



QM8700

Mellanox Quantum™ HDR Edge Switch

40-port Non-blocking Managed HDR 200Gb/s InfiniBand Smart Switch

Mellanox provides the world's smartest switch, enabling in-network computing through the Co-Design Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)[™] technology. QM8700 has the highest fabric performance available in the market with up to 16Tb/s of non-blocking bandwidth with sub-90ns port-to-port latency.

SCALING-OUT DATA CENTERS WITH HDR 200G INFINIBAND

Faster servers, combined with high-performance storage and applications that use increasingly complex computations are causing data bandwidth requirements to spiral upward. As servers are deployed with next generation processors, High-Performance Computing (HPC) environments and Enterprise Data Centers (EDC) will need every last bit of bandwidth delivered with Mellanox's next generation of HDR InfiniBand, high-speed, smart switches.

WORLD'S SMARTEST SWITCH

Built with Mellanox's Quantum InfiniBand switch device, the QM8700 provides up to forty 200Gb/s ports, with full bi-directional bandwidth per port. These stand-alone switches are an ideal choice for top-of-rack leaf connectivity or for building small to extremely large sized clusters.

QM8700 is the world's smartest network switch, designed to enable in-network computing through the Co-Design Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)[™] technology. The Co-Design architecture enables the usage of all active data center devices to accelerate the communications frameworks using embedded hardware, resulting in an order of magnitude application performance improvements.

QM8700 enables efficient computing with features such as static routing, adaptive routing, congestion control and enhanced VL mapping to enable modern topologies (SlimFly, Dragonfly+, 6DT). These ensure the maximum effective fabric bandwidth by eliminating congestion hot spots.

COLLECTIVE COMMUNICATION ACCELERATION

Collective communication describes communication patterns in which all members of a group of communication endpoints participate. Collective communications are commonly used in HPC protocols such as MPI and SHMEM.

The Quantum switch improves the performance of selected collective operations by processing the data as it traverses the network, eliminating the need to send data multiple times between endpoints.



HIGHLIGHTS

BENEFITS

- Industry-leading switch platform in performance, power, and density
- Designed for energy and cost savings
- Collective communication acceleration
- Maximizes performance by removing fabric congestions
- Backward compatible to previous speeds
- Quick and easy setup and management

KEY FEATURES

- Performance
 - 40 x HDR 200Gb/s ports in a 1U switch
 - 80 x HDR100 100Gb/s ports (using splitter cables)
 - 16Tb/s aggregate switch throughput
 - Sub-90ns switch latency
- Optimized design
 - 1+1 redundant & hot-swappable power
 - N+1 redundant & hot-swappable fans
 - 80 gold+ and energy star certified power supplies
 - x86 ComEx Broadwell CPU
- Advanced design
 - Adaptive routing
 - Congestion control
 - Collective offloads (Mellanox SHARP technology)
 - VL mapping (VL2VL)

It also supports the aggregation of large data vectors at wire speed to enable MPI large vector reduction operations, which are crucial for machine learning applications.

HDR100

QM8700 together with the Mellanox ConnectX[®]-6 adapter card support HDR100. By utilizing two pairs of two lanes per port, the QM8700 can support up to 80 ports of 100G to create the densest TOR switch available in the market. This is a perfect solution for double dense racks with more than 40 servers per rack and also helps small-medium deployments with the need to scale to 3-level fat-tree, to lower power, latency and space.

MANAGEMENT

The QM8700's x86 ComEx Broadwell CPU comes with an on-board subnet manager, enabling simple, out-of-the-box bring-up for up to 2K nodes in the fabric. Running the MLNX-OS[®] software package, it delivers full chassis management through CLI, WebUI, SNMP or JSON interfaces.

QM8700 also incorporates Mellanox Unified Fabric Manager software for managing scale-out, InfiniBand, computing environments to enable efficient provisioning, health indications and monitoring of the cluster. UFM[®] ensures that the fabric is up and running at maximum performance at all times.

FEATURES

Mellanox QM8700

- 19" rack mountable 1U chassis
- 40 QSFP56 non-blocking ports with aggregate data throughput up to 16 Tb/s (HDR)

Switch Specifications

- Compliant with IBTA 1.21 and 1.3
- 9 virtual lanes: 8 data + 1 management
- 256 to 4Kbyte MTU
- Adaptive Routing
- Congestion control
- Port Mirroring

- VL2VL mapping
- 4X48K entry linear forwarding database

Management Ports

- 100/1000 RJ45 Ethernet port
- RS232 console port
- USB port
- DHCP
- Industry standard CLI
- Management over IPv6
- Management IP
- SNMP v1,v2,v3
- WebUI

Fabric Management

- On-board Subnet Manager (SM) supporting fabrics of up to 2K nodes
- Unified Fabric Manager (UFM) agent

Connectors and Cabling

- QSFP56 connectors
- Passive copper or active fiber cables
- Optical modules

Indicators

- Per port status LED Link, Activity
- System LEDs: System, fans, power supplies
- Unit ID LED

Power Supply

- Dual redundant slots
- Hot plug operation
- Input range: 100-127VAC, 200-240VAC
- Frequency: 50-60Hz, single phase AC, 4.5A, 2.9A

Cooling

- Front-to-rear or rear-to-front cooling option
- Hot-swappable fan unit

Power Consumption

- Contact Mellanox Sales

COMPLIANCE

Safety

- CB
- cTUVus
- CE
- CU

EMC (Emissions)

- CE
- FCC
- VCCI
- ICES
- RCM

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 3200m

Acoustic

- ISO 7779
- ETS 300 753

Others

- RoHS compliant
- Rack-mountable, 1U
- 1-year warranty

Table 1 - Part Numbers and Descriptions

OPN	Description
MQM8700-HS2F	Mellanox Quantum™ HDR 200Gb/s InfiniBand switch, 40 QSFP56 ports, non-blocking switching with capacity of 16Tb/s, 2 AC PSUs, standard depth, x86 CPU, P2C airflow, Rail Kit
MTEF-PSF-AC-C	200G 1U systems 850W AC Power Supply w/ P2C air flow
MTEF-PSR-AC-C	200G 1U systems 850W AC Power Supply w/ C2P air flow
MTEF-FANF-C	200G 1U systems fan module w/ P2C air flow
MTEF-FANR-C	200G 1U systems fan module w/ C2P air flow

*P2C is connector side outlet, C2P is connector side inlet.



350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085
 Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com