



Mellanox HPC Presentation

Mellanox Confidential

April 2013

 **Mellanox**
TECHNOLOGIES
Connect. Accelerate. Outperform.™

Leading Supplier of End-to-End Interconnect Solutions



Comprehensive End-to-End InfiniBand and Ethernet Portfolio

ICs	Adapter Cards	Switches/Gateways	Host/Fabric Software	Cables

- InfiniBand has become the de-facto interconnect solution for HPC
 - The most used interconnect on the TOP500 list – 224 systems
- InfiniBand connects 10 out of 23 of the world's sustained Petaflop systems
- The most used interconnect solution in the TOP100, TOP200, TOP300, TOP400
 - Connects 52% (52 systems) of the TOP100 while Ethernet only 4% (4 system)
 - Connects 54% (108 systems) of the TOP200 while Ethernet only 12.5% (25 systems)
 - Connects 49.7% (149 systems) of the TOP300 while Ethernet only 25% (75 systems)
 - Connects 46% (184 systems) of the TOP400 while Ethernet only 34% (136 systems)
- Mellanox FDR InfiniBand is the fastest interconnect solution on the TOP500
 - Used in 45 systems on the TOP500 list – more than double since the June 2012 list
 - Connects the fastest InfiniBand-based supercomputers – 2.9 Petaflops (LRZ), 2.6 Petaflops (TACC)



Mellanox InfiniBand Paves the Road to Exascale



PetaFlop
Mellanox Connected



Mellanox 56Gb/s FDR Technology

FDR INFINIBAND TECHNOLOGY

THE NEXT GENERATION OF
HIGH-PERFORMANCE SCALABLE CONNECTIVITY

FDR InfiniBand Features and Capabilities

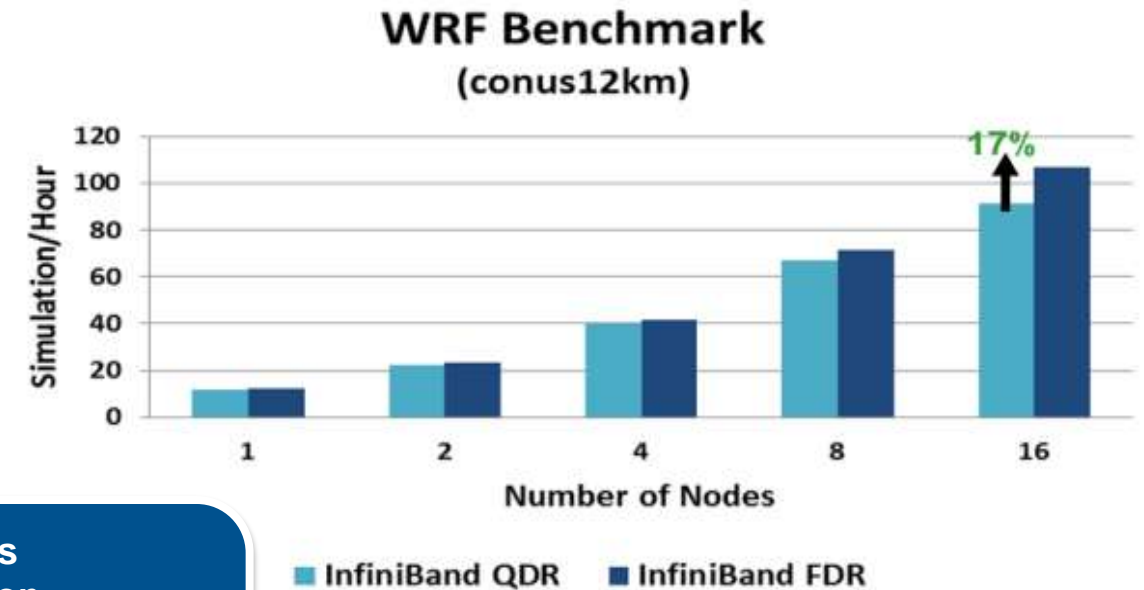
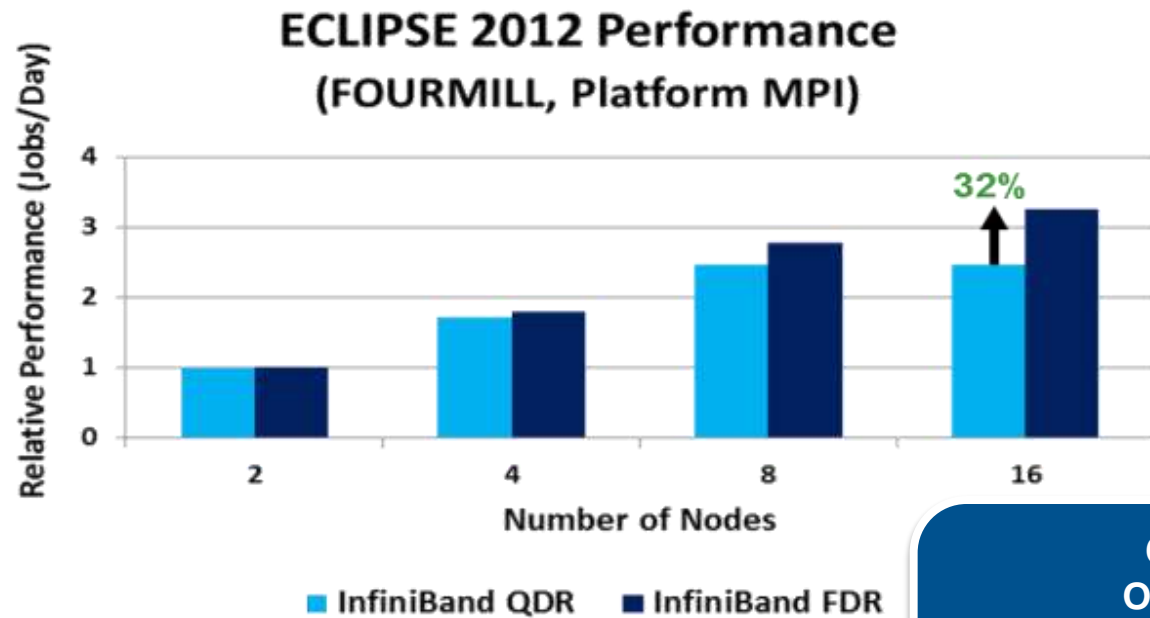
Performance / Scalability

