# Mellanox Interconnect Solutions

Michael Kagan
Chief Technology Officer

# Company Overview

Ticker: MLNX

- Leading provider of high-throughput, low-latency server and storage interconnect
  - FDR 56Gb/s InfiniBand and 10/40GbE
  - Reduces application wait-time for data
  - Dramatically increases ROI on data center infrastructure
- Company headquarters:
  - Yokneam, Israel; Sunnyvale, California
  - ~950 employees; worldwide
- Solid financial position
  - Record revenue in FY'11; $259.3M
  - Q1'12 revenue = $88.7M; up 61.2% Y-o-Y
  - Q2'12 guidance ~$125.0M to $130.0M
  - Cash + investments @ 3/31/12 = $276.6M

## Fortune 100 Penetration

**5 of top 10** Global Banks

**10 of top 10** Automotive Manufacturers

**5 of top 10** Pharmaceutical Companies

**9 of top 10** Oil and Gas Companies

# The Vision – Everyone's Job

**Forecast – tell the future**

**Vision – tell the future**

***and make it happen!***

# 2003: tell the Future – Performance Leadership

## Mellanox HCA Roadmap – September 2003

### (1st silicon schedule)

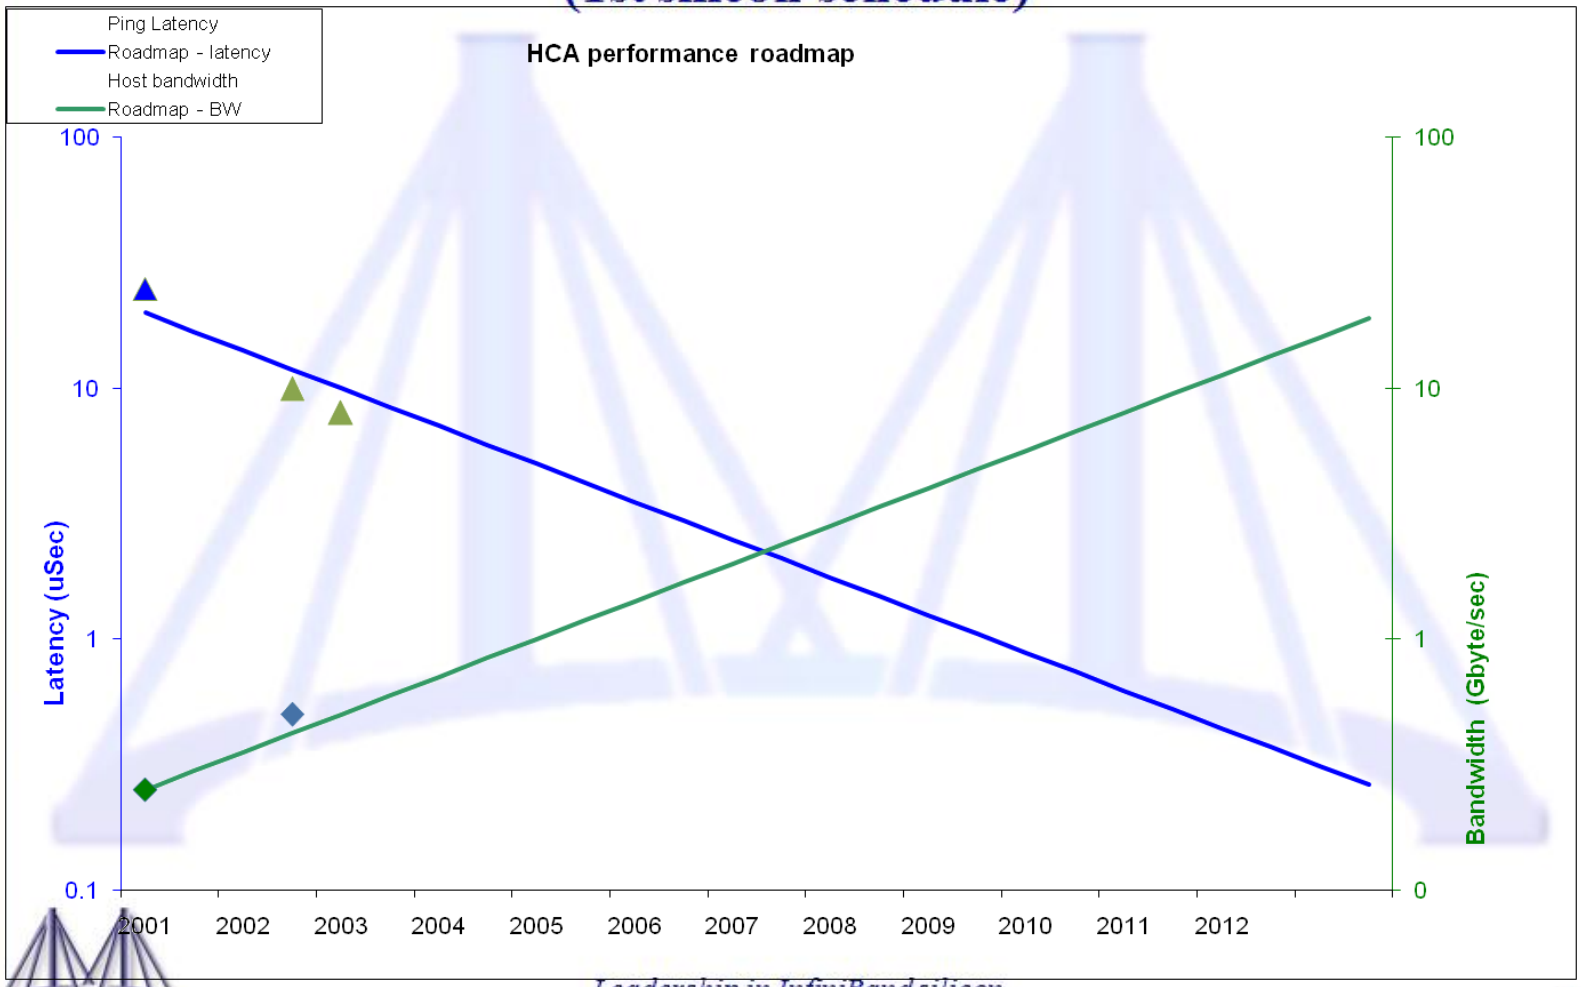


*Leadership in InfiniBand silicon*

*Mellanox confidential*

3

# 2012: …and Make it Happen

## Mellanox HCA performance track record (March 2011)

### (1st silicon schedule)

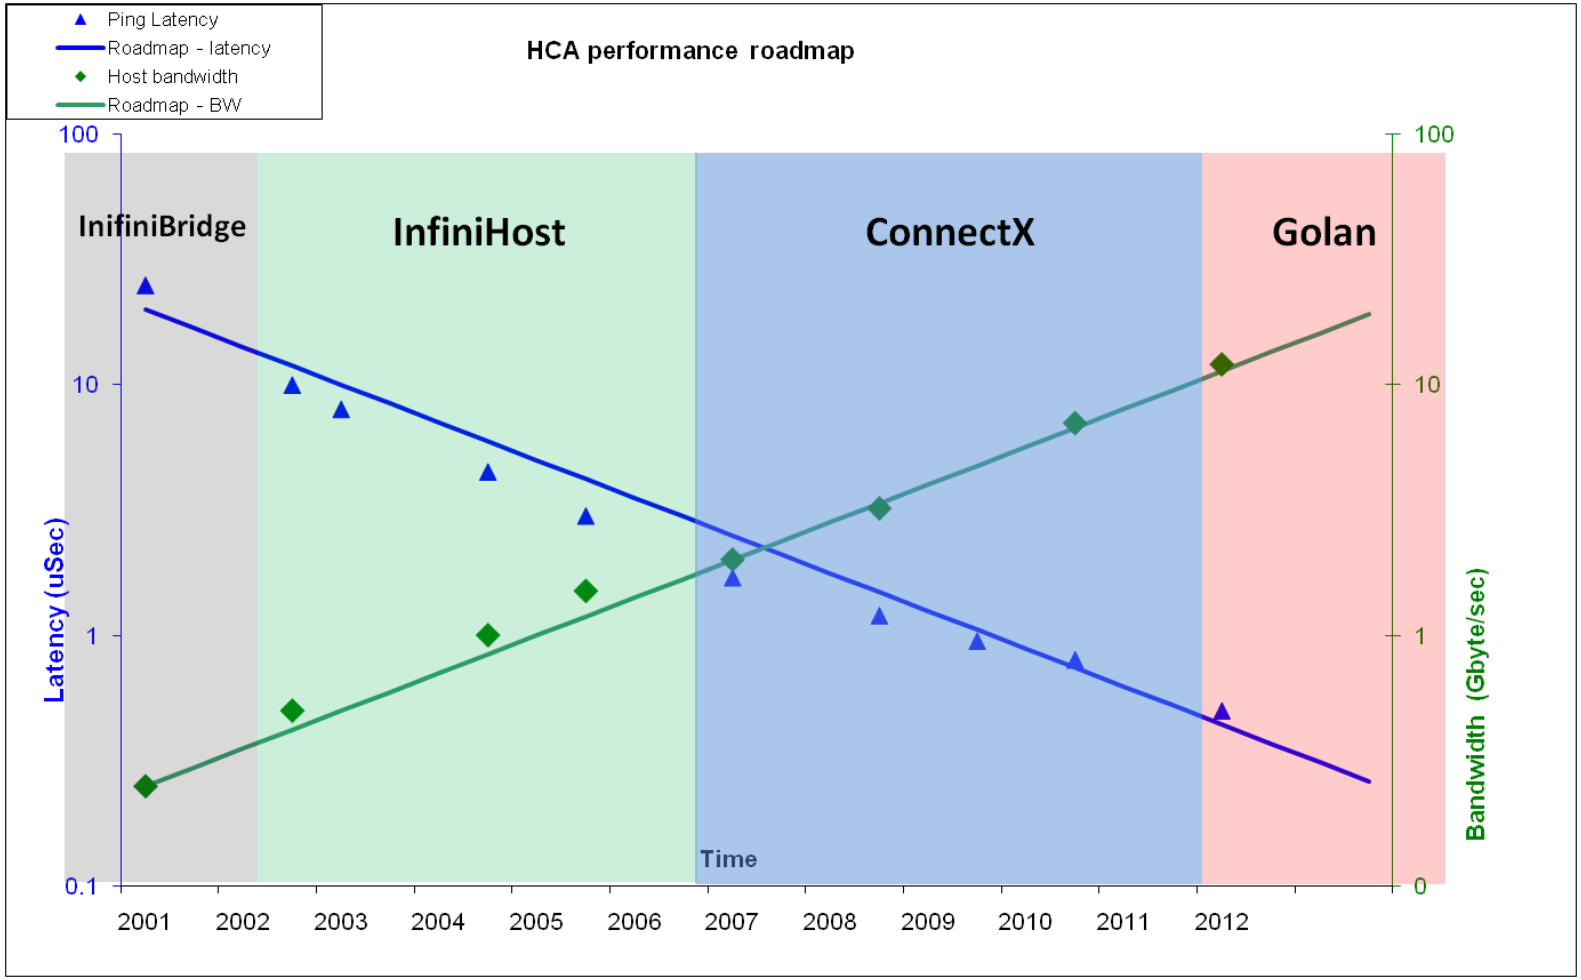

# Rule #1 – Deliver

***Execution is God***

*Andy Grove, 1992*

