

Mellanox Quantum™ HDR Switch Silicon

40-ports, supporting HDR 200Gb/s InfiniBand throughput per port, with a total of 16Tb/s bidirectional throughput and 15.6 billion messages per second

Mellanox continues its interconnect leadership by providing the world's smartest switch, enabling in-network computing through Co-Design 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 Quantum is the world's smartest 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.

Quantum offers industry-leading integration of 160 SerDes, with rich speed flexibility ranging from 2.5Gb/s to 50Gb/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 Quantum'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.

INFINIBAND INTERCONNECT

Quantum 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.

Quantum 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.

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 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 end-points.

It also supports aggregation of large data vectors at wire speed to enable MPI large vector reduction operations, which are crucial for machine learning applications.



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 16Tb/s of switching capacity
- Up to 15.6 billion messages per second
- Up to 40 HDR 200Gb/s InfiniBand ports
- Collective communication acceleration
- Hardware-based adaptive routing
- Hardware-based congestion control

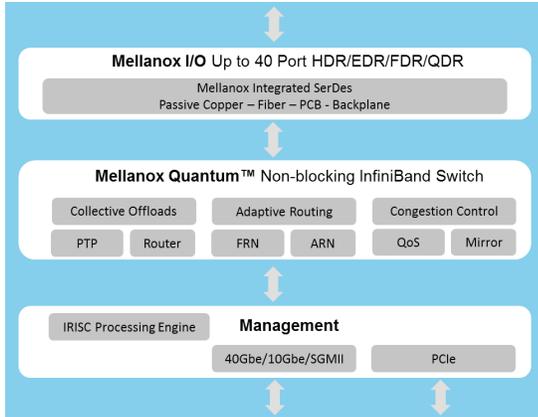


Figure 1. Mellanox Quantum Architecture

TELEMETRY

Visibility is a critical component to a more efficient network. Capturing what a network is ‘thinking’ or ‘doing’ is the basis for true network automation and analytics.

In particular, today’s HPC and cloud networks require fine-grained visibility into:

- Network state in real-time
- Dynamic workloads in virtualized and containerized environments
- Advanced monitoring and instrumentation for troubleshooting

Mellanox Quantum was designed for maximum visibility using such features as mirroring, sFlow, congestion based mirroring and histograms.

SWITCH PRODUCT DEVELOPMENT

The Quantum 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 QSFP55 interfaces for verifying InfiniBand functionality. In addition, SMA connectors are available for SerDes characterization. The Quantum 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 Quantum switch devices.

CONFIGURATIONS

Quantum allows OEMs to deliver:

- 40-port 1U HDR 200Gb/s InfiniBand switch
- Modular chassis switch with up to 800 HDR 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, 100 or 200Gb/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

- SPI Flash interface, I²C
- IEEE 1149.1/1149.6 boundary-scan JTAG
- LED driver I/Os
- General purpose I/Os
- 55 x 55mm HFCBGA

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

- 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
MT54240A0-FCCR-H	Mellanox Quantum™ 40 port 200G HDR InfiniBand Switch ASIC	190W



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