



SB7800

InfiniBand EDR 100Gb/s Switch System

36-port Non-blocking Managed EDR 100Gb/s InfiniBand Smart Switch

Mellanox provides the world's first smart switch, enabling in-network computing through the Co-Design Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)[™] technology. SB7800 has the highest fabric performance available in the market with up to 7.2Tb/s of non-blocking bandwidth with 90ns port-to-port latency.

Scaling-Out Data Centers with EDR 100G InfiniBand

Faster servers based on PCIe 3.0, 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 Extended Data Rate (EDR) InfiniBand high-speed smart switches.

World's First Smart Switch

Built with Mellanox's latest Switch-IB[®] 2 InfiniBand switch device, EDR uses efficient 64/66 encoding while increasing the per lane signaling rate to 25Gb/s. SB7800 provides up to thirty-six 100Gb/s 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.

SB7800 is the world's first smart 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, resulting in order of magnitude applications performance improvements.

SB7800 enables efficient computing with features such as static routing, adaptive routing, congestion control and enhanced VL mapping to enable modern topologies (SlimFly, Dragonfly+, Torus). These ensure the maximum effective fabric bandwidth by eliminating congestion hot spots.

The SB7800 switch has best-in-class design to support low power consumption. ATIS weighted power consumption is 136W for a fully-populated switch. Power is further reduced if not all ports are used or if partially utilized.

Collective Communication Acceleration

Collective communication is a term used to describe communication patterns in which all members of a group of communication endpoints participate. Collective communications are commonly used in HPC communication protocols such as MPI and SHMEM (Open-SHMEM).

Switch-IB[®] 2[†]

HIGHLIGHTS

BENEFITS

- Industry-leading switch platform in performance, power, and density
- Collective communication acceleration
- Quick and easy setup and management
- Maximizes performance by removing fabric congestions
- Backward compatible to FDR technology

FEATURES

- Performance
 - 36 x EDR 100Gb/s ports in a 1U switch
 - 7.2Tb/s aggregate switch throughput
 - Up to 7.02 billion messages-per-second
 - 90ns switch latency
 - 136W typical power consumption
- Optimized design
 - 1+1 redundant & hot-swappable power
 - N+1 redundant & hot-swappable fans
 - 80 gold+ and energy star certified power supplies
 - Dual-core x86 CPU
- Advanced design
 - Adaptive routing
 - Congestion control
 - VL mapping (VL2VL)
 - Collective offloads (Mellanox SHARP technology)

Collective operations have implications on overall application performance and scale. SB7800 introduces the Co-design SHARP technology, which enables the switch to manage collective communications using embedded hardware. Switch-IB 2 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 end-points. This decreases the amount of data traversing the network and additional benefit of freeing up CPU resources for computation rather than using them to process communication.

The network portion of the reduction operation on a fully-populated three-level fat-tree can be completed in less than three microseconds.

Management

SB7800, dual-core x86 CPU, comes with an onboard subnet manager, enabling simple, out-of-the-box fabric bring-up for up to 2k nodes. SB7800 switch runs the same MLNX-OS[®] software package as Mellanox FDR products to deliver complete chassis management, to manage the firmware, power supplies, fans and ports.

FEATURES

Mellanox SB7800

- 19" rack mountable 1U chassis
- 36 QSFP28 non-blocking ports with aggregate data throughput up to 7.2Tb/s (EDR)

Switch Specifications

- Compliant with IBTA 1.21 and 1.3
- 9 virtual lanes:
8 data + 1 management
- 256 to 4Kbyte MTU
- Adaptive Routing
- VL2VL mapping
- 4x48K entry linear forwarding database

Management Ports

- 10/100/1000Mb/s Ethernet ports
- RS232 port over DB9
- USB port
- DHCP
- Familiar 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

- QSFP28 connectors
- Passive copper or active fiber cables
- Optical modules

Indicators

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

Physical Characteristics

- Dimensions: 1.7" (43.6 mm) H x 16.85" (428mm) W x 27" (685.8mm) D
- Weight: 11kg (24.2lb)

Power Supply

- Dual redundant slots
- Hot plug operation
- Input range:
100-127 VAC, 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

- Typical Power with Passive Cables (ATIS): 136W

COMPLIANCE

Safety

- CB
- cTUVus
- CE
- CU

EMC (Emissions)

- CE
- FCC

Safety

- VCCI
- ICES
- RCM

Operating Conditions

- Temperature:
– Operating 0°C to 45°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
MSB7800-ES2F	Switch-IB 2 based 36-port QSFP28 EDR 1U Managed InfiniBand switch system with a non-blocking switching capacity of 7.2Tb/s. 2PS, Standard depth, P2C airflow
MTEF-PSF-AC-A	460W AC Power Supply w/ P2C air flow
MTEF-PSR-AC-A	460W AC Power Supply w/ C2P air flow
MTEF-FANF-A	Fan module w/ P2C air flow
MTEF-FANR-A	Fan module w/ C2P air flow
LIC-Fabric-Inspector	Enhanced InfiniBand Diagnostics license

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