# Execution is God – Deliver Products

- InfiniBridge – IB TCA + switch
  - Tape out – Q4/00
  - Samples – Q1/01
  - Production – Q2/01
- InfiniScale – 8-port SDR IB switch
  - Tape out – Q3/01
  - Samples – Q4/01
  - Production – Q1/01
- InfiniHost – PCIX HCA
  - Tape out – Q2/03
  - Samples – Q3/03
  - Production – Q4/03
- InfiniScale-3 – 24-port DDR IB switch
  - Tape out – Q3/03
  - Samples – Q4/03
  - Production – Q1/04
- InfiniHost-III EX – PCIe HCA, 2xDDR
  - Tape Out – Q4/03
  - Samples – Q1/04
  - Production – Q2/04
- InfiniHost-III LX – PCIe HCA,1xDDR
  - Tape Out – Q4/04
  - Samples – Q1/05
  - Production – Q2/05
- ConnectX – VPI HCA, 2xQDR/10GE
  - Tape Out – Q4/06
  - Samples – Q1/07
  - Production – Q2/07
- InfiniScale-IV – 36-port QDR switch
  - Tape Out – Q1/08
  - Samples – Q2/08
  - Production – Q3/08
- BridgeX – IB/FC/Eth bridge
  - Tape Out – Q4/08
  - Samples – Q1/09
  - Production – Q3/09
- ConnectX2 – VPI HCA, 40GE Eth
  - Tape Out – Q1/09
  - Samples – Q2/09
  - Production – Q3/09
- SwitchX – VPI FDR/40/10/1GE/FC switch
  - Tape Out – Q4/10
  - Samples – Q1/11
  - Production – Q2/11
- ConnectX3 – VPI HCA, 40GE Eth, 56G IB
  - Tape Out – Q1/11
  - Samples – Q2/11
  - Production – Q3/11
- ConnectIB – 100Gbit HCA
  - TapeOut - Q1/12
  - Samples – Q2/12
  - Production – Q3/12 (E)

