

# Mellanox Technologies

## Maximize World-Class Cluster Performance

April, 2008

Gilad Shainer – Director of Technical Marketing

The logo for the HPC User Forum, featuring the text "HPC" in a large, stylized font with a grid pattern, and "USER FORUM" in a smaller, solid font below it.



CONFIDENTIAL

# Powerful Trends Driving Large Markets



- Connectivity Requirements
- Industry Standard
- Performance
- Price
- Power
- Reliable
- Efficient
- Scalable
- Ease



# ConnectX: Leading IB and 10GigE Adapters



- **Server and Storage Interconnect**
  - Highest InfiniBand and 10GigE performance
- **Single chip or slot optimizes cost, power, footprint, reliability**
- **One device for 40Gb/s IB, FCoIB, 10GigE, CEE, FCoE**
  - One SW stack for offload, virtualization, RDMA, storage

## 10/20/40 Gb/s IB Adapter

10/20/40 Gb/s  
PCIe Gen 1/2  
Dual Port  
Hardware IOV



## 10GigE Adapter

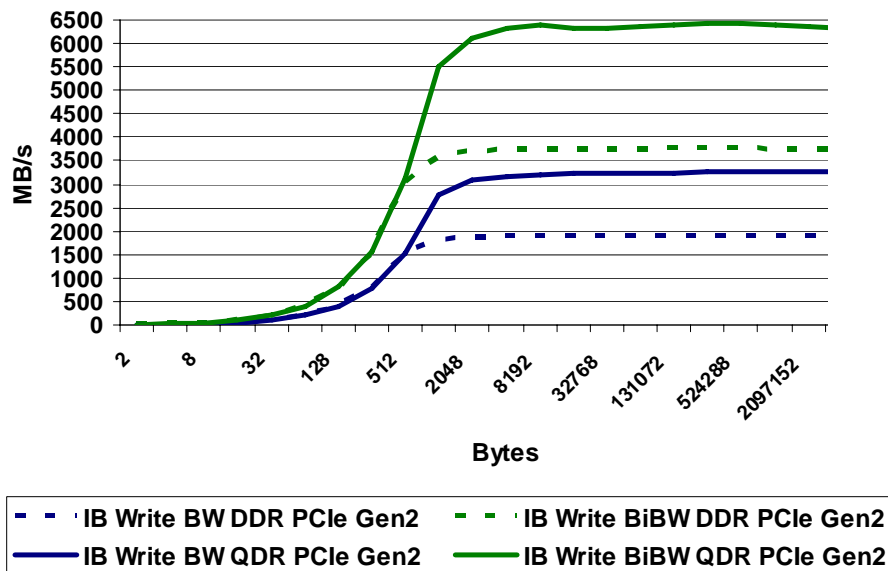
Dual Port  
PCIe Gen1/2  
Hardware IOV  
FCoE, CEE



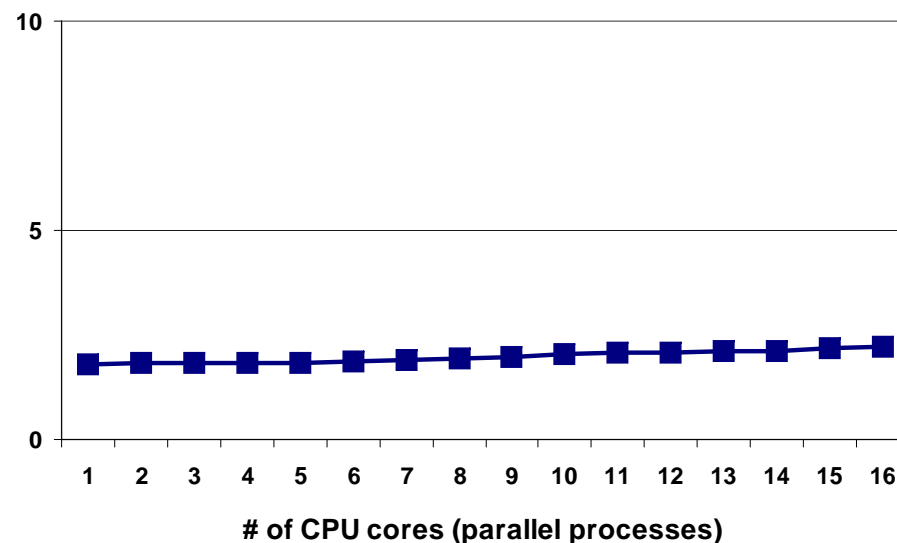
# ConnectX Multi-core MPI Scalability



ConnectX IB QDR versus DDR Bandwidth  
(PCIe Gen2)



