NVIDIA QUANTUM-2 CS9500 MODULAR SWITCH SERIES

Deliver Unparalleled Data Throughput
and Density with 400G InfiniBand

Boost the Availability and Data Capacity of High-Performance Computing (HPC) and Hyperscale Data Centers for Research, Innovation, and Product Development

As artificial intelligence (AI) and increasingly complex applications demand faster, smarter, and more scalable networks, NVIDIA 400 gigabits per second (400Gb/s) InfiniBand provides the fastest networking solution available, offered on the world's only fully offloadable in-network computing platform.

NVIDIA 400Gb/s InfiniBand's massive throughput, smart acceleration engines, flexibility, and robust architecture let HPC, AI, and hyperscale cloud infrastructures achieve unmatched performance, with less cost and complexity. Providing up to 2,048 400Gb/s InfiniBand ports, the high-density NVIDIA CS9500 modular switches enable an extremely high-performance, low-latency fabric solution for exascale computing and hyperscale cloud data centers.

The NVIDIA Quantum-2 CS9500 modular switch offers 2.5X the port density of the preceding switch generation while boosting AI acceleration by 32X. In addition, it surges the previous generation modular switch system's aggregated bidirectional throughput by 5X, to 1.64 petabits per second, enabling users to run larger workloads with fewer constraints.

The Era of Data-Driven Computing

Complex workloads demand ultra-fast processing of high-resolution simulations, extreme-size datasets, and complex, highly parallelized algorithms that require real-time information exchanges. As the highest-performing fabric solution in a 29U form factor, the CS9500 delivers 1600 terabits per second (Tb/s) of full bidirectional bandwidth with ultra-low port latency. The CS9500 modular switches create the highest scalability for large data aggregation through the network, with the highest application performance of complex computations while data moves through the data center network.

World's Only Fully In-Network Acceleration Platform

Offloading operations is crucial for AI workloads. In keeping with previous generations of InfiniBand switches, the seventh generation includes CPU offloads and remote direct-memory access (RMDA) that enable ultra-low latency and double the data throughput, while lowering CPU utilization. It also comes with new NVIDIA In-Network Computing engines to provide additional acceleration.

SYSTEM SPECIFICATIONS

Switch system options	**2048 400Gb/s InfiniBand ports or 4096 200Gb/s InfiniBand ports delivering 1600 Tb/s total throughput**
	1024 400Gb/s InfiniBand ports or 2048 200Gb/s InfiniBand ports delivering 800 Tb/s total throughput
	512 400Gb/s InfiniBand ports or 1048 200Gb/s InfiniBand ports delivering 400 Tb/s total throughput
Cooling	**Liquid-cooled**
	> Liquid-to-air AHX
	> Liquid-to-liquid CDU

NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™ technology improves collective operations for HPC protocols such as Message Passing Interface (MPI) and Symmetric Hierarchical Memory (SHMEM), resulting in an order-of-magnitude improvement in application performance. The third-generation NVIDIA SHARP allows deep learning training operations to be offloaded and accelerated by the InfiniBand network, resulting in 32X higher AI acceleration power. When combined with the NVIDIA Magnum IO™ software stack, it provides out-of-the-box accelerated scientific computing.

Maximizing Computing Efficiencies

The CS9500 ensures the maximum effective fabric performance under all types of traffic conditions, eliminating congestion hot spots. Features like self-healing networking provide resiliency and fast rerouting of traffic in case of port failure. Adaptive routing optimizes load balancing to eliminate in-network congestion, and quality of service (QoS) enables traffic prioritization.

World-Class, Flexible Design

The modular CS9500 chassis switch is designed for flexibility, speed, serviceability, energy savings, and high availability. Unlike traditional InfiniBand chassis switches, the CS9500 modular switch contains 2U spine switches. Benefits include:

> Three versions of switches to choose from based on number of ports: 2,048, 1,024, or 512

> Flexible pay-as-you-grow design enables data center infrastructures to start small and scale according to need

> Single point of management for each switch

With its excellent price-performance ratio for medium- to extremely large-size clusters, the modular switch's leaf blades, spine blades, management modules, power supplies, and fan units are all hot-swappable to help eliminate downtime.

As an eco-friendly solution, the CS9500 is cooled solely by a liquid closed loop to enable a low noise level and reduce IT opex. It arrives with a liquid cooling data unit (CDU) or air heat exchange unit (AHX) to best fit any data center. In addition, NVIDIA Quantum-2-based switches provide maximum flexibility, supporting a wide variety of network topologies, including Fat Tree, DragonFly+, multi-dimensional Torus, and more. NVIDIA modular switches are also backwards compatible to previous generations and include extensive software ecosystem support.

Enhanced Management

Running the NVIDIA MLNX-OS® software package, the subnet manager delivers full chassis management through command-line interface (CLI), web-based user (WebUI), Simple Network Management Protocol (SNMP), or JavaScript Object Notation (JSON) interfaces.

The Quantum-2-based switches can also utilize the advanced NVIDIA Unified Fabric Manager (UFM®) feature sets to empower data center operators to efficiently provision, monitor, manage, preventatively troubleshoot, and maintain the modern data center fabric, to realize higher utilization and reduce overall opex.

Ordering Information

ORDERABLE PART NUMBER (OPN)	DESCRIPTION
MCS9500	NVIDIA 1,600Tb/s, 2,048-port 400Gb/s InfiniBand chassis and internal connectivity
MCS9510	NVIDIA 800Tb/s, 1,024-port 400Gb/s InfiniBand chassis and internal connectivity
MCS9520	NVIDIA 400Tb/s, 512-port 400Gb/s InfiniBand chassis and internal connectivity
MCS9505-AHX	NVIDIA MCS95XX modular systems liquid-to-air heat exchanger
MCS9500-KIT-AHX	NVIDIA MCS9500, 2 AHXs equipment kit
MCS9500-REDKIT-AHX	NVIDIA MCS9500, 3 AHXs redundancy equipment kit
MCS9510-KIT-AHX	NVIDIA MCS9510, 1 AHX equipment kit
MCS9510-REDKIT-AHX	NVIDIA MCS9510, 2 AHXs redundancy equipment kit
MQM9510-N	Quantum-2 400Gb/s InfiniBand, 2U leaf blade, 128 NDR ports, 64 OSFP ports
MQM9520-N	NVIDIA Quantum-2 400Gb/s InfiniBand, 2U spine blade, 128 NDR ports
MTDF-LIQ-D	NVIDIA MCS95XX chassis 19 liters PG25 Coolant
MMB9500	NVIDIA MCS95XX management switch
MTDF-PDU-A	NVIDIA MCS95XX modular system PDU
GPS-CS9500-OST	CS9500 (2,048-port) 400Gb/s chassis installation, including onsite rack and liquid cooling system deployment. Price per chassis
GPS-CS9510-20-OST	CS9510 (1,024-port) 400Gb/s chassis installation, including onsite rack and liquid cooling system deployment. Price per chassis

Learn more

Learn more about NVIDIA InfiniBand switches at **nvidia.com/en-us/networking/infiniband-switching/**

Learn more about NVIDIA UFM software at **nvidia.com/en-us/networking/management-software/**

NVIDIA.