- 56Gb/s bandwidth
- <0.7usec latency
- PCI Express 3.0
- InfiniBand Routing and IB-Ethernet Bridging

Reliability / Efficiency


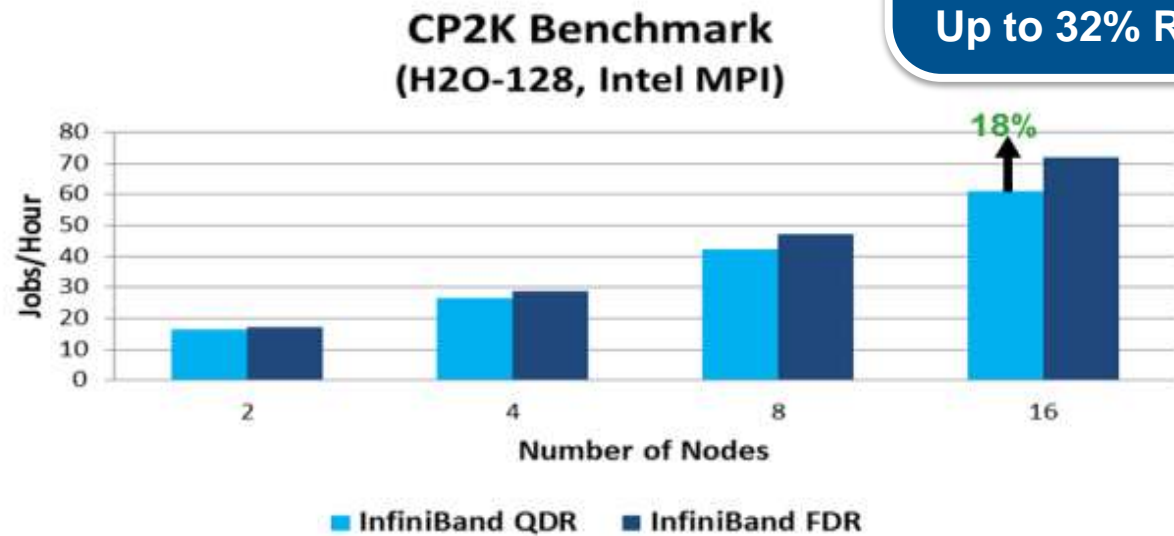
- Unbreakable Link Technology
 - Forward Error Correction
 - Link quality auto-sensing
 - Link Level Retransmission
- Link bit encoding – 64/66
- Lower power consumption