**13/13 A-step silicon in production**

# Rule #2 – Serve your Customer

***Always do your best supporting your customer***

# Rule #2 – Make Your Customer Successful

***It is not always enough that we do our best; sometimes we have to do what is required***

*Winston Churchill*

# Rule #3 – Maintain Leardership

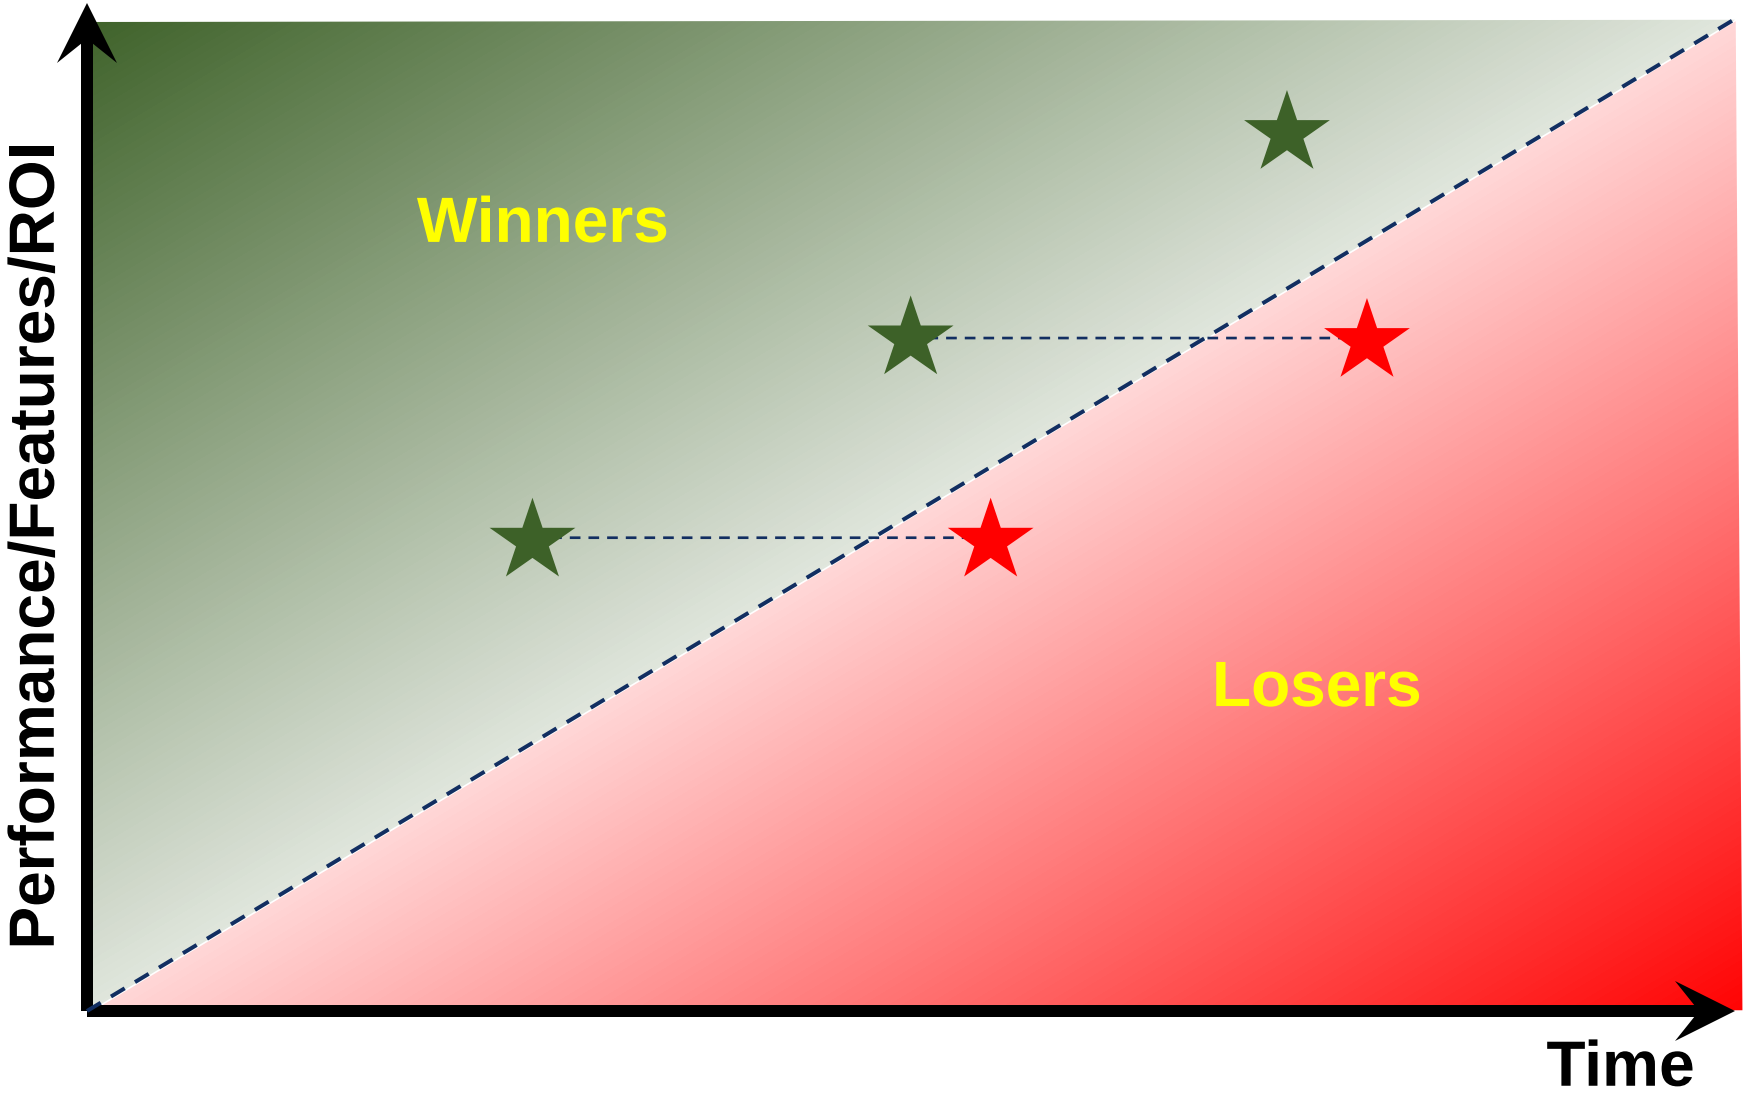

# Leading Supplier of End-to-End Interconnect Solutions


## Comprehensive End-to-End InfiniBand and Ethernet Portfolio

### ICs

### Adapter Cards

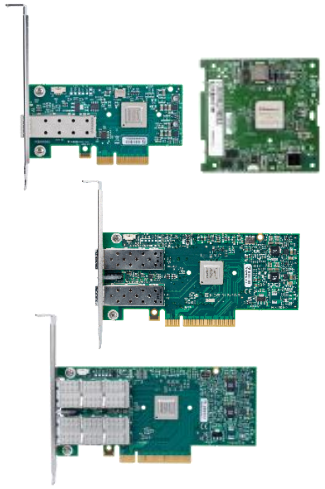

### Switches/Gateways

### Host/Fabric Software


### Cables

# Top Tier OEMs, ISVs and Distribution Channels


## Hardware OEMs

### Servers

HP, IBM, ORACLE, DELL, ASUS, BULL, FUJITSU, NEC, sgi, Sugon, SUPERMICRO

### Storage

EMC², HP, IBM, NetApp, ORACLE, DataDirect Networks, LSI, sgi, xyratex

### Embedded

GE FANUC, DALSA, FUJITSU Medical, Computer Systems MERCURY, SeaChange

## Software Partners

Microsoft, redhat, Novell, SOLARIS, AIX L, HP-UX 11i, vmware, ORACLE, IBM DB2, REUTERS, SYNOPSYS, Autodesk

## Selected Channel Partners

ALTIMA, ARROW ENTERPRISE COMPUTING SOLUTIONS, atlantik systeme, AVNET, DSCon, hammer, Netweb TECHNOLOGIS, SAMBO SYSTEM, SERVANTS INTERNATIONAL, SysOne, Tech Data

# The Debut – Top of the Top500, 2003

| | Machine | CPUs | Gflops |
|---|---|---|---|
| 1. | Earth Simulator | 5120 | 35,860 |
| 2. | ASCI (Quadrics) | 8160 | 13,880 |
| **3.** | **Apple (InfiniBand)** | **2200** | **10,280** |
| 4. | Xeon (Myrinet) | 2500 | 9,819 |

**Performance of Various Computers Using Standard Linear Equations Software**

Jack J. Dongarra*

Computer Science Department
University of Tennessee
Knoxville, TN 37996-1301

and

Computer Science and Mathematics Division
Oak Ridge National Laboratory
Oak Ridge, TN 37831


If sleep-deprived pizza-fed students can do it, so can YOU!

