Mellanox BridgeX™ VPI

10/20/40Gb/s InfiniBand to 10 GigE and 2/4/8Gb/s FC Gateway
10 GigE (FCoE) to 2/4/8G FC Gateway

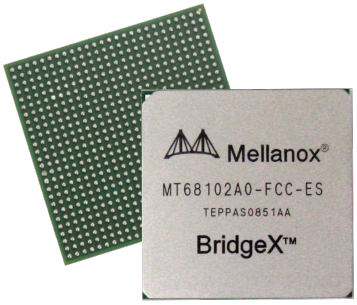

Mellanox BridgeX is the first VPI (Virtual Protocol Interconnect) gateway allowing the OEMs to design I/O consolidation solutions using InfiniBand or Ethernet as the convergence fabric of choice. A unified server I/O, where multiple traffic types can run over a single physical connection can help cut I/O cost and power significantly while reducing total cost of ownership through reduced number of ports to manage, reduced cabling complexity and simpler fabric management. At the same time, connectivity to IP/Ethernet based LAN and NAS infrastructures, and Fibre Channel based SAN infrastructures must remain seamless. The Mellanox BridgeX™ silicon enables the delivery of utility computing by consolidating server I/O using InfiniBand and exposing 10 GigE interfaces for connectivity to LAN/NAS infrastructures, and Fibre Channel interfaces for connectivity to SAN infrastructures.

Unparalleled Performance

The Mellanox BridgeX device supports two 10/20/40Gb/s InfiniBand ports that connect to an InfiniBand switch. The cumulative uplink bandwidth of 80Gb/s is shared by the six 10 GigE downlink ports or eight Fibre Channel downlink ports supporting data rates from 2Gb/s to 8Gb/s.

Efficient and Scalable Design

The Mellanox BridgeX device is built using a novel stateless approach where the gateway functions are accomplished as a simple packet relay and without requiring maintenance of any stateful information. Ethernet and Fibre Channel sessions are initiated and terminated at end points (initia-

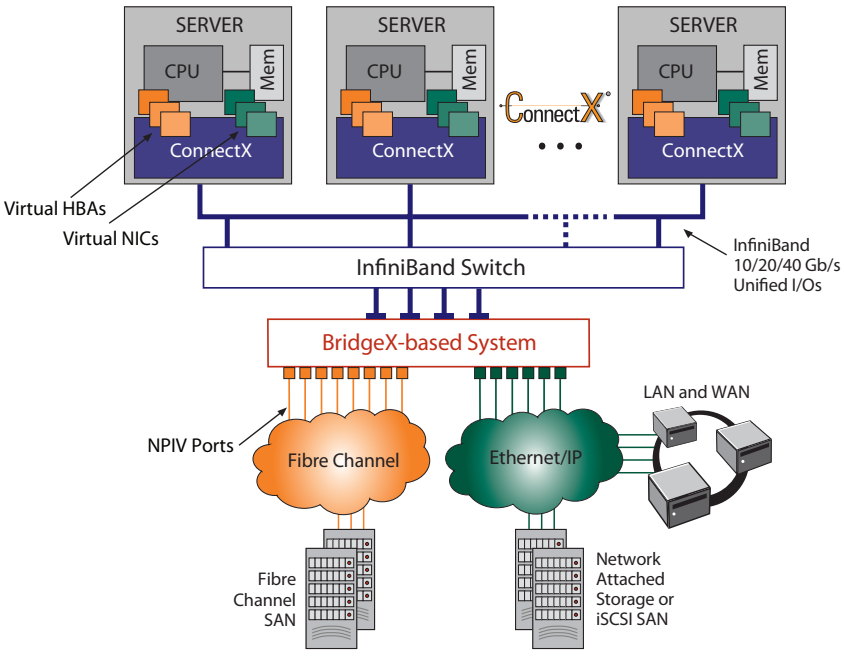

BridgeX System Deployment Scenario

BENEFITS

- Delivers I/O consolidation on high performance 40G InfiniBand or on 10GigE
- Single chip solution for I/O consolidation
- Low power, low cost and high performance
- Virtualized 10GigE and 8G FC ports

KEY FEATURES

HARDWARE FEATURES

- Two InfiniBand 10/20/40Gb/s or Six 10GigE uplink ports
- Six 10GigE downlink ports, or eight 2/4/8Gb/s FC downlink ports, or three 10GigE and four FC downlink ports
- 1024 virtual NICs per Ethernet port
- 1024 virtual HBAs per FC port
- Total of 8K MAC, VLAN addresses
- Total of 8K WWN addresses
- Interoperable with InfiniBand, Ethernet and Fibre Channel Switches
- PCI Express

SOFTWARE FEATURES

- Virtual NIC / HBA configuration
- WWN address, NPIV, VLAN configuration
- APIs for interfacing to third-party management
- Virtual topology maintenance
- FC MIBs, Ethernet MIBs and Bridge MIBs

tors on the server side and targets on the LAN/SAN side) only. The result is a high-performance, low cost and low power design.