FDR Application Benchmarks Highlight ROI



Computational Fluid Dynamics
Oil and Gas Reservoir Simulation
Molecular Dynamics
Weather and Earth Sciences

Up to 32% ROI on equipment/operating costs



Connect-IB

Architectural Foundation for Exascale Computing

- A new interconnect architecture for compute intensive applications
- World's fastest server and storage interconnect solution
- Enables unlimited clustering (compute and storage) scalability
- Accelerates compute-intensive and parallel-intensive applications
- Optimized for multi-tenant environments of 100s of Virtual Machines per server

Enter the World of Scalable Performance

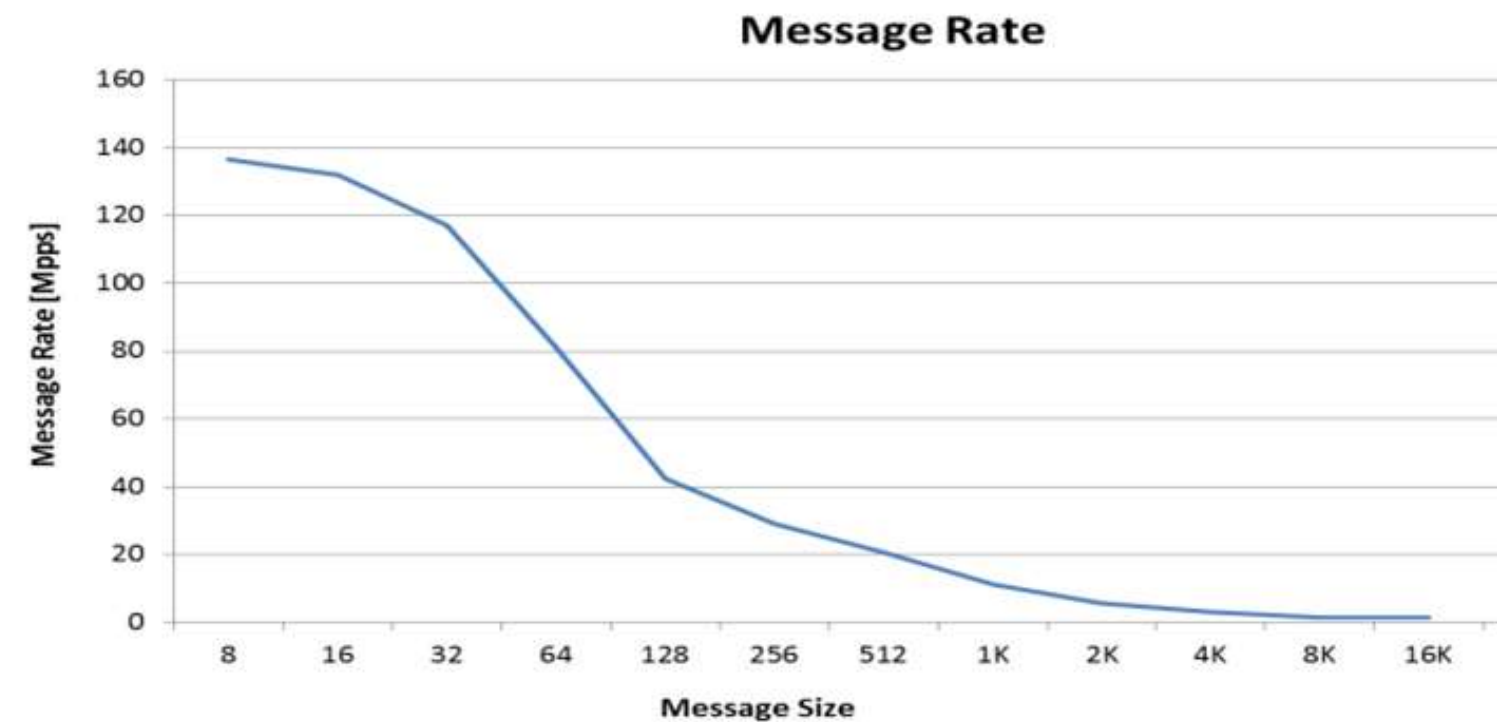
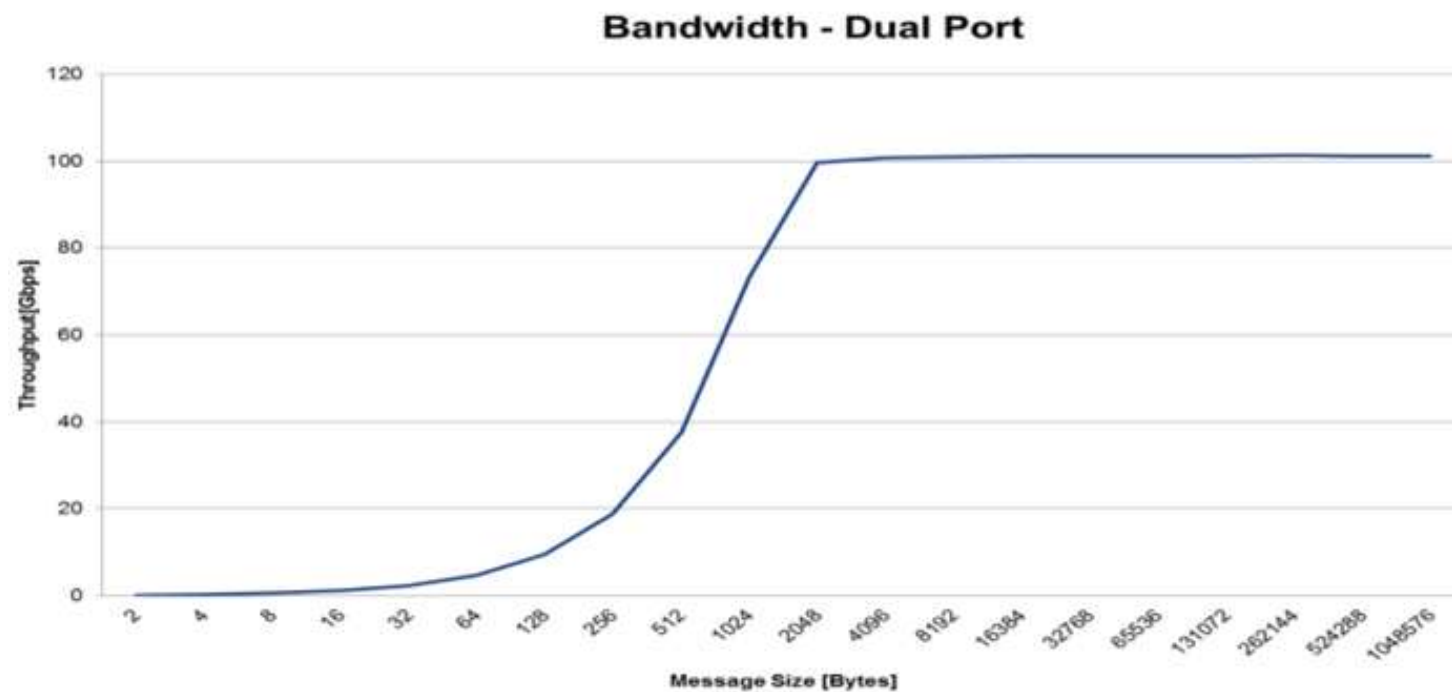
Connect **IB**[™]



