

SB7890 - Switch-IB™ 2 Externally Managed EDR Switch

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

Mellanox provides the world's first smart switch, enabling in-network computing through the Co-Design SHArP technology. SB7890 has the highest fabric performance available in the market with up to 7Tb/s of non-blocking bandwidth with 90ns port-to-port latency.

Scaling-Out Data Centers with Extended Data Rate (EDR) 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 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. The SB7890 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.

SB7890 is the world's first smart network switch, designed to enable in-network computing through the Co-Design Scalable Hierarchical Aggregation 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.

SB7890 enables efficient computing with features such as static routing, adaptive routing, and congestion control. These ensure the maximum effective fabric bandwidth by eliminating congestion hot spots.

The SB7890 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 end-points participate. Collective communications are commonly used in HPC communication protocols such as MPI and SHMEM (OpenSHMEM).

Collective operations have implications on overall application performance and scale. SB7890 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.

Building Efficient Clusters

The SB7890 is the industry's most cost-effective building block for deploying high performance clusters and data centers. Whether looking at price-to-performance or energy-to-performance, the SB7890 offers superior performance, low power and scale reducing capital and operating expenses providing the best return-on-investment.

Switch-IB™ 2



HIGHLIGHTS

BENEFITS

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

KEY FEATURES

Performance

- 36 X EDR 100Gb/s ports in a 1U switch
- 7Tb/s aggregate switch throughput
- 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

Advanced Design

- Adaptive Routing
- InfiniBand Router*
- Collective offloads (SHArP)

FEATURES

MELLANOX SB7890

- 19" rack mountable 1U chassis
- 36 QSFP28 non-blocking ports with aggregate data throughput up to 7 Tb/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
- Congestion control
- Port Mirroring
- IB Router*
- 4X48K entry linear forwarding data base

MANAGEMENT PORTS

- I2C (RJ45)
- System reset button

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: 1.72"H x 16.84"W x 27"D
- Weight: 11kg (24.2 Lb)

POWER SUPPLY

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

COOLING

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

POWER CONSUMPTION

- ATIS Weighted Power Consumption: 136W
- Max Power Consumption (class 4) : 281W

COMPLIANCE

SAFETY

- US/Canada: cTUVus
- EU: IEC60950
- International: CB

EMC (EMISSIONS)

- USA: FCC, Class A
- Canada: ICES, Class A
- EU: EN55022, Class A
- EU: EN55024, Class A
- EU: EN61000-3-2, Class A
- EU: EN61000-3-3, Class A
- Japan: VCCI, Class A

ENVIRONMENTAL

- EU: IEC 60068-2-64: Random Vibration
- EU: IEC 60068-2-29: Shocks, Type I / II
- EU: IEC 60068-2-32: Fall Test

OPERATING CONDITIONS

- Operating 0°C to 45°C,
- Non-Operating -40°C to 70°C
- Humidity: Operating 5% to 95%
- Altitude: Operating -60 to 3200m

ACOUSTIC

- ISO 7779
- ETS 300 753

OTHERS

- RoHS-6 compliant
- Rack-mountable, 1U
- 1-year warranty

*Available in future release

Ordering Part Number	Description
MSB7890-ES2F2	Switch-IB™ 2 based 36-port QSFP28 EDR 1U Managed InfiniBand switch system with a non-blocking switching capacity of 7Tb/s. 2PS, Standard depth, P2C airflow, RoHS-6
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

*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