



Mellanox
Quantum

CS8500

Mellanox Quantum™ HDR Modular Switch

800-port HDR 200Gb/s InfiniBand Smart Modular Switch

Mellanox provides the world's smartest switch, enabling in-network computing through the Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™ technology. CS8500 has the highest performing fabric solution in a 29U form factor by delivering 320Tb/s of full bi-directional bandwidth with ultra-low port latency.

SCALING-OUT DATA CENTERS WITH HDR 200G INFINIBAND

Faster servers based on PCIe 3.0 and above, 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 devices, the CS8500 provides up to 800 200Gb/s full bi-directional bandwidth per port. The CS8500 modular chassis switch provides an excellent price performance ratio for medium to extremely large size clusters, along with the reliability and manageability expected from a modular-class switch.

CS8500 is the world's smartest network switch designed to enable in-network computing through Scalable Hierarchical Aggregation and Reduction Protocol (SHARP) technology. Mellanox SHARP architecture enables the usage of all active data center devices to accelerate the communications frameworks, resulting in an order of magnitude improvement in applications performance.

CS8500 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 features ensure the maximum effective fabric bandwidth by eliminating congestion hot spots.

WORLD-CLASS DESIGN

CS8500 is an elegant modular switch designed for performance, serviceability, energy savings and high availability. The leaf blades, spine blades and management modules, as well as the power supplies and fan units, are all hot-swappable to help eliminate down time. CS8500 is an eco-friendly solution cooled solely by a liquid closed loop to enable a low noise level and reduce IT OPEX. CS8500 arrives with liquid and air coolant data units (CDUs) to best fit all data centers.



HIGHLIGHTS

BENEFITS

- Industry-leading switch platform in performance, power and density
- Collective communication acceleration
- Maximizes performance by removing fabric congestions
- Backward compatible to previous InfiniBand speeds

FEATURES

- Performance
 - 800 HDR 200Gb/s ports;
 - 1600 HDR100 100Gb/s ports
 - 320Tb/s aggregate switch throughput
 - Ultra low switch latency
- Optimized design
 - N+N redundant & hot-swappable power
 - 80 gold+ and energy star certified power supplies
 - x86 ComEx Broadwell CPU
 - Liquid-cooled system

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 Mellanox 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. 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.

CS8500 together with the Mellanox ConnectX®-6 adapter card support HDR100. By utilizing two pairs of two lanes per port, the CS8500 can support up to 1600 ports of 100G to create the densest modular switch available in the market. This is a perfect solution for medium-large deployments that require lower power, latency and space.

FEATURES

Mellanox CS8500

- 29U modular chassis
- 40 QSFP56 HDR 200Gb/s InfiniBand ports per leaf blade

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

Management Ports

- 100/1000Mb/s RJ45 Ethernet port

- RS232 console port
- USB port
- DHCP
- Industry standard CLI
- Management over IPv6
- Management IP
- SNMP v1,v2,v3
- Web UI

Fabric Management

- On-board Subnet Manager 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

Physical Characteristics

- Dimensions: TBD
- Weight: TBD

Power Supply

- Hot swappable with N+N redundancy

Cooling

- Ordering options:
 - Liquid to Air CDU
 - Liquid to Liquid CDU

Power Consumption

- TBD

COMPLIANCE

Safety

- CB
- cTUVus
- CE
- CU

EMC (Emissions)

- CE
- FCC
- VCCI
- ICES
- RCM

Operating Conditions

- Operating 0°C to 40°C
- Non-Operating -40°C to 70°C
- Humidity: Operating 10% to 85%, non-condensing
- Altitude: Operating -60 to 3200m

Acoustic

- ISO 7779
- ETS 300 753

Others

- RoHS-6 compliant
- 1-year warranty

Table 1 - CS8500 Series Part Numbers and Descriptions

OPN	Description
MCS8500	320Tb/s, 800-port HDR InfiniBand chassis, includes 16 power supplies (N+N), RoHS6
MSB8510	Mellanox Quantum™ HDR InfiniBand leaf blade, 40 QSFP56 ports, RoHS6
MSB8520	Mellanox Quantum™ HDR InfiniBand spine blade, 40 QSFP56 ports, RoHS6
MMB8500	MCS85XX modular systems X86 dual-core chassis management module, RoHS R6
MTDF-PS-B	MCS85xx modular systems 3KW AC power supply w/P2C airflow
MTDF-KIT-C	MCS85xx modular systems rail kit (not including hoses)
MTDF-LIQ-A	MCS85xx modular systems liquid-to-liquid heat exchanger
MTDF-LIQ-B	MCS85xx modular systems liquid-to-air heat exchanger