Connect-IB Performance Highlights




- World's first 100Gb/s InfiniBand interconnect adapter
 - PCIe 3.0 x16, dual FDR 56Gb/s InfiniBand ports to provide >100Gb/s
- Highest InfiniBand message rate: 137 million messages per second
 - 4X higher than other InfiniBand solutions



Unparalleled Throughput and Message Injection Rates

■ New innovative transport – Dynamically Connected Transport service

- The new transport service combines the best of:
 - Unreliable Datagram (UD) – no resources reservation
 - Reliable Connected Service – transport reliability
- Scale out for unlimited clustering size of compute and storage
- Eliminates overhead and reduces memory footprint
- Interconnect resources requirement is constant, not influenced by number of nodes

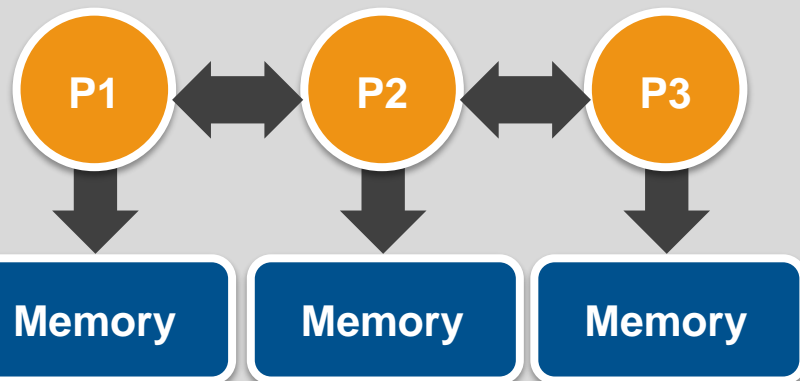


Scalable Transport Mechanisms


Mellanox ScalableHPC

Accelerations for Parallel Programs


MPI



SHMEM



PGAS



MXM

- Reliable Messaging Optimized for Mellanox HCA
- Hybrid Transport Mechanism
- Efficient Memory Registration
- Receive Side Tag Matching

