InfiniScale® III

24-port 4X (or 8-port 12X) InfiniBand Switch supporting up to 60Gb/s per port

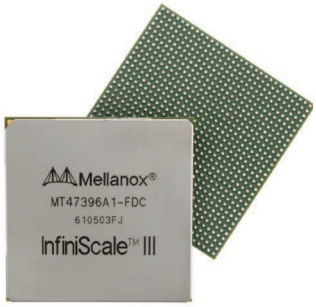

Terabit Switching Capacity in a Single-Chip

InfiniScale III is Mellanox's third generation InfiniBand switch silicon device used to create reliable, scalable, and easy to manage interconnect fabrics for compute, communication, storage, and embedded applications. The device supports twenty-four 4X ports or eight 12X InfiniBand ports or a combination of both port types.

The single data rate (SDR) version of the InfiniScale III supports 10Gb/s per 4X port and 30Gb/s per 12X port delivering 480Gb/s of aggregate bandwidth. The double data rate (DDR) version of the InfiniScale III supports 20Gb/s per 4X port and 60Gb/s per 12X port delivering 960Gb/s of aggregate bandwidth.

The architecture features an intelligent non-blocking, packet switch design with an advanced scheduling engine that provides quality of service with switching latencies of less than 200 nanoseconds. The InfiniScale III architecture also scales efficiently to allow designers to create InfiniBand switch systems that support non-blocking clusters over 10,000 nodes.

This combination of unmatched bandwidth, inherent low latency design, and the reliability of InfiniBand fabrics, makes this device an excellent building block for high bandwidth clusters for data center applications and High Performance Computing (HPC) clusters.

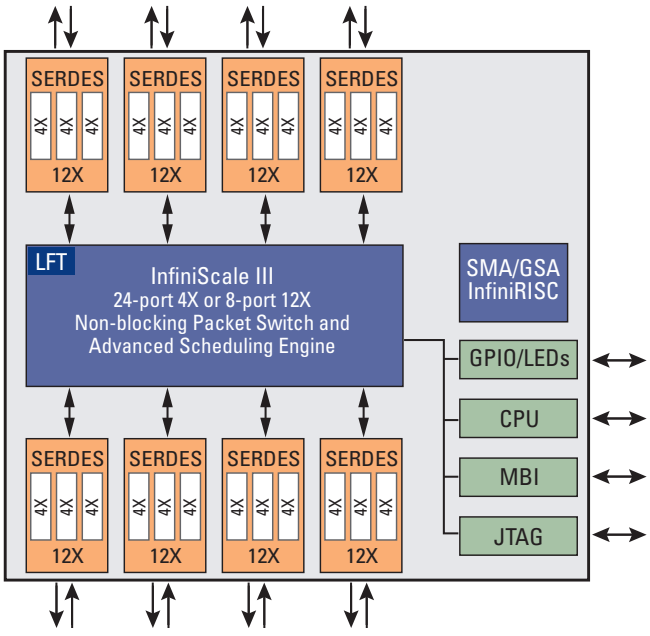

InfiniScale III Block Diagram

BENEFITS

- World-class cluster performance
- High-performance networking and storage access
- Guaranteed bandwidth and low-latency services
- Reliable transport
- I/O consolidation
- Virtualization acceleration
- Scales to tens-of-thousands of nodes

KEY FEATURES

- Twenty-four 10 or 20Gb/s InfiniBand 4X ports or eight 30 or 60Gb/s InfiniBand 12X ports (or any combination)
- 480Gb/s (SDR version) or 960Gb/s (DDR version) of total switching bandwidth
- 96 integrated 2.5Gb/s (SDR version) or 5Gb/s (DDR version) SerDes interfaces
- Auto-negotiation of port link speed
- Ultra low latency cut-through switching
- Integrated Subnet Management Agent
- Inbound and Outbound Partition (P_KEY) Checking
- Programmable port mirroring
- IBTA v1.2 compatible
- Eight virtual lanes supported plus management lane
- Multicast support for up to 1K entries
- Hardware Credit Based Link Level Flow Control
- Hardware CRC Checking and Generation
- Bad Packet Filtering
- Per port Service Level (SL) to VL mapping
- Integrated Link Packet Buffer

Complete Integrated Physical Layer

InfiniScale III integrates ninety-six 2.5Gb/s (SDR version) or 5Gb/s (DDR version) SerDes interfaces in a single 961-ball package– reducing power, system cost, and PCB size and complexity. With a single InfiniScale III device, system vendors can build a complete 24-port 10 or 20Gb/s InfiniBand switch system or server blade switch module. Multiple InfiniScale III devices can be interconnected to build much larger port count, single chassis systems.

Additional Architectural Advancements

InfiniScale III features significant advancements in its management architecture and interfaces. The local management interface is transaction-based using a small set of registers called the Command Interface. This offers a simplified API for switch management using embedded processors connected to the on-chip CPU interface. The CPU bus interface includes master and slave DMA capabilities, which increases speed and lowers CPU utilization of the embedded management solution. In addition, the internal channel adapter port has been significantly updated for performance and includes extended management features.

InfiniScale III-based InfiniBand Switch Products from Mellanox as Production Development Kits (PDK)

24-port 4X 1U InfiniBand Switch

144-port 4X InfiniBand Switch

Switch Silicon Family

Product	Switch Capacity	Switch Ports	Typical Power	Part Number
InfiniScale III SDR	480Gb/s	24-10Gb/s (4X) or 8-30Gb/s (12X)	25W	MT47396A1-FDC
InfiniScale III DDR	960Gb/s	24-20Gb/s (4X) or 8-60Gb/s (12X)	34W	MT47396A1-FDC-D

KEY APPLICATIONS

- Virtualized data centers that require a high-bandwidth, low-latency interconnect for server and storage grids
- High-performance parallelized computing leveraging Message Passing Interface (MPI) based applications such as molecular modeling, oil and gas exploration, car crash simulations, etc.
- Clustered database applications, parallel RDBMS queries, high-throughput data warehousing
- Performance storage applications such as backup, restore, mirroring, etc.
- High bandwidth streaming content such as video-on-demand and HDTV
- Electronic Design Automation (EDA)
- Networking, telecom and industrial data acquisition

PACKAGES

- Package Size: 40mm x 40mm, 1.27mm ball pitch
- Package: HFCBGA, with 961 balls
- ROHS-R5 Compliant
- Typical Power: 25 Watts (SDR version), 34W (DDR version)

INTERFACES

- PowerQUICC 8/16-bit CPU Interface
- MBI (I^2C Compatible Management Bus Interface)
- Flash interface
- Boundary Scan (JTAG) interface
- In-band InfiniBand IBML management support

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