



NVIDIA Quantum-2 QM9700

Also known as: MQM9700-NS2F · MQM9790-NS2F · QM9790 · Quantum-2 NDR switch

The Quantum-2 QM9700 is a 1U, fixed-configuration NDR InfiniBand switch built around the Quantum-2 ASIC: 64 ports of 400 Gb/s NDR across 32 twin-port OSFP cages, 51.2 Tb/s of non-blocking aggregate throughput, and a 64-way radix that makes it the standard building block for H100 and H200 GPU fabrics. It ships as a managed switch (onboard x86 subnet manager, up to 2,000 nodes) or as an externally managed switch that hands subnet management to NVIDIA UFM. It supports SHARPV3 in-network reduction, adaptive routing and SHIELD self-healing link recovery. The same chassis serves as either leaf or spine in a two-tier fat-tree design, which is how most H100/H200 clusters still spec their fabric in 2026.

64 × 400 Gb/s NDR Port configuration · 32 physical twin-port OSFP cages, splittable to 128 × 200 Gb/s	51.2 Tb/s Aggregate throughput · bidirectional, non-blocking	66.5 billion pps Packet rate
130 ns Port-to-port latency	747 W typical Power · passive cables; up to 1,720 W with active cables	1U Form factor · 19 in rack, air-cooled

SWITCHING		01/05
ASIC	NVIDIA Quantum-2	
Switching capacity	51.2 Tb/s · bidirectional aggregate; 25.6 Tb/s one direction	
Packet processing	66.5 billion pps	
Radix	64 ports · at 400 Gb/s NDR	
In-network computing	SHARPV3 · up to 64 parallel flows, 32x AI acceleration vs prior generation	
Topologies supported	Fat Tree, DragonFly+, SlimFly, multi-dimensional Torus	

PORTS & CABLING		02/05
Connectors	32 × OSFP · twin-port cages, each carrying 2 × 400 Gb/s	
Port split option	400G → 2 × 200G · OSFP to 2 × QSFP56 DAC/AOC splitter	
Legacy interop	OSFP → 2 × QSFP56 · connects QM9700 to QM8700 HDR spine switches	
Cable types	DAC, AOC, active/passive optics	
Speed compatibility	40 / 56 / 100 / 200 / 400 Gb/s per port	
POWER & THERMAL		03/05
Power, QM9700 (managed)	747 W typical / 1,720 W max · typical with passive cables, max with active cables	
Power, QM9790 (unmanaged)	640 W typical / 1,610 W max	
Input	200-240 VAC, 10 A, 50/60 Hz	
Power supplies	1+1 redundant, hot-swappable	
Fans	6+1 hot-swappable	
Operating temperature	0-35°C forward airflow / 0-40°C reverse airflow	
Acoustic	78.4 dBA · at room temperature	
PHYSICAL		04/05
Dimensions (H x W x D)	43.6 mm × 438 mm × 660 mm · 1.7 in × 17.2 in × 26 in	
Weight	13.6-14.8 kg · single vs dual PSU configuration	
Rack mount	19 in standard rack	
Humidity	10%-85% non-condensing	
Altitude	3,050 m max	
MANAGEMENT & FEATURES		05/05
Subnet manager (managed SKU)	onboard x86 · up to 2,000 nodes	
Subnet manager (unmanaged SKU)	NVIDIA UFM · externally managed, no practical node cap	
Software	MLNX-OS	
Interfaces	CLI, WebUI, SNMP (v1/v2c/v3), JSON	
Self-healing	SHIELD · link recovery ~5,000x faster than software-based failover	
GPU acceleration	GPU Direct RDMA, NCCL support	

PROCUREMENT

Typical lead time

Quote-driven; no fixed published lead time. Reported allocation windows for NVIDIA InfiniBand switches run from several weeks to a few months depending on model, region and distributor channel, and stretch further during AI build-out demand surges.

Pricing

\$33,000 to \$53,000, reported. FS.com listed the managed MQM9700-NS2F at \$38,400; NADDOD listed the same part bundled with 5-year service at \$53,448; other resellers (Connection, ServerSupply, Optdex) clustered in the low-to-mid \$30,000s to high \$30,000s. Treat any single figure as a snapshot, not a list price.

