



KEY FEATURES

Scale and Bandwidth

- 36 40Gb/s InfiniBand 4X ports
- 12 80Gb/s InfiniBand 8X ports
- 12 120Gb/s InfiniBand 12X ports
- Port speed auto-negotiation
- Line rate switching
- 2.88 Tb/s total switching capacity
- 144 integrated SerDes interfaces, each operate up to 10+10Gb/s

Latency

- 60ns switch latency
- Reduced number of hops for large scale clusters

Performance Enhancements


- Adaptive Routing
- Congestion Control
- Quality of Service

Manageability

- PCIe 2.0 x4 Management Processor Interface
- InfiniBand Standard In-band Management
- Management processor is optional
- Integrated Subnet Management Agent (SMA)
- Integrated General Service Agent (GSA)
- Advanced Port Mirroring

InfiniScale™ IV

Fourth Generation InfiniBand Switch Architecture



3 Terabits per Second Switching

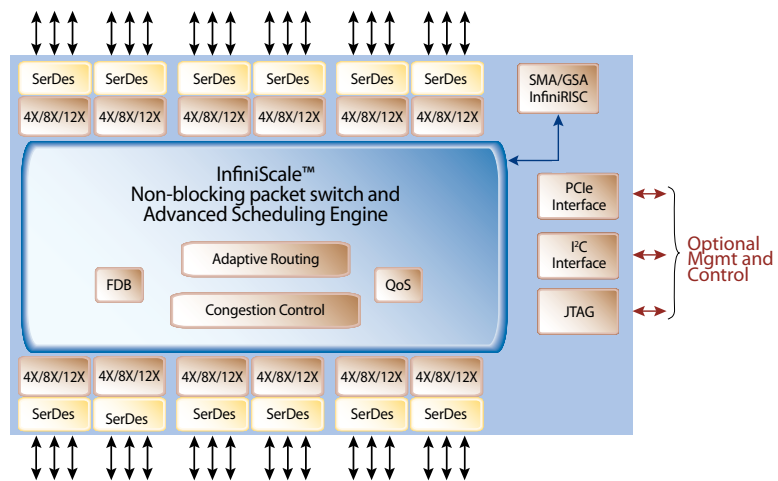
InfiniBand provides the highest bandwidth, lowest latency, and most scalable interconnect for servers and storage. The fourth generation switching architecture from Mellanox, InfiniScale IV, improves these leading attributes further, making InfiniBand an obvious choice for the most demanding applications. InfiniScale IV silicon products using this architecture will be used by system architects to construct some of the world's most powerful computing systems, and for carrying converged traffic with the best combination of bandwidth and granular quality of service.

The InfiniScale IV architecture extends the performance and capabilities of InfiniBand switching across several dimensions: more bandwidth, lower latency, more scale, and improved features. Extensive experience in building systems using the InfiniScale III and previous generations of InfiniBand switching products provide a wealth of field-tested experience, leading to insights in how

to optimize performance and extend functionality.

With 36 high-performance 40Gb/s ports, fewer switch chips are needed than with customary 24 port building blocks. Besides the obvious benefit of lowering parts count, another benefit is gained – constant bisectional bandwidth (CBB) fabrics can use fewer tiers (layers) of switches, reducing the hop count and latency. Coupled with the lower latency inherent in the InfiniScale IV architecture, large clusters run with lower overall latency and better overall application performance.

Adaptive Routing has been added to the static routing previously provided by the architecture. InfiniBand is optimized to move traffic via multiple parallel paths, and adaptive routing uses immediate feedback within the switch chip to move traffic from congested paths to less congested ones. This is done at in-chip hardware speeds, to avoid the problems created if this feedback loop is too slow. In cases where traffic patterns are more predictable, static routing has been shown



InfiniScale IV Block Diagram

InfiniScale™ IV Architecture

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

INFINIBAND BENEFITS

- Industry-standard technology
- Unified computing, storage and management
- High-bandwidth, low-latency
- Highly-efficient clustering
- Ultimate reliability and scalability
- Multi-platform support
- Congestion management and QoS
- Virtualized I/O fabric
- World-Class price/performance

