



Paving The Road to Exascale Computing

brandon@mellanox.com

HPC User Forum

April 16, 2012

Enterprise HPC



High-end HPC




HPC Clouds

Mellanox Interconnect Networking Solutions

ICs



Adapter Cards



Host/Fabric Software



Switches/Gateways



Cables




Introducing FDR 56Gb/s InfiniBand Solutions



2002
10Gb/s



2005
20Gb/s



2008
40Gb/s



2011
56Gb/s



Highest Performance, Reliability, Scalability, Efficiency

Mellanox Virtual Protocol Interconnect Solutions for PCIe 3.0


ConnectX[®] 3



Applications



SW Acceleration Products



- **Industry leader**
 - PCI Express 3.0
 - Dual-port FDR IB or 40GbE
 - Native RDMA
 - CORE-Direct

SwitchX[™]

Unified Fabric Manager

Switch OS Layer



- **Industry leader**
 - 36 x FDR IB or 40GE 64 x 10GbE
 - Integrated routers and bridges
 - 4Tbit switching capacity
 - Ultra-low latency
 - Switch systems: from 36-port to 648-port

FDR INFINIBAND TECHNOLOGY

THE NEXT GENERATION OF
HIGH-PERFORMANCE SCALABLE CONNECTIVITY

FDR InfiniBand New Features and Capabilities

Performance / Scalability

- >12GB/s bandwidth, <0.7usec latency
- PCI Express 3.0
- InfiniBand Router and Ethernet Bridging

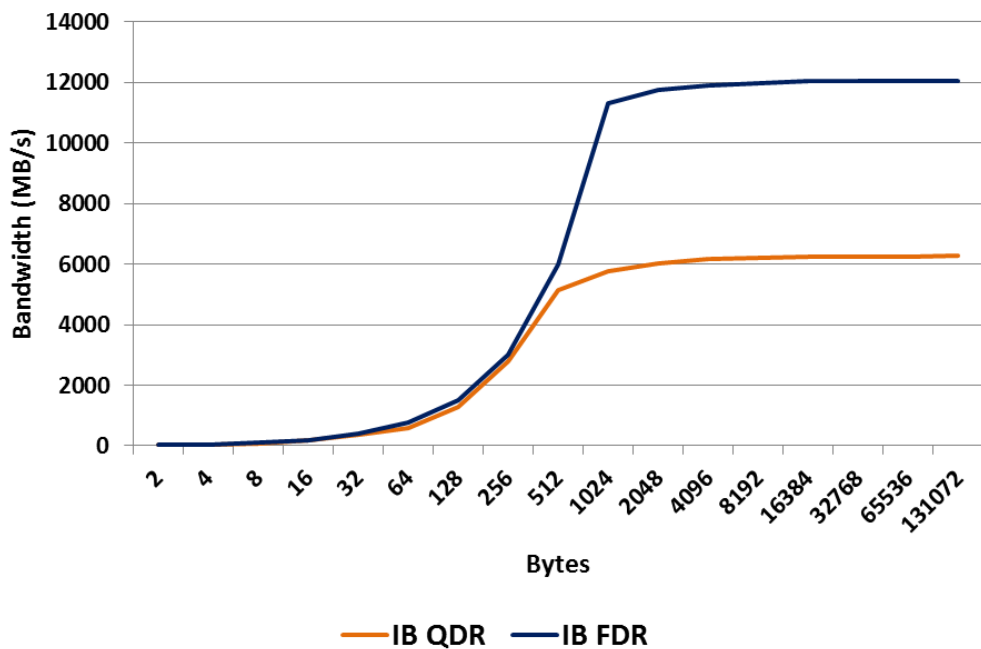
Reliability / Efficiency

- Link bit encoding – 64/66
- Forward Error Correction
- Lower power consumption

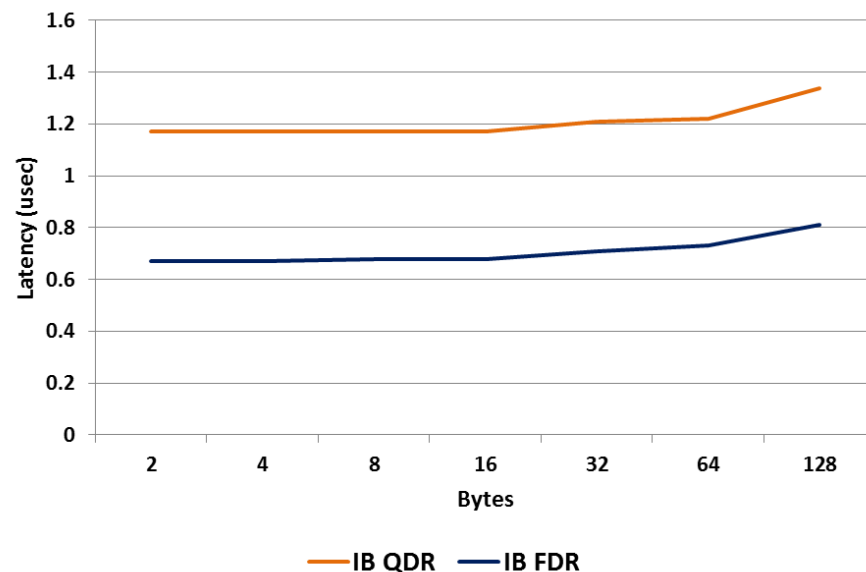
FDR InfiniBand PCIe 3.0 vs QDR InfiniBand PCIe 2.0