Mellanox ConnectX  
MPI Latency - Multi-core Scaling



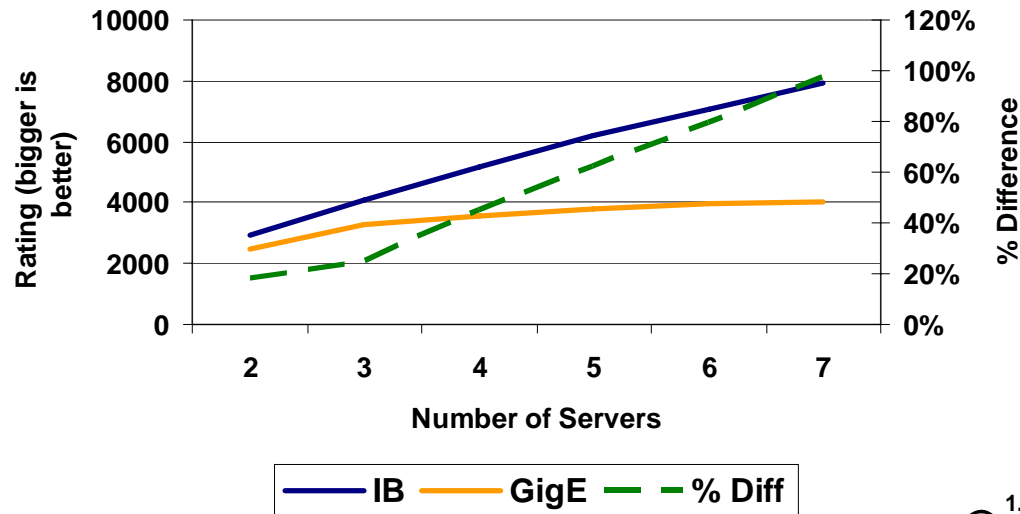
- Scalability to 64+ cores per node, to 20K+ nodes per cluster
- Guarantees same low latency regardless of the number of cores
- Guarantees linear scalability for real applications



# ConnectX InfiniBand FLUENT Benchmarks



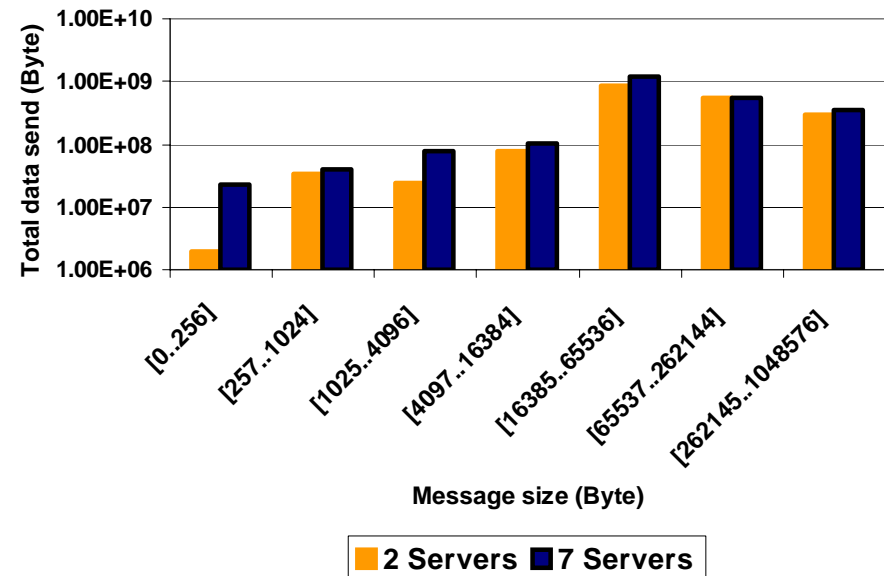
Fluent Turbo\_500k Benchmark Results



## Scalability Mandates Bandwidth

4-node InfiniBand cluster demonstrates higher performance versus any cluster size with GigE


Fluent Message Size Profiling



# InfiniBand Address the Needs for Petascale Computing



- **Balanced random network streaming**
  - "One to One" random streaming
  - Solution: Dynamic routing (InfiniScale IV)
- **Balanced known network streaming**
  - "One to One" known streaming
  - Solution: Static routing (Now)
- **Un-balanced network streaming**
  - "Many to one" streaming
  - Solution: Congestion control (Now)
- **Faster network streaming propagation**
  - Network speed capabilities
  - Solution: InfiniBand QDR (InfiniScale IV)
    - 40/80/120Gb/s