# Top500 Interconnects

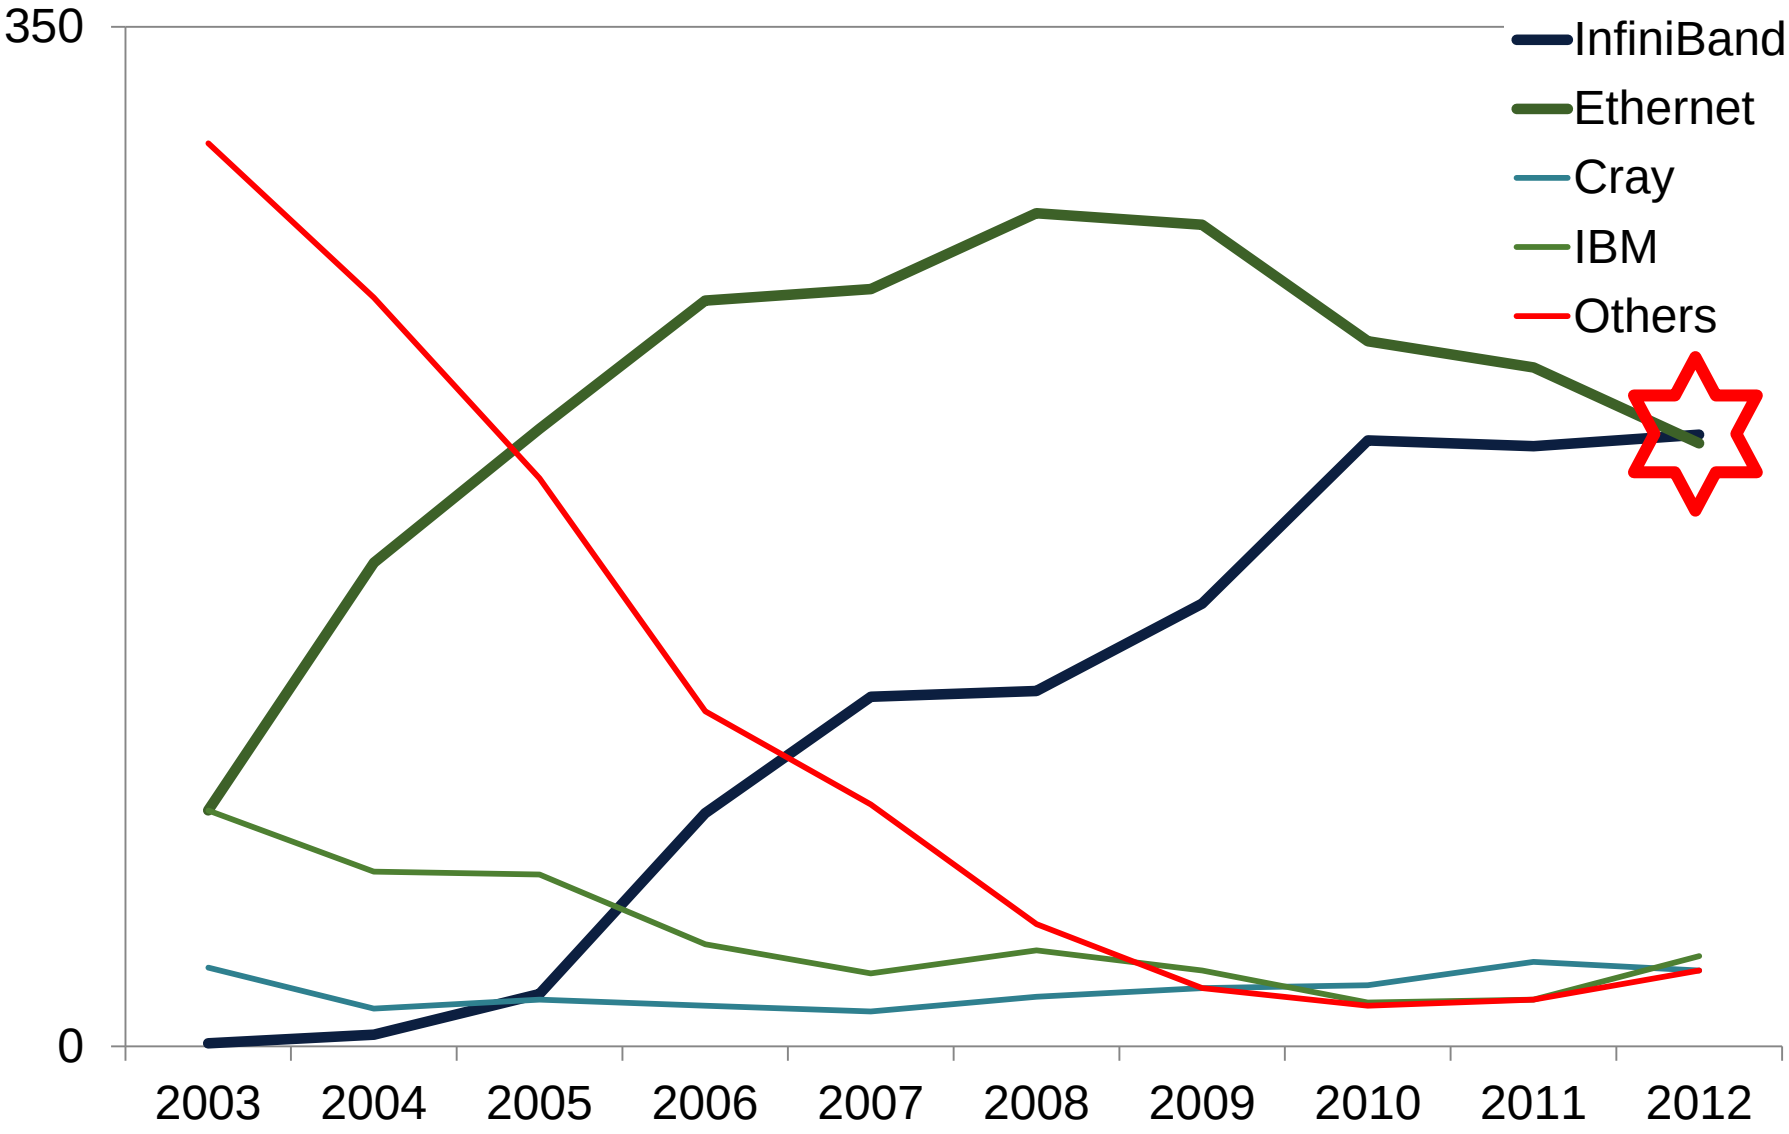

# TOP500 Performance Trend

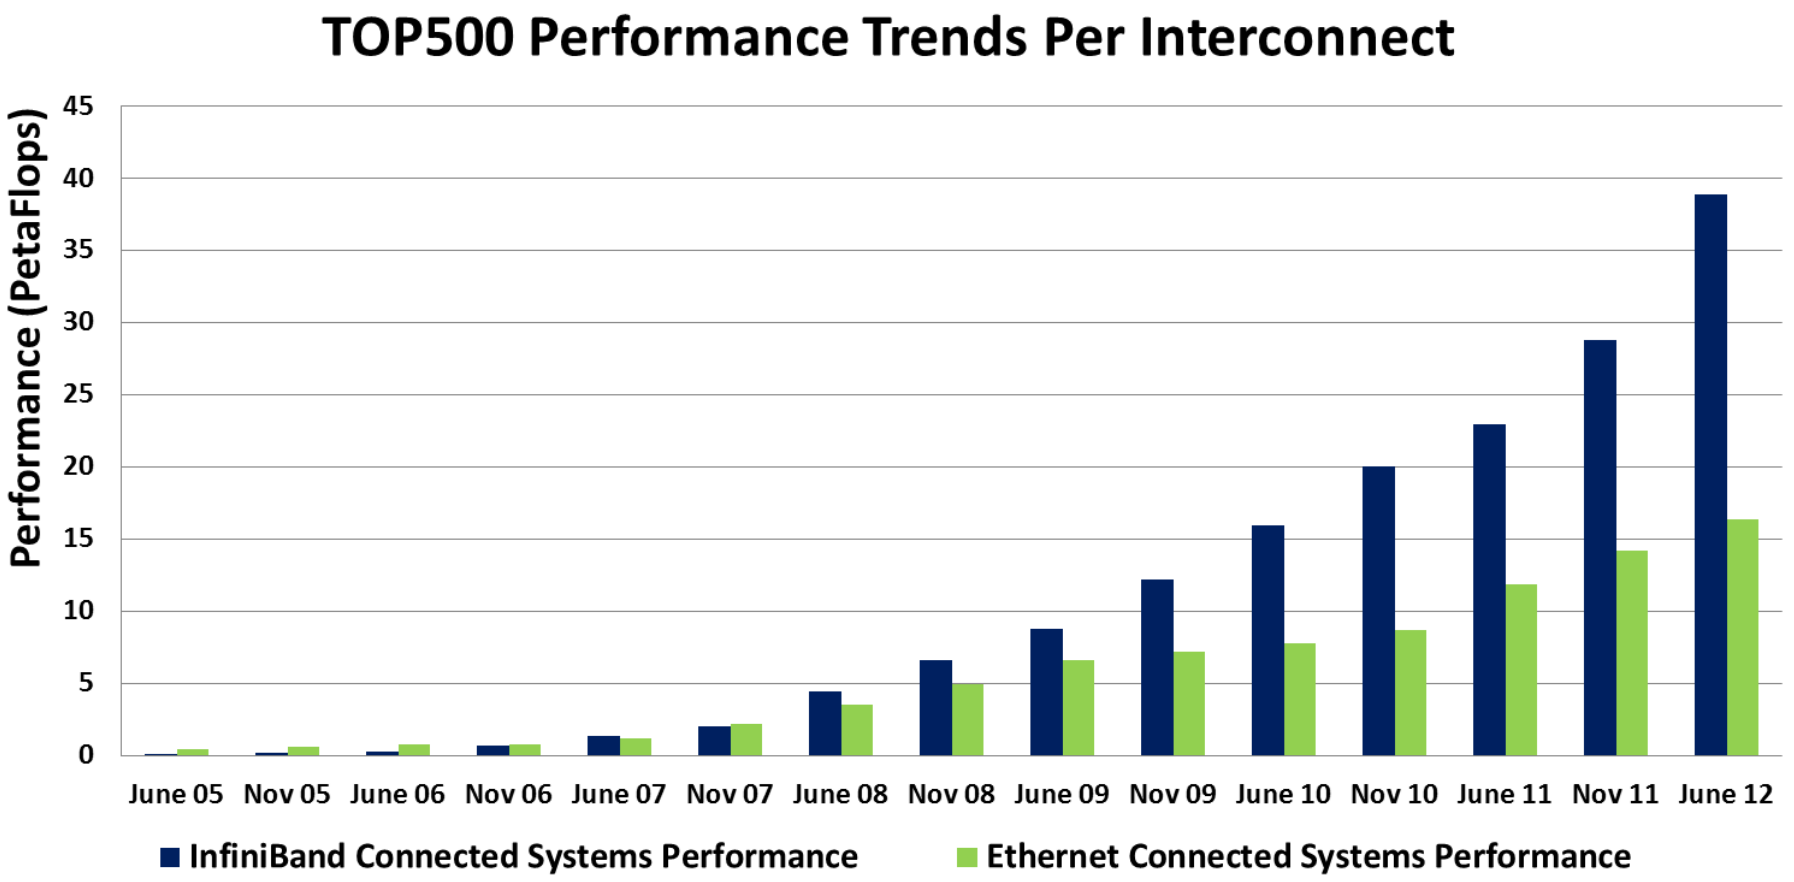


**InfiniBand-based clusters outpaced Ethernet in 3 years**

# InfiniBand's Unsurpassed System Efficiency

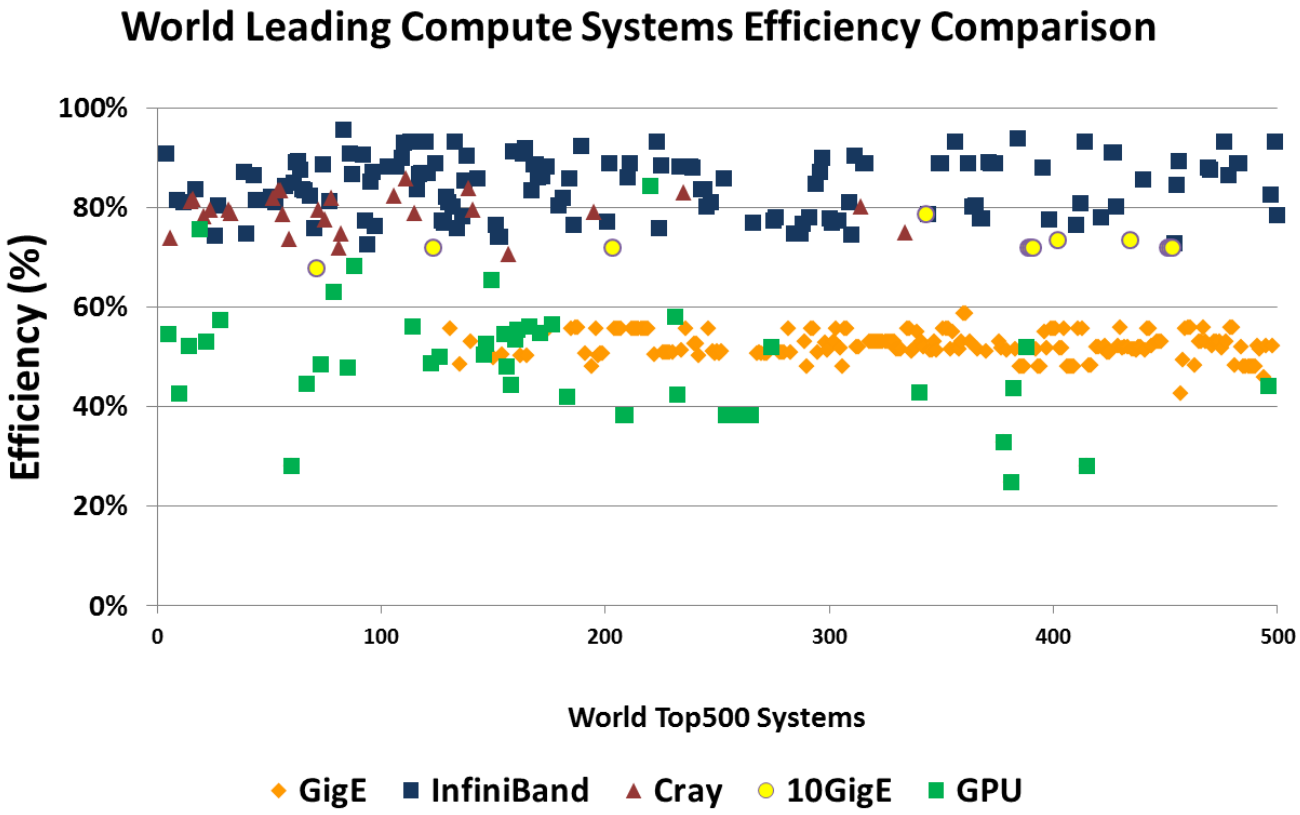


**With InfiniBand the cluster is 1.5X more efficient**

# InfiniBand Roadmap

**2002**
**10Gb/s**

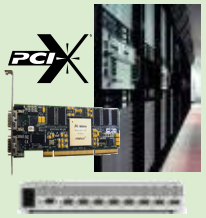