Virtualized I/O

BridgeX device supports I/O edge virtualization. A layer of abstraction is inserted between the InfiniBand ports on the servers and the external networks that connect to the downlink Ethernet and Fibre Channel ports. In a data center environment, the server administrator can assign I/O connections to the downlink Ethernet and Fibre Channel ports, and the LAN and SAN administrator treat these downlink ports as the end points of their networks. This cleanly separates server enclosure administration (such as in blade servers that includes the BridgeX silicon) from LAN and SAN administration. It also facilitates independent scaling of compute, LAN and SAN capacity.

Comprehensive Management Interface

Host management is simplified because the software driver and application interfaces are identical to Ethernet/IP and Fibre Channel/SCSI interfaces. Configuration of the Mellanox BridgeX silicon is simplified by using Mellanox provided device drivers, APIs, and command-line interface (CLI) that allow configuration of the uplink and downlink ports and mapping of resources and forwarding tables related to the ports. By using the APIs and CLI, OEMs can build management applications and GUI to configure ports and their connectivity to external networks and servers.

SPECIFICATIONS

UPLINK PORTS
- 2 ports of 10/20/40G InfiniBand or 6 ports of 1/10 GigE

DOWNLINK PORTS
- 6 ports of 1/10 GigE or 8 ports of 2/4/8G FC ports

INFINIBAND PORTS
- IBTA specification 1.2 compliant
- 10Gb/s (SDR), 20Gb/s (DDR) and 40Gb/s (QDR) transfer rate
- Automatic link speed negotiation
- Hardware-based congestion control

ETHERNET PORTS
- IEEE 802.3z Gigabit Ethernet
- XAUI, XFI and SGMII support
- IEEE 802.3ak 10GBASE-CX4
- IEEE 802.3ae 10GBASE-SR/LR
- IEEE 802.3ap 10GBASE-KR (including FEC), 10GBASE-KX4, 1000BASE-KX
- Automatic link speed negotiation 1/10 Gbps
- IEEE 802.3x Pause
- IEEE 802.3ad Link Aggregation and Failover
- IEEE 802.1p and DIFFSERV Priorities
- IEEE 802.1Q VLAN tags
- Data Center Bridging (DCB)
  - IEEE 802.1Qbb Per Priority Flow Control (PFC)
  - IEEE 802.1Qau Congestion Notification
  - IEEE 802.1Qaz Enhanced Transmission Selection
- 1K MAC/VLAN ID addresses per port
- 1K Ethernet virtual end points (Virtual NICs) per port
- Total of 8K (MAC, VLAN ID) addresses
- Multicast support
- IEEE 802.1D Spanning Tree
- IEEE 802.1AB Link Layer Discovery

FC PORTS
- SCSI-3, SCSI-FCP, FC-FS-3, FC-PI-4
- 2/4/8 Gb/s per port
- Full-duplex 64 Gb/s aggregate throughput
- Speed negotiation
- Configured as Node port (N_Port) or Fabric port (F_Port / E_Port)
- 1K virtual end points (Virtual HBAs) per Fibre Channel port
- Class 2 and Class 3 Fibre Channel services
- Class F Fibre Channel services (E_Port only)
- Fibre Channel over Ethernet (FCoE) protocol encapsulation (T11 compliant)

PCI EXPRESS INTERFACE
- PCIe Base 2.0 compliant, 1.1 compatible
- 2.5GT/s or 5.0GT/s link rate (per lane)
- Auto-negotiation to x4, x2, or x1
- Support for MSI-X mechanisms

OTHER INTERFACES
- Flash Memory Interface
- I2C Interface
- GPIO
- MDIO
- LEDs Interface

PACKAGE
- 31 x 31 HFCBGA
- RoHS5 compliant

*Only six 2/4/8G FC ports when the uplink ports are 10GigE

Gateway Silicon

Ordering Part Number	Uplink Port	Downllink Port	Power (Typ)
MT68102A0-FCC-VEF	2 x 10/20/40G InfiniBand or 6 x 1/10GigE	6x 10GigE, 8x FC, 3x 10GigeE + 4x FC	13.83*W *40G IB to 10G XFI

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