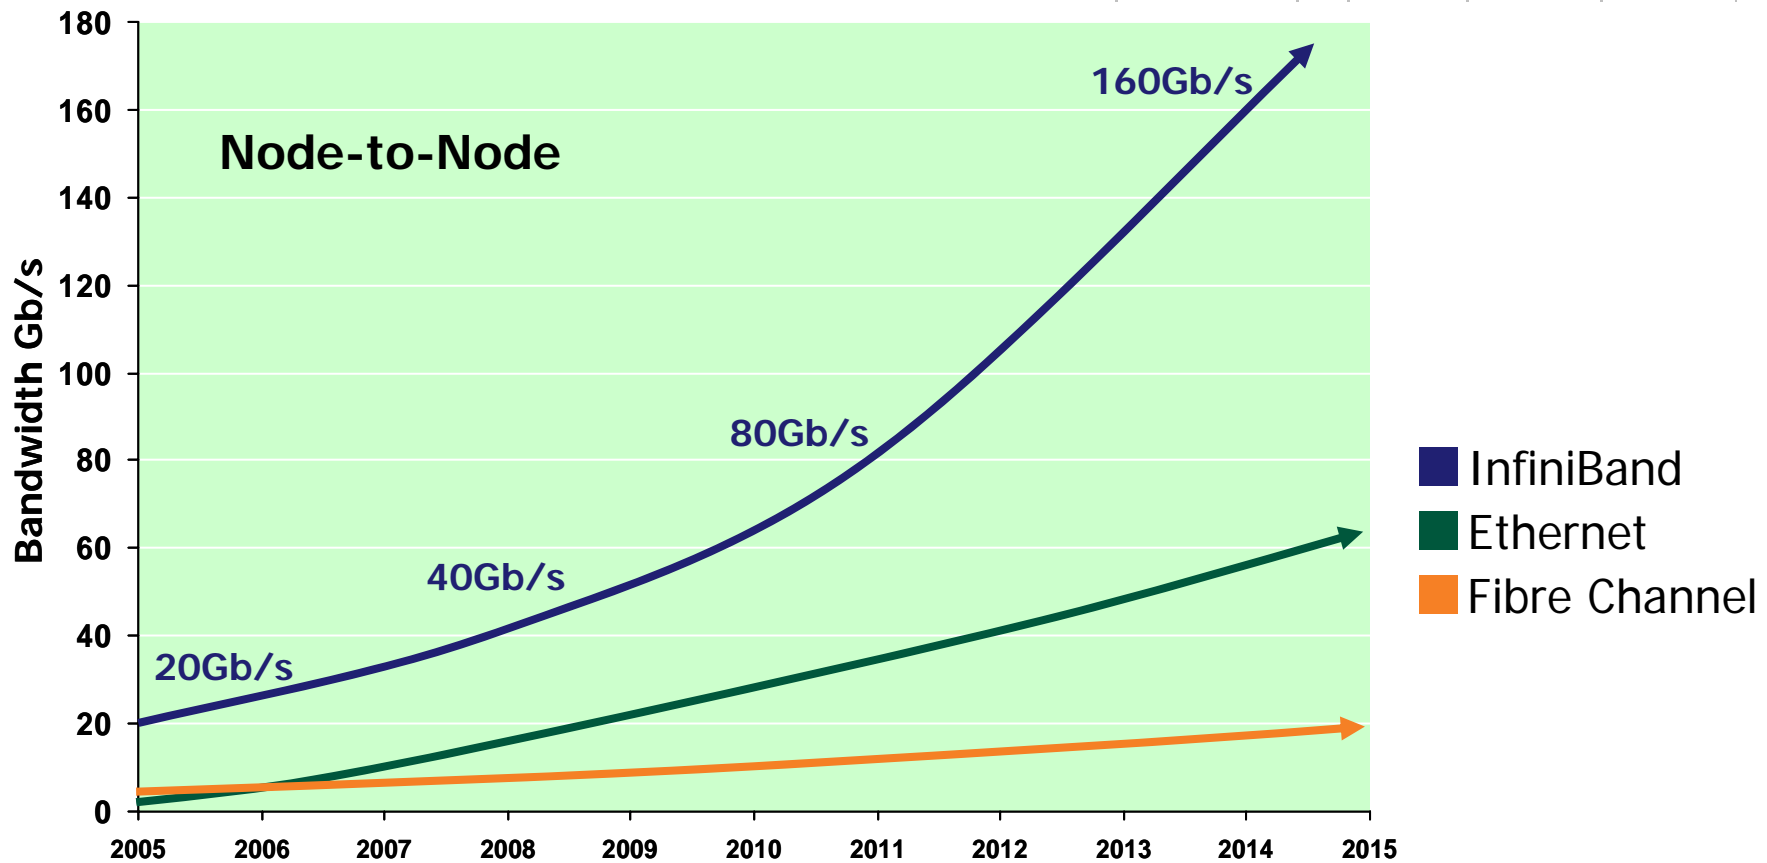
**IB designed to handle all communications in HW**

# InfiniBand Roadmap



InfiniBand roadmap outpaces all other I/O technologies in both pure performance and price/performance

| Number of IB Lanes | Per Lane Bandwidth |         |         |         |
|--------------------|--------------------|---------|---------|---------|
|                    | 5Gb/s              | 10Gb/s  | 20Gb/s  | 40Gb/s  |
| 4X                 | 20Gb/s             | 40Gb/s  | 80Gb/s  | 160Gb/s |
| 8X                 | 40Gb/s             | 80Gb/s  | 160Gb/s | 320Gb/s |
| 12X                | 60Gb/s             | 120Gb/s | 240Gb/s | 480Gb/s |



Thank You

Gilad Shainer

[shainer@mellanox.com](mailto:shainer@mellanox.com)

