




## NVIDIA MELLANOX QUANTUM HDR 200G INFINIBAND SWITCH SILICON



NVIDIA® Mellanox® Quantum™ switch silicon offers 40 ports supporting HDR 200 Gb/s InfiniBand throughput per port, with a total of 16 Tb/s bidirectional throughput and 15.6 billion messages per second.

Mellanox Quantum is the world's smartest network switch that enables in-network computing through the co-design Scalable Hierarchical Aggregation and Reduction Protocol (SHARP)™ technology. Its co-design architecture enables the usage of all active data center devices to accelerate communication frameworks, resulting in an order of magnitude improvement in application performance and enabling the highest performing server and storage system interconnect solutions for Enterprise Data Centers, Cloud Computing, High-Performance Computing, and Embedded environments.

Mellanox Quantum embeds an innovative solution called SHIELD™ (Self-Healing Interconnect Enhancement for Intelligent Datacenters) that makes the fabric capable of self-healing autonomy. So, the speed at which communications can be corrected in the face of a link failure can be increased by 5000X, making it fast enough to save expensive retransmissions or absolute communications failure.

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

### Key Features

- > 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 16 Tb/s of switching capacity
- > Up to 15.6 billion messages per second
- > Up to 40 HDR 200 Gb/s InfiniBand ports
- > Collective communication acceleration
- > Hardware-based adaptive routing
- > Hardware-based congestion control
- > Mellanox SHARP™v2 collective offloads support streaming for Machine Learning
- > SHIELD-enabled self-healing technology

## INFINIBAND INTERCONNECT

Mellanox Quantum InfiniBand devices enable industry standard networking, clustering, storage, and management protocols to seamlessly operate over a single “one-wire” converged network. Combined with the Mellanox 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. Mellanox 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 Mellanox 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 endpoints.

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

## TELEMETRY

Visibility is a critical component of an 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 is designed for maximum visibility using such features as mirroring, sFlow, congestion based mirroring, and histograms.

## SWITCH PRODUCT DEVELOPMENT

The Mellanox Quantum Evaluation Board (EVB) and Software Development Kit (SDK) are available to accelerate an OEM’s time to market and for running benchmark tests. These rack-mountable evaluation systems are equipped with QSFP56 interfaces for verifying InfiniBand functionality. In addition, SMA connectors are available for SerDes characterization. The Mellanox 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 developers 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 Mellanox SwitchX®-2, Switch-IB®, Switch-IB 2, Mellanox Spectrum®, and Mellanox Quantum switch devices.

## Compatibility

### CPU

- > PowerPC, Intel x86, AMD x86, MIPS

### PCI Express Interface

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

### Connectivity

- > Interoperability with InfiniBand adapters and switches
- > Passive copper cables, fiber optics, PCB or backplanes

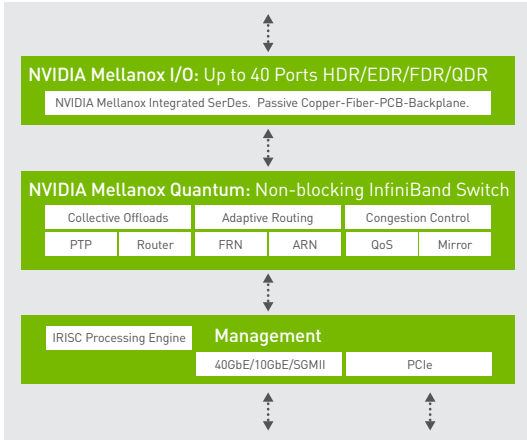
### Management & Tools

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

## CONFIGURATIONS

Mellanox Quantum allows OEMs to deliver:

- > 40-port 1U HDR 200 Gb/s InfiniBand switch
- > 80-port 1U HDR100 100 Gb/s InfiniBand switch
- > Modular chassis switch with up to 800 HDR InfiniBand ports
- > Modular chassis switch with up to 1600 HDR100 InfiniBand ports



## NVIDIA MELLANOX ADVANTAGE

NVIDIA 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.

## ORDERING INFORMATION

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

## Specifications

### InfiniBand

- > IBTA Specification 1.4 compliant
- > 10, 20, 40, 56, 100 or 200 Gb/s per 4X port
- > Integrated SMA/GSA
- > Hardware-based congestion control
- > 256 to 4 KB MTU
- > 9 virtual lanes: 8 data +1 management

### I/O Specifications

- > SPI Flash interface, I<sup>2</sup>C
- > IEEE 1149.1/1149.6 boundary scan JTAG
- > LED driver I/Os
- > General purpose I/Os
- > 55 x 55 mm HFCBGA

Learn more at [www.mellanox.com/products/infiniband-switches-ic/quantum](http://www.mellanox.com/products/infiniband-switches-ic/quantum)

© 2020 Mellanox Technologies. All rights reserved. NVIDIA, the NVIDIA logo, Mellanox, Mellanox Quantum, Mellanox Spectrum, SwitchX, SwitchIB, ConnectX, Scalable Hierarchical Aggregation and Reduction Protocol (SHARP) and SHIELD are trademarks and/or registered trademarks of Mellanox Technologies Ltd. and/or NVIDIA Corporation in the U.S. and in other countries. Other company and product names may be trademarks of the respective companies with which they are associated. Sep20/53361PB-R2