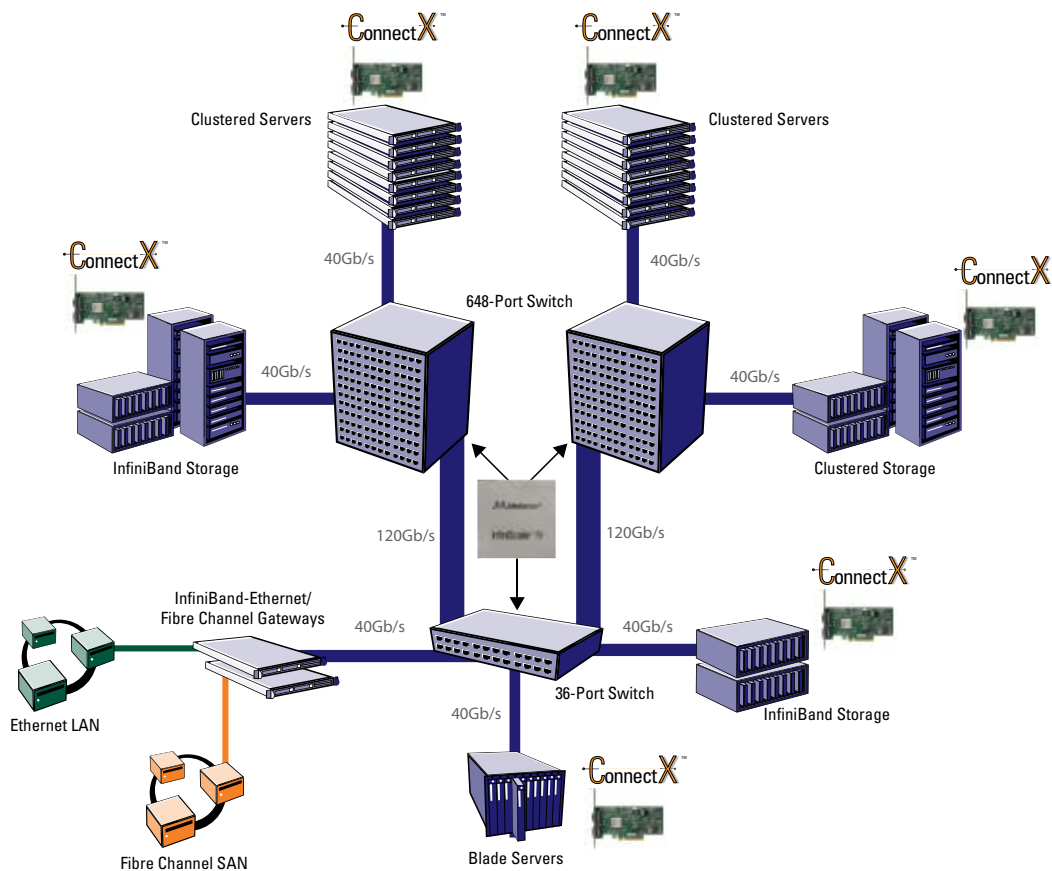
to produce superior results. The InfiniScale IV architecture provides the best of both static and adaptive routing.

Some switching contention is unavoidable – for example when two sources are trying to reach a single destination. Congestion control, using InfiniBand 1.2 standard mechanisms, is the only proven way to control hotspots in large fabrics. InfiniScale IV works in conjunction with ConnectX IB adapters to control sources when hotspots are detected, keeping traffic moving while keeping latency low and constant.

For converged traffic, the combination of high bandwidth, adaptive routing, and congestion control provide the industry's best traffic carrying capacity. Adding Quality of Service makes sure that traffic classes can be protected, guaranteeing the delivery of critical traffic

How to Build the Best Fabric

Start with the fastest fabric. Add Adaptive Routing, Congestion Control, and Quality of Service. Reduce latency until it's nearly non-existent. That's the compelling benefit of the InfiniScale IV architecture.



40Gb/s and 120Gb/s Connected InfiniBand Server and Storage Diagram



2900 Stender Way, Santa Clara, CA 95054
 Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com