- Managed (MQM9700-NS2F) includes an onboard x86 subnet manager for up to 2,000 nodes; externally managed (MQM9790-NS2F) is the cheaper SKU when NVIDIA UFM is already running the fabric.
- OSFP transceivers and DAC/AOC cabling are frequently the largest single line item in an NDR fabric build, not the switch chassis itself: one documented 1,024-GPU (128-node) H100 build saw transceivers alone exceed \$1.5M at NVIDIA list pricing before third-party module substitution cut that by roughly 55%.
- Splitter cables let one 400G OSFP port fan out to 2 × 200G, useful for HDR-generation interop or building 128-port fabrics from a 64-port chassis, but they consume switch radix budget that a straight 400G run does not.
- The 32 physical OSFP cages are twin-port: the 64-port spec sheet number is double the cage count, worth flagging early when planning patch panels and cable runs.

FIELD NOTES

- In rail-optimized H100/H200 topologies, a common reference design uses 32 leaf switches (one per rail group) with a full-mesh spine layer; a two-tier, non-blocking fat tree built from 64-port QM9700 leaves and spines is commonly sized in the range of ~1,000 nodes / ~8,000 GPUs before oversubscription or a third tier is needed. Treat this as an illustrative topology size, not a vendor-guaranteed capacity number.
- Pick MQM9700 (managed) if you want the onboard subnet manager and are not already running UFM; pick MQM9790 (externally managed) if UFM is already part of the fabric management stack, it is typically the lower-cost SKU per unit.
- Budget cable and optics cost separately from the switch line item, and budget it early: OSFP DAC/AOC and transceivers routinely dominate NDR fabric spend.
- QM9700 remains the safe default for H100/H200 NDR clusters going into new builds in 2026: it is a mature, broadly stocked platform with a deep third-party optics ecosystem. Quantum-X800 is the newer step up for 800G, Blackwell-class endpoints, not a drop-in replacement.

QUESTIONS WE GET ON THIS PART

How many GPUs can a single QM9700 fabric support?

One switch has 64 ports of 400 Gb/s NDR. Fabric-wide capacity scales with the number of switches in a leaf/spine design, not a single unit: the onboard subnet manager on a managed QM9700 tops out at 2,000 nodes, and two-tier non-blocking fat trees built from 64-port leaves and spines are commonly sized in the low thousands of GPUs before a third tier is required.

Quantum-2 QM9700 vs Quantum-X800: which one for a new build?

QM9700 runs NDR at 400 Gb/s per port and pairs naturally with H100/H200 and ConnectX-7; it is the mature, widely available choice for 2026 procurement. Quantum-X800 doubles port speed to XDR 800 Gb/s, targets Blackwell-class GPUs with ConnectX-8/9, and is the platform to specify for new 800G-endpoint builds.

Managed or unmanaged QM9700?

Managed (MQM9700-NS2F) ships with an onboard x86 subnet manager good for up to 2,000 nodes and is the simpler out-of-box option. Unmanaged/externally-managed (MQM9790-NS2F) relies on NVIDIA UFM running elsewhere and is typically the lower-priced unit, the right call once UFM is already deployed.

Is the port count 64 or 32?

Both numbers are accurate for different things. The chassis has 32 physical OSFP cages, but each cage is a twin-port connector carrying two independent 400 Gb/s lanes, giving 64 usable 400G ports, or up to 128 ports if split down to 200G.

REFERENCES

- [1] <https://networking-docs.nvidia.com/qm97x0hw/introduction>
- [2] <https://networking-docs.nvidia.com/qm97x0hw/specifications>
- [3] <https://www.naddod.com/products/nvidia-networking/102404>
- [4] <https://www.fs.com/products/194710.html>
- [5] <https://www.advancedhpc.com/pages/nvidia-quantum-2-qm9700-series-switches>
- [6] <https://www.fibermall.com/blog/nvidia-infiniband-switches.htm>
- [7] <https://ascentoptics.com/blog/osfp-nvidia-infiniband-ndr/>
- [8] <https://www.fibermall.com/blog/osfp-ai-networking-guide.htm>