InfiniBand Bandwidth Bi-Directional




InfiniBand Latency



**Double the Bandwidth, Half the Latency
120% Higher Application ROI**

**187Tflops with Only 648 Nodes
#54 on the TOP500 list**

PURDUE
UNIVERSITY



**And
More!**

FDR InfiniBand is The New Deal!

FDR 56Gb/s InfiniBand Solutions Portfolio



Modular Switch



648p





324p



216p

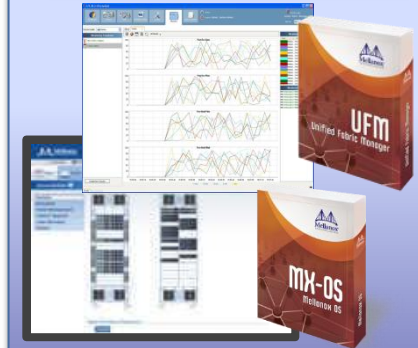
SX6500
Up to 648 ports FDR

Adapters



Single and dual port
FDR adapter card

Management



**Mellanox M-1
E-2-E Cluster
Support Services**

Edge Switch

SX60XX
1U 36 port




SX6036 – 36 ports managed



SX6025 – 36 ports externally managed

Cables



Feature License Keys




UFM Diagnostics
Virtual Protocol Interconnect

EMC 1000-Node Analytic Platform


Accelerates Industry's Hadoop Development

24 PetaByte of physical storage

Mellanox VPI Solutions (FDR InfiniBand, 40GbE)



- *High Bandwidth*
- *Low Latency*
- *RDMA*
- *Mellanox UDA*
- *Mellanox VPI*



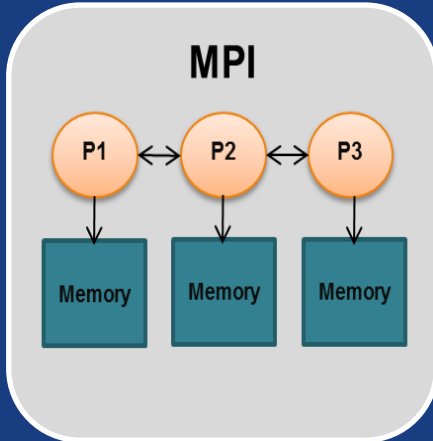
**Hadoop
Acceleration**

*2X Faster Hadoop
Job Run-Time*

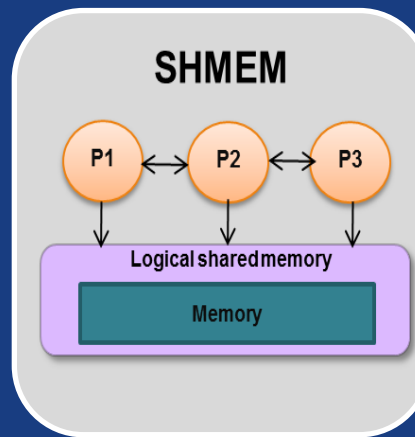
**MemCached
Acceleration**

*Up to 13X More
Transactions/Sec*

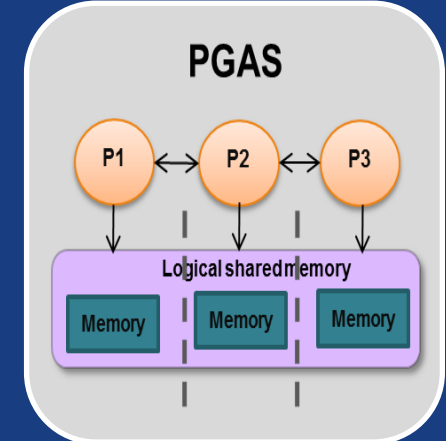
Expanding HPC Frontiers!



