

Switch-IB™ 2 EDR Switch Silicon

36-ports, supporting EDR 100Gb/s InfiniBand throughput per port, with a total of 7.2Tb/s bidirectional throughput and 7.02 billion messages per second

Mellanox continues its interconnect leadership by providing the world's first smart switch, enabling in-network computing through the Co-Design Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™ technology and the highest performing server and storage system interconnect solutions for Enterprise Data Centers, Cloud Computing, High-Performance Computing, and Embedded environments.

Mellanox Switch-IB 2 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.

Switch-IB 2 offers industry-leading integration of 144 SerDes, with rich speed flexibility ranging from 1 to 25 Gb/s per lane, making this Mellanox switch an obvious choice for OEMs that must address end-user requirements for faster and more robust applications. Network architects can utilize Switch-IB 2's reduced power and footprint, and a fully integrated PHY capable of connectivity across PCBs, backplanes, and passive and active copper/fiber cables, to deploy leading, fabric-flexible computing and storage systems with the lowest total cost of ownership.

EDR

EDR InfiniBand technology uses the efficient 64/66 encoding while increasing per lane signaling rate to 25Gb/s. By using 4-lane (4X) EDR ports and the Mellanox ConnectX® family of InfiniBand adapters, users can increase end-to-end bandwidth up to 78% over previous-generation FDR solutions.

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. Switch-IB 2 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 as aggregation nodes are reached. Implementing the collective communication algorithms in the network has the 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 that is built with 36-port switches can be completed in less than three microseconds.



Switch-IB™ 2 †

HIGHLIGHTS

- Industry-leading switch silicon in performance, power and density
- Industry-leading cut-through latency
- Low-cost solution
 - Single-chip implementation
- Fully integrated PHY
 - Backplane and cable support
 - 1, 2 and 4 lanes
- Up to 7.2Tb/s of switching capacity
- Up to 7.02 billion messages-per-second
- Up to 36 EDR 100Gb/s InfiniBand ports
- Hardware-based adaptive routing
- Congestion control
- Up to 8 switch partitions
- Mirroring

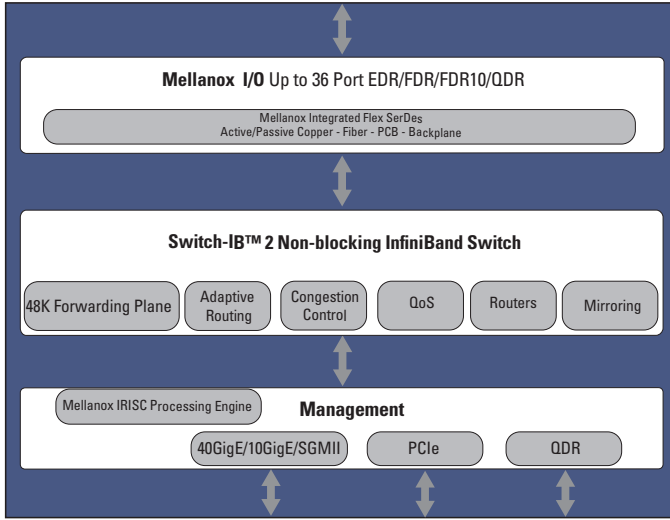


Figure 1. Mellanox Switch-IB 2 Architecture

INFINIBAND INTERCONNECT

Switch-IB 2 InfiniBand devices enable industry standard networking, clustering, storage, and management protocols to seamlessly operate over a single “one-wire” converged network. Combined with Mellanox’s ConnectX family of adapters, on-the-fly fabric repurposing can be enabled for Cloud, Web 2.0, EDC and Embedded environments providing “future proofing” of fabrics independent of protocol.

Switch-IB 2 enables IT managers to program and centralize their server and storage interconnect management and dramatically reduce their operations expenses by completely virtualizing their data center network.

SWITCH PRODUCT DEVELOPMENT

The Switch-IB 2 Evaluation Board (EVB) and Software Development Kit (SDK) are available to accelerate OEMs’ time to market and for running benchmark tests. These rack-mountable evaluation systems are equipped with QSFP interfaces for verifying InfiniBand functionality. In addition, SMA connectors are available for SerDes characterization. The Switch-IB 2 SDK provides customers the flexibility to implement InfiniBand connectivity using a single switch device.

The SDK includes a robust and portable device driver with two levels of APIs, so the user can choose their level of integration. A minimal set of code is implemented in the kernel to allow for easy porting to various CPU architectures and operating systems such as x86 and PowerPC architectures utilizing the Linux operating system. Within the SDK, the device driver and API libraries are written in standard ANSI “C” language for easy porting to additional processor architectures and operating systems. The same SDK supports the SwitchX®-2, Switch-IB, Spectrum™ and Switch-IB 2 switch devices.

CONFIGURATIONS

Switch-IB 2 allows OEMs to deliver:

- 36-port 1U EDR InfiniBand switch
- Modular chassis switch with up to 648 EDR InfiniBand ports

MELLANOX ADVANTAGE

Mellanox is the leading supplier of industry standard InfiniBand and Ethernet network adapter silicon and cards (HCAs and NICs), switch silicon and systems, interconnect products, and driver and management software. Mellanox products have been deployed in clusters scaling to tens of thousands of nodes and are being deployed end-to-end in data centers and TOP500 systems around the world.

SPECIFICATIONS

Infiniband

- IBTA Specification 1.3 compliant
- 10, 20, 40, 56 or 100Gb/s per 4X port
- Integrated SMA/GSA
- Hardware-based congestion control
- 256 to 4KB MTU
- 9 virtual lanes:
8 data +1 management

I/O Specifications

- 36 4X SDR/DDR/QDR/FDR10/FDR/ EDR InfiniBand ports
- SPI Flash interface, I2C
- IEEE 1149.1/1149.6 boundary-scan JTAG
- LED driver I/Os
- General purpose I/Os
- 52.5 x 52.5 mm FCBGA

COMPATIBILITY

CPU

- PowerPC, Intel x86, AMD x86, MIPS

PCI Express interface

- PCIe 3.0, 2.0, and 1.1 compliant
- 2.5GT/s, 5GT/s or 8GT/s x4 link rate

Connectivity

- Interoperability with InfiniBand adapters and switches

- Active/passive copper cables, fiber optics, PCB or backplanes

Management & Tools

- Support for Mellanox and IBTA compliant Subnet Managers (SM)
- Diagnostic and debug tools
- Fabric Collective Accelerator (FCA) software library

Table 1 - Part Number and Description

OPN	Description	Typical Power
MT53236A0-FDCR-E	Mellanox Switch-IB 2, 36 Port EDR InfiniBand Switch IC	83W



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