**2005**
**20Gb/s**

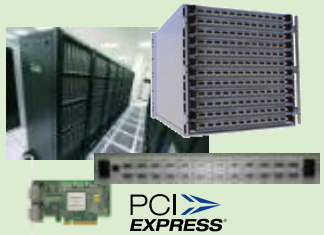


**2008**
**40Gb/s**


**2011**
**56Gb/s**


## Highest Performance, Reliability, Scalability, Efficiency


# FDR InfiniBand New Features and Capabilities

**Performance / Scalability**

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

**Reliability / Efficiency**

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

# The Building Blocks

## ConnectX®-3 VPI Adapter

redhat. Novell SUSE Microsoft vmware OPENFABRICS SOFTWARE

Applications

Networking | Storage | Clustering | Management

Acceleration Engines

PCI EXPRESS™ 3.0

Ethernet: 10/40 Gb/s

InfiniBand:10/20/40/56 Gb/s

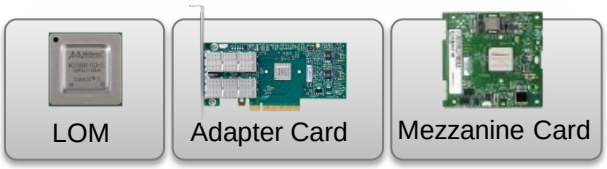


## SwitchX™ VPI Switch

Unified Fabric Manager

Switch OS Layer


64 ports 10GbE
36 ports 40GbE
48 10GbE + 12 40GbE
36 ports IB up to 56Gb/s
8 VPI subnets