ScableMPI



ScalableSHMEM




ScalableUPC

Mellanox Co-Design

MellanoX Messaging Accelerations (MXM) , Collectives Offloads, GASnet , OpenSHMEM , RDMA

Today - Servers with GPUs



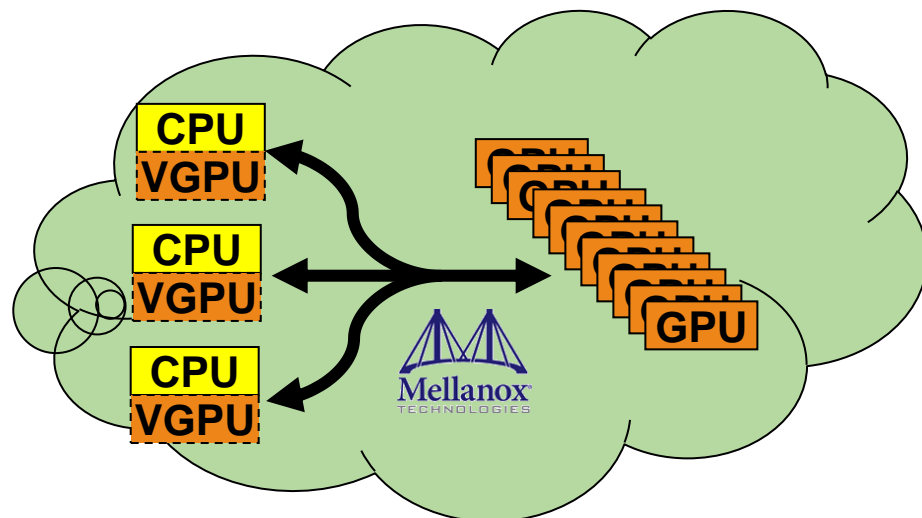
PCIe-equivalent performance

- 56Gb/s bandwidth
- 0.7usec latency


RDMA dwarfs overhead

- Maintains local access model
- Supports memory management

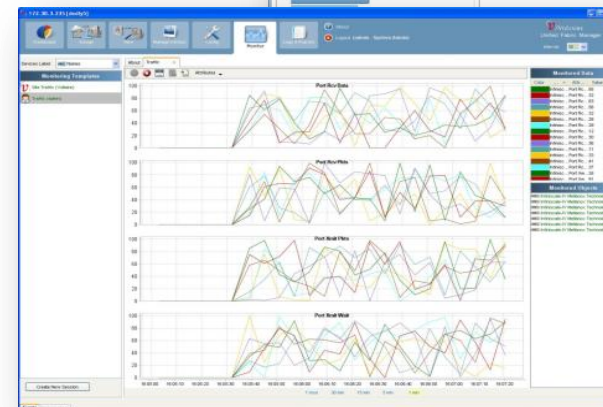
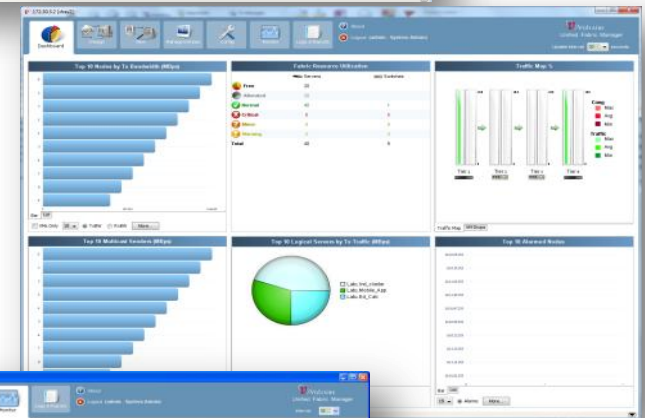
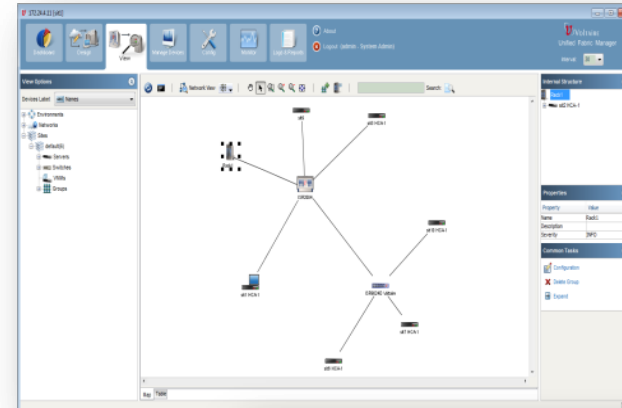
GPU as a Service



Time for matrix-matrix product (4096x4096)



- ***Monitor & analyze fabric performance***
- ***Real-time fabric-wide health monitoring***
- ***Innovative congestion analysis***
- ***Threshold based alerts***
- ***Scales out to Tens of Thousands of nodes***



Thank You

HPC@mellanox.com

PAVING THE ROAD
TO **EXASCALE**

ADVANCING NETWORK PERFORMANCE,
EFFICIENCY, AND SCALABILITY.