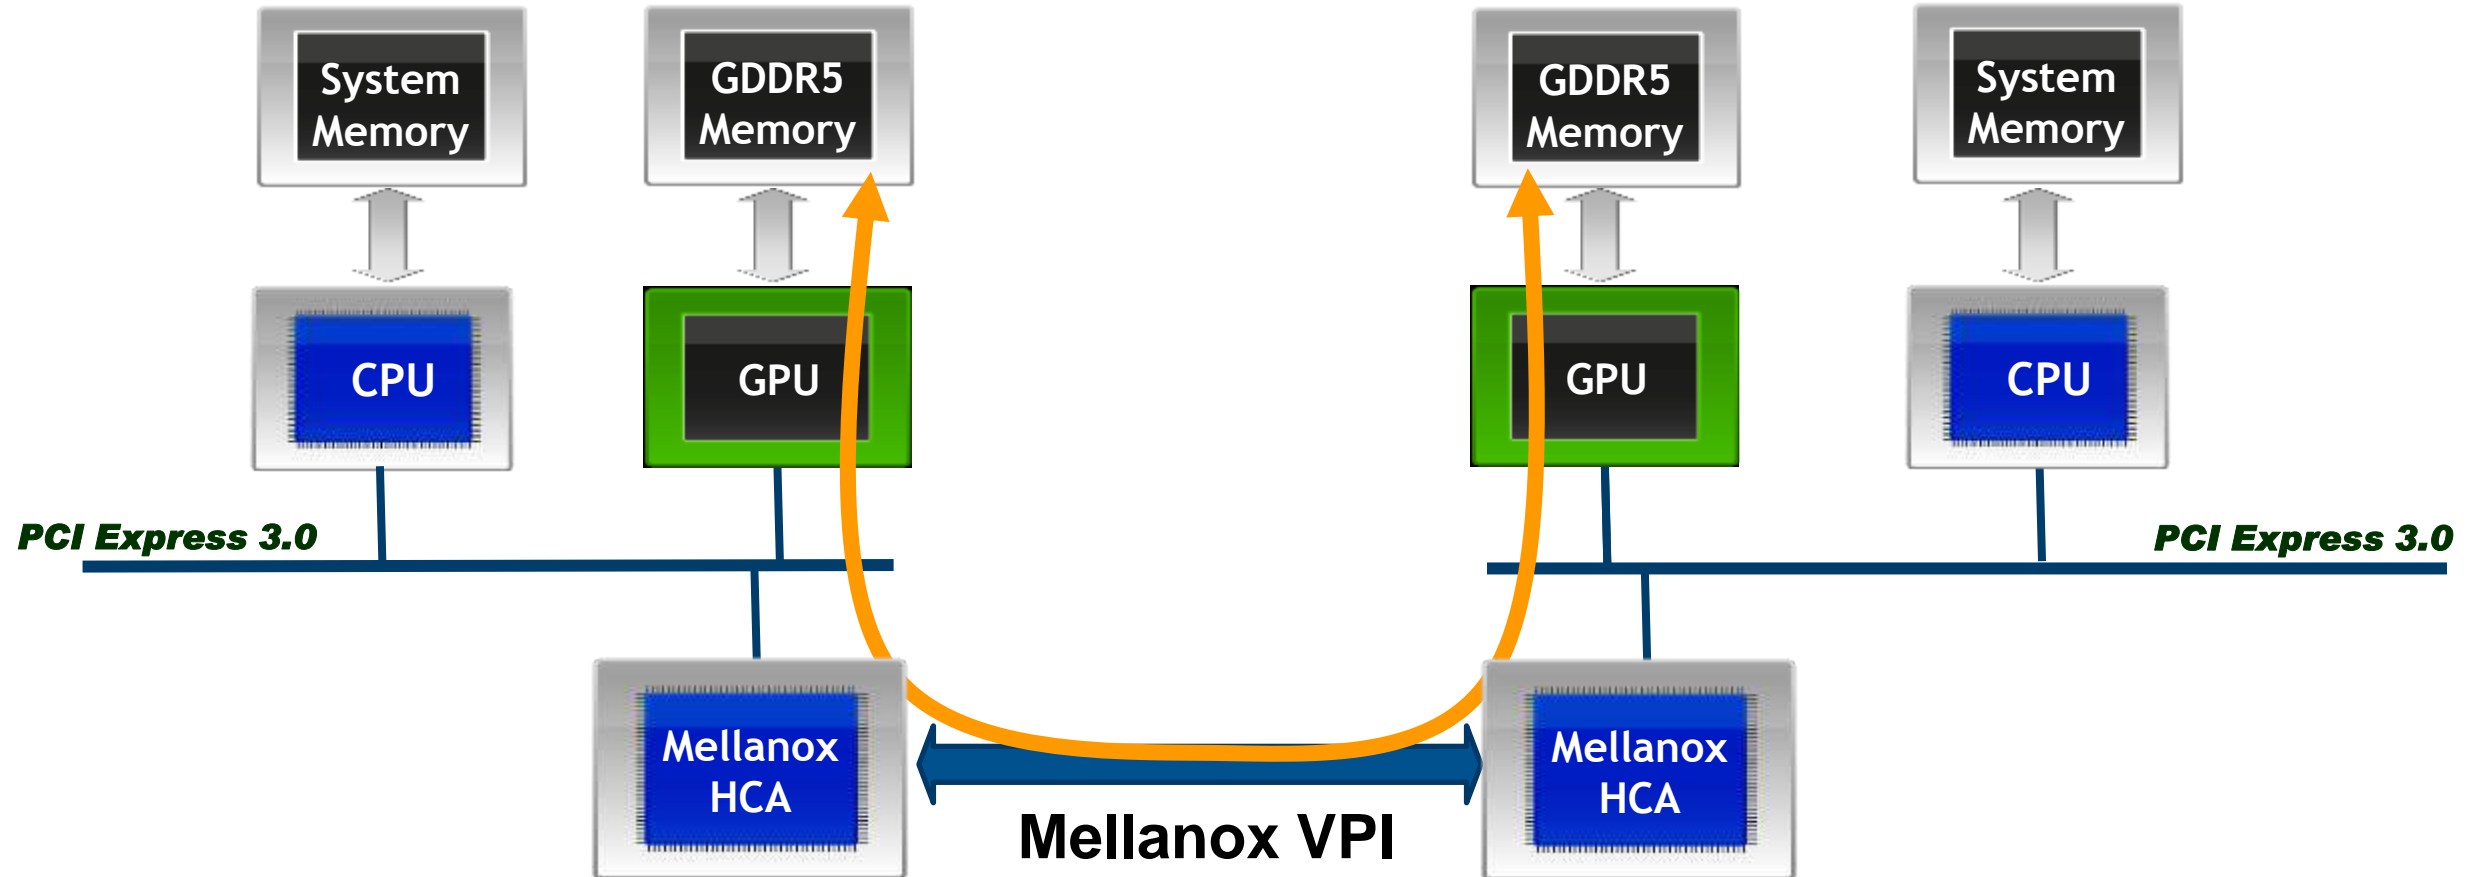
FCA

- Topology Aware Collective Optimization
- Hardware Multicast
- Separate Virtual Fabric for Collectives
- CORE-Direct Hardware Offload

InfiniBand Verbs API

Accelerator and GPU Offloads

- GPU Direct RDMA (previously known as GPU Direct 3.0)
- Enables peer to peer communication directly between HCA and GPU
- Dramatically reduces overall latency for GPU to GPU communications



Thank You

