



NVIDIA QUANTUM-2 QM9700 SERIES

Scaling Out Data Centers with
NDR 400G InfiniBand Smart Switches

Accelerate Research, Innovation, and Product Development with Greater Data Capacity and In-Network Computing

As high-performance computing (HPC) and artificial intelligence (AI) applications become more complex, the demand for the most advanced high-speed networking is critical for extreme-scale systems. NVIDIA Quantum-2 is the industry-leading switch platform in power and density, with NDR 400 gigabit per second (Gb/s) InfiniBand throughput that provides AI developers and scientific researchers with the highest networking performance available to take on the world's most challenging problems.

Advanced Computing Needs Advanced Networking

The NVIDIA Quantum-2-based QM9700 and QM9790 switch systems deliver an unprecedented 64 ports of NDR 400Gb/s InfiniBand per port in a 1U standard chassis design. A single switch carries an aggregated bidirectional throughput of 51.2 terabits per second (Tb/s), with a landmark of more than 66.5 billion packets per second (BPPS) capacity. Supporting the latest NDR technology, NVIDIA Quantum-2 brings a high-speed, extremely low-latency and scalable solution that incorporates state-of-the-art technologies such as Remote Direct Memory Access (RDMA), adaptive routing, and NVIDIA Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™.

Unlike any other networking solution, NVIDIA InfiniBand provides self-healing network capabilities, as well as quality of service (QoS), enhanced virtual lane (VL) mapping, and congestion control to provide the highest overall application throughput. As an ideal rack-mounted InfiniBand solution, the QM9700 and QM9790 NDR InfiniBand fixed-configuration switches allow maximum flexibility, as they enable a variety of topologies, including Fat Tree, SlimFly, DragonFly+, multi-dimensional Torus, and more. They're also backwards compatible to previous generations and include expansive software ecosystem support.

The Era of Data-Driven Computing

Today's complex research demands ultra-fast processing of high-resolution simulations, extreme-size datasets, and complex, highly parallelized algorithms that need to exchange information in real time. The QM9700 NDR InfiniBand switches extend NVIDIA In-Network Computing technologies and introduce the third generation of NVIDIA SHARP technology, SHARPv3. Creating virtually unlimited scalability for large data aggregation through the

exertis | ENTERPRISE

a **DCC** business

Speak to an expert
define@exertisenterprise.com



network, SHARPV3 enables support for up to 64 parallel flows—32X higher AI acceleration power compared to the previous generation. SHARPV3 dramatically boosts application performance of complex computations while data moves through the data center network, participating in the application’s runtime and reducing the amount of data needed to traverse the network.

Streamlining Network Design and Topologies

By implementing NVIDIA port-split technology, the QM9700 and QM9790 switches provide a double-density radix for 200Gb/s (NDR200) data speeds, reducing the cost of network design and network topologies. Supporting up to 128 ports of 200Gb/s, NVIDIA delivers the densest top-of-rack (TOR) switch available on the market. The QM9700 family of switches enables small to medium-sized deployments to scale with a two-level Fat Tree topology while reducing power, latency, and space requirements.

Enhanced Management

The internally managed QM9700 switch features an on-board subnet manager that enables simple, out-of-the-box bringup for up to 2,000 nodes. 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 externally managed QM9790 switch can 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.

ORDERABLE PART NUMBER (OPN)	DESCRIPTION
MQM9700-NS2F	64-ports NDR, 32 OSFP ports, managed, power-to-connector (P2C) airflow (forward)
MQM9700-NS2R	64-ports NDR, 32 OSFP ports, managed, connector-to-power (C2P) airflow (reverse)
MQM9700-NS2FS	64-ports NDR, 32 OSFP ports, managed, P2C airflow (forward), secure boot
MQM9700-NS2RS	64-ports NDR, 32 OSFP ports, managed, C2P airflow (reverse), secure boot
MQM9790-NS2F	64-ports NDR, 32 OSFP ports, unmanaged, P2C airflow (forward)
MQM9790-NS2R	64-ports NDR, 32 OSFP ports, unmanaged, C2P airflow (reverse)
MTEF-PSF-AC-J	NVIDIA power-supply unit, 2,000W AC, P2C airflow; for QM97xx switches, power cord included
MTEF-PSR-AC-J	NVIDIA power-supply unit, 2,000W AC, C2P airflow; for QM97xx switches, power cord included
MTEF-FANF-L	400G 1U systems fan module with P2C air flow
MTEF-FANR-L	400G 1U systems fan module with C2P air flow

SYSTEM SPECIFICATIONS

Performance	400Gb/s per port
Switch radix	64 NDR non-blocking ports with aggregate data throughput up to 51.2Tb/s
Connectors and cabling	32 octal small form-factor pluggable (OSFP) connectors; passive or active copper or active fiber cable; optical module
System power usage	Maximum power with passive cables: 1,084W Maximum power with optical cables (with each port consuming 18.5W): 1,720W
Power supply	1+1 redundant and hot-swappable power Input range: 200-240Vac 80 Gold+ and ENERGY STAR certified
Cooling	Front-to-rear or rear-to-front Cooling option: hot-swappable fan unit
Management ports	1x USB 3.0 x1 1x USB for I2C channel 1x RJ45 1x RJ45 (UART)
CPU	x86 Coffee Lake i3
System memory	Single 8GB, 2,666 mega transfers per second (MT/s), DDR4 SO-DIMM
Storage	M.2 SSD SATA 16GB 2242 FF
Software	MLNX-OS
System weight	14.5kg
System dimensions	Height: 1.7in (43.6mm) Width: 17.0in (438mm) Depth: 26.0in (660.4mm)
Rack mount	1U rack mount
Operating conditions	Temperature: > operating 0°C to 40°C > non-operating -40°C to 70°C Humidity: > operating 10% to 85% non-condensing > non-operating 10% to 90% non-condensing Altitude: Up to 3050m
EMC (emissions)	CE, FCC, VCCI, ICES, and RCM
Product safety compliant/certified	RoHS, CB, cTUVus, CE, and CU
Warranty	1 year