	0°C to 40°C Notes: SN4700M 0°C to 35°C
Non-Operating temperature	-40°C to 70°C
NEBs and ETSI operating temperature	-5°C to 55°C
Operating relative humidity (operational)	
Noncondensing	10% to 85%
Operating Altitude	0 - 3050m
Compliant	RoHS
Safety/ EMC	CB, cTUVus, CE, CU, S_Mark, CE, FCC, VCCI, ICES, RCM, BSMI, KCC, CCC

Electrical characteristics

Frequency	50/60 Hz
Voltage	100-264 Vac

Ethernet Ports Maximum High Power Support

- **SN3420M**
 - Ports 1-6 Power Class 3 (2.5W)
 - Ports 7-48 Power Class 1 (1.5W)
 - Ports 49-52, 54, 56, 58, 60 Power Class 4 (3.5W)
 - Ports 53, 55, 57, 59 Power Class 7 (5W)
- **SN3700cM**
 - Ports 1, 2, 31, 32 Power class 7 (5W)
 - Ports 3-30 Power class 4 (3.5W)
- **SN3700M**
 - Ports 1-32 Power class 7 (5W)
 - Ports 1, 2, 21, 22 Power class 8 (6.5W)
- **SN4600cM**
 - Ports 1-48 Power Class 4 (3.5W)
 - Ports 49-64 Power Class 7 (5W)
- **SN4700M**
 - All Ports 1-32 Power Class 48 (12W) – QSFP-DD new method Watt *4
- **SN2010M**
 - Ports 1,2,17,18 Power Class 3 (2.5W)
 - Ports 3-16 Power Class1 (1.5W)
 - Ports 19-22 Power Class (4.5W)
- **SN2100M**
 - Ports 1,2,15,16 Power Class 7 (5W)
 - Ports 3-14 Power Class 4 (3.5W)

Technical Specifications

Safety CB, CE, cTUVus, CU

EN 60950-1:2006+A11:2009+A1:2010+A12:2011+AC:2011+A2:2013, IEC 60950-1:2005 (Second Edition) + AMD 1:2009 + AMD 2:2013 and EN 60950-1:2006 + A11:2009 + A1:2010 + A12:2011 + AC:2011 + A2:2013, UL 60950-1:2007 R10.14, CAN/CSA C22.2 No. 60950 -1-07+A1:2011+A2:2014, IEC 60950-1 Ed. 2.0 :2005 + Am 1:2009+ Am 2:2013, LV CU TR 004/2011 and EMC CU TR 020/2011

EMC CE, ICES, FCC, RCM, VCCI

EN 55032:2012 + AC(13) class A, EN 55024:2010, EN 61000-3-2:2014, EN 61000-3-3: 2013, EN 61000-4-2: 2002, EN 61000-4-3: 2006+A1(08)+A2(10), EN 61000-4-4: 2004+A1(10), EN 61000-4-5: 2006, EN 61000-4-6: 2014 EN 61000-4-11:2004, CFR 47, FCC Part 15:2017, Sub-part B:2019 Class A, ICES-003, Issue 6: 2016 Class A, VCCI-CISPR 32: 2016, Class A, AZ/NZS CISPR 32:2015 Class A, KN22:2009 class A/ KN24:2009

Acoustic

High-speed fan:

- SN3420M – 67.6 dB(A)

Typical power with passive cables (ATIS):

- SN3420M
 - 689.25 BTUs per Hour (202 watts)

HPE Power Advisor

To address a need to accurately estimate power requirements and to ensure the appropriate levels of power and cooling and power-related operating costs, HPE created the **HPE Power Advisor utility**. The HPE Power Advisor utility provides accurate and meaningful estimates of the power needs for HPE servers, storage, and switches including M-series Ethernet switches.

Standards

- 802.1D Bridging and Spanning Tree
- 802.1p QOS
- 802.1Q VLAN Tagging
- 802.1w Rapid Spanning Tree
- 802.1s Multiple Spanning Tree Protocol
- 802.1AB Link Layer Discovery Protocol
- 802.1Qaz ETS
- 802.1Qbb PFC
- 802.3ad Link Aggregation with LACP
- 802.3ba
- 802.3x Flow Control
- 1000BASE-KX
- 802.3ae 10 Gigabit Ethernet



Technical Specifications

SNMP MIBs

Due to licensing restrictions, Cumulus Linux does not install all MIBs. For the MIBs that Cumulus Linux does not install, you must add the “non-free” archive to `/etc/apt/sources.list`. To see which MIBs are on your switch, run `ls /usr/share/snmp/mibs/`.