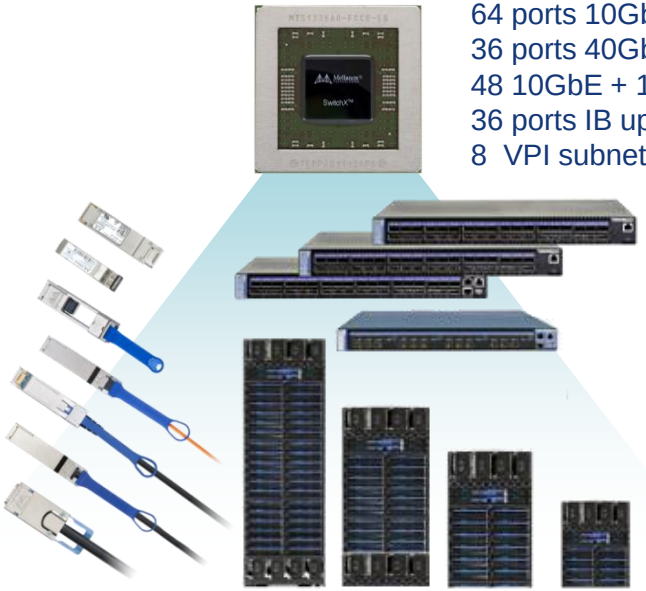

# FDR InfiniBand PCIe 3.0 vs QDR InfiniBand PCIe 2.0

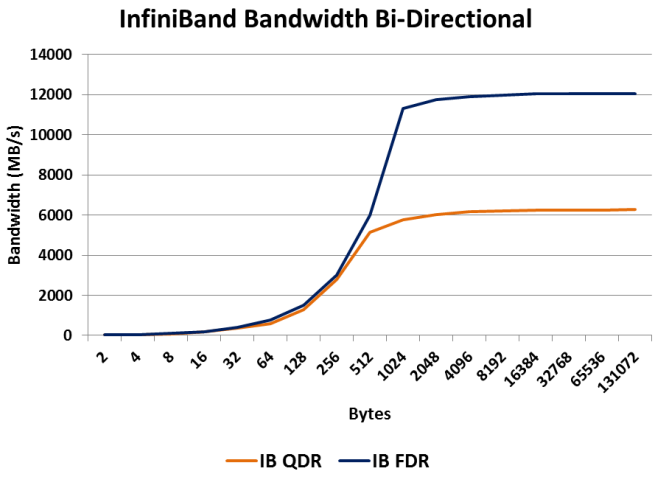


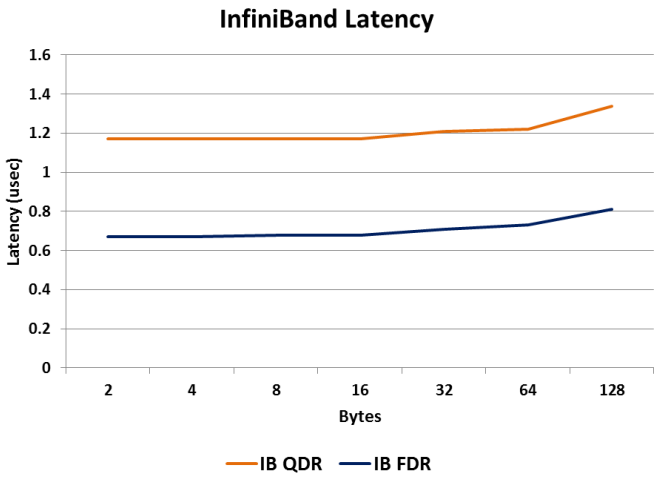


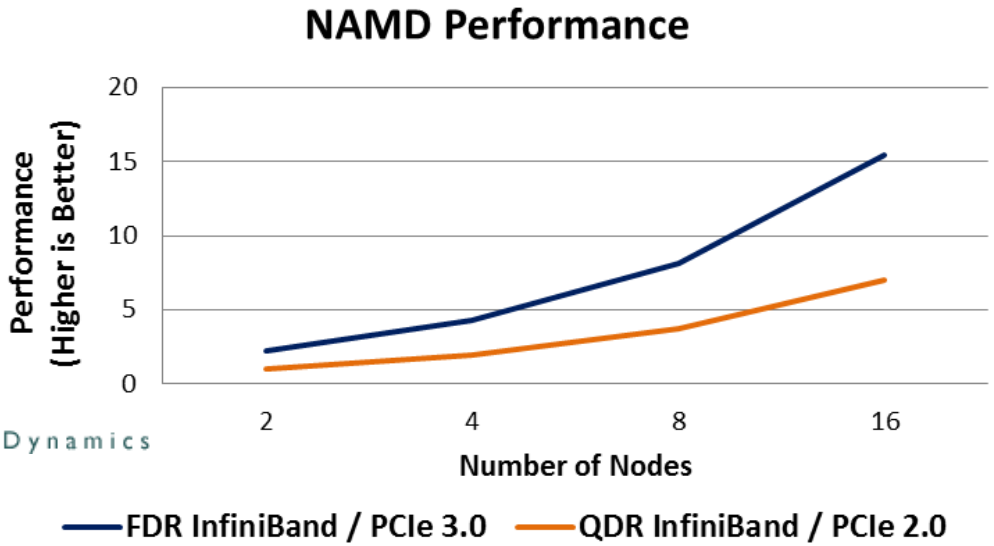


NAMD
Scalable Molecular Dynamics

**Double the Bandwidth, Half the Latency**

**Higher Application ROI**

# FDR InfiniBand Application Example: LS-DYNA

- LS-DYNA Automotive Crash Simulator
  - General purpose structural and fluid analysis simulation software package
  - Used for automobile, aerospace and manufacturing simulations
- FDR 56Gb/s InfiniBand improves runtime by nearly 40%
  - 24 nodes,  384 core SandyBridge based cluster

LSTC
Livermore Software Technology Corp.

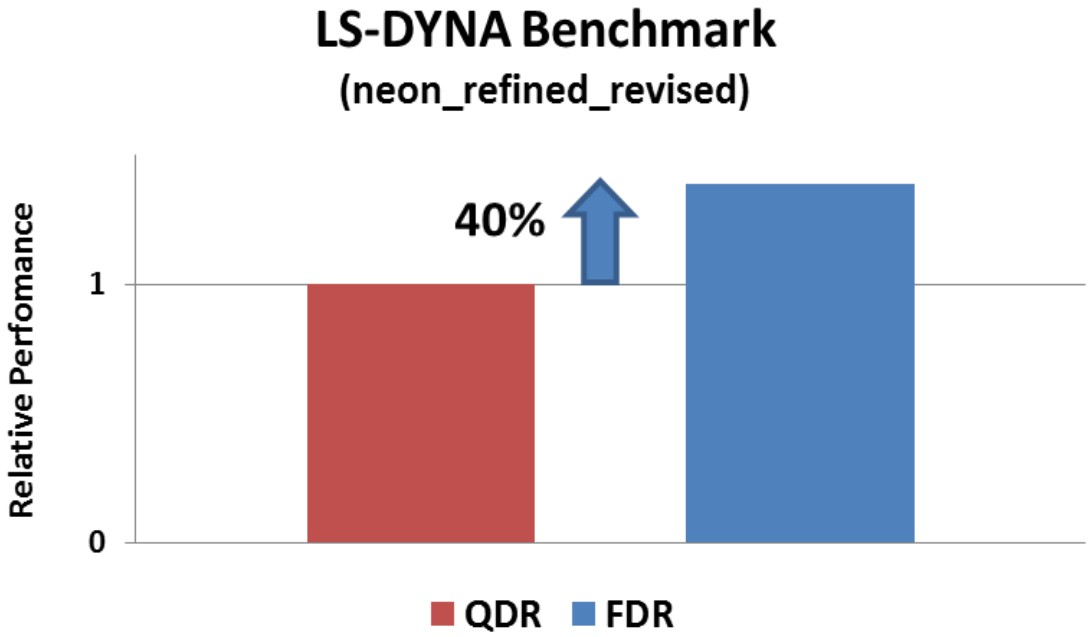

# Co-Design Strategy