BGP4-MIB OSPFv2-MIB OSPFv3-MIB RIPv2-MIB	You can enable FRR SNMP support to provide support for OSPF-MIB (RFC-1850), OSPFv3-MIB (RFC-5643), and BGP4-MIB (RFC-1657).
CUMULUS-BGPVRF-MIB	Cumulus Linux also includes its own BGP unnumbered MIB for BGP unnumbered peers, defined in <code>/usr/share/snmp/mibs/Cumulus-BGPUN-MIB.txt</code> , which has the OID <code>1.3.6.1.4.1.40310.7</code> .
CUMULUS-COUNTERS-MIB	Discard counters: Cumulus Linux also includes its own counters MIB, defined in <code>/usr/share/snmp/mibs/Cumulus-Counters-MIB.txt</code> . It has the OID <code>1.3.6.1.4.1.40310.2</code> .
CUMULUS-RESOURCE-QUERY-MIB	Cumulus Linux includes its own resource utilization MIB, which is similar to using <code>cl-resource-query</code> . This MIB monitors layer 3 entries by host, route, nexthops, ECMP groups, and layer 2 MAC/BDPU entries. <code>/usr/share/snmp/mibs/Cumulus-Resource-Query-MIB.txt</code> defines this MIB, which has the OID <code>1.3.6.1.4.1.40310.1</code> .
CUMULUS-SNMP-MIB	SNMP counters. For information on exposing CPU and memory information with SNMP, see this knowledge base article.
DISMAN-EVENT-MIB	Trap monitoring
ENTITY-MIB	Cumulus Linux supports the temperature sensors, fan sensors, power sensors, and ports from RFC 4133. Notes: The ENTITY-MIB does not show the chassis information in Cumulus Linux.
ENTITY-SENSOR-MIB	Physical sensor information (temperature, fan, and power supply) from RFC 3433.
HOST-RESOURCES-MIB	Users, storage, interfaces, process info, run parameters.
BRIDGE-MIB Q-BRIDGE-MIB	The <code>dot1dBasePortEntry</code> and <code>dot1dBasePortIfIndex</code> tables in the BRIDGE-MIB and <code>dot1qBase</code> , <code>dot1qFdbEntry</code> , <code>dot1qTpFdbEntry</code> , <code>dot1qTpFdbStatus</code> , and <code>dot1qVlanStaticName</code> tables in the Q-BRIDGE-MIB tables. You must uncomment the <code>bridge_pp.py</code> <code>pass_persist</code> script in <code>/etc/snmp/snmpd.conf</code> .
IEEE8023-LAG-MIB	Implementation of the IEEE 8023-LAG-MIB includes the <code>dot3adAggTable</code> and <code>dot3adAggPortListTable</code> tables. To enable this, edit <code>/etc/snmp/snmpd.conf</code> and uncomment or add the following lines: view systemonly included 1.2.840.10006.300.43 pass_persist 1.2.840.10006.300.43 <code>/usr/share/snmp/ieee8023_lag_pp.py</code>
IF-MIB	Interface description, type, MTU, speed, MAC, admin, operation status, counters. Notes: Cumulus Linux disables the IF-MIB cache by default. The non-caching code path in the IF-MIB treats 64-bit counters like 32-bit counters (a 64-bit counter rolls over after the value increments to a value that extends beyond 32 bits). To enable the counter to reflect traffic statistics using 64-bit counters, remove the <code>-y</code> option from the <code>SNMPDOPTS</code> line in the <code>/etc/default/snmpd</code> file. The example below first shows the original line, commented out, then the modified line without the <code>-y</code> option: cumulus@switch:~\$ cat /etc/default/snmpd # SNMPDOPTS='-y -LS 0-4 d -Lf /dev/null -u snmp -g snmp -l -smux -p /run/snmpd.pid' <code>SNMPDOPTS='-LS 0-4 d -Lf /dev/null -u snmp -g snmp -l -smux -p /run/snmpd.pid'</code>
IP-FORWARD-MIB	IP routing table.
IP-MIB (includes ICMP)	IPv4, IPv4 addresses counters, netmasks.
IPv6-MIB	IPv6 counters.
LLDP-MIB	Layer 2 neighbor information from <code>lldpd</code> (you need to enable the SNMP subagent in LLDP). You need to start <code>lldpd</code> with the <code>-x</code> option to enable connectivity to <code>snmpd</code> (AgentX).

Technical Specifications

LM-SENSORS MIB	Fan speed, temperature sensor values, voltages. The ENTITY-SENSOR MIB replaces this MIB.
NET-SNMP-AGENT-MIB	Agent timers, user, group config.
NET-SNMP-VACM-MIB	Agent timers, user, group config.
NOTIFICATION-LOG-MIB	Local logging.
SNMP-FRAMEWORK-MIB	Users, access.
SNMP-MPD-MIB	Users, access.
SNMP-TARGET-MIB	SNMP-TARGET-MIB.
SNMP-USER-BASED-SM-MIBS	Users, access.
SNMP-VIEW-BASED-ACM-MIB	Users, access.
TCP-MIB	TCP-related information.
UCD-SNMP-MIB	System memory, load, CPU, disk IO.
UDP-MIB	UDP-related information.



Summary of Changes

Date	Version History	Action	Description of Change
03-Jun-2024	Version 3	Changed	Overview, Configuration Information and Technical Specifications sections were updated
01-Apr-2024	Version 2	Changed	Overview and Configuration Information sections were updated
04-Mar-2024	Version 1	New	New QuickSpecs




Copyright

Make the right purchase decision.
Contact our presales specialists.

 Chat now (sales)

 Call now



Get updates

© Copyright 2024 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

a50007023enw - 17133 - Worldwide - V3 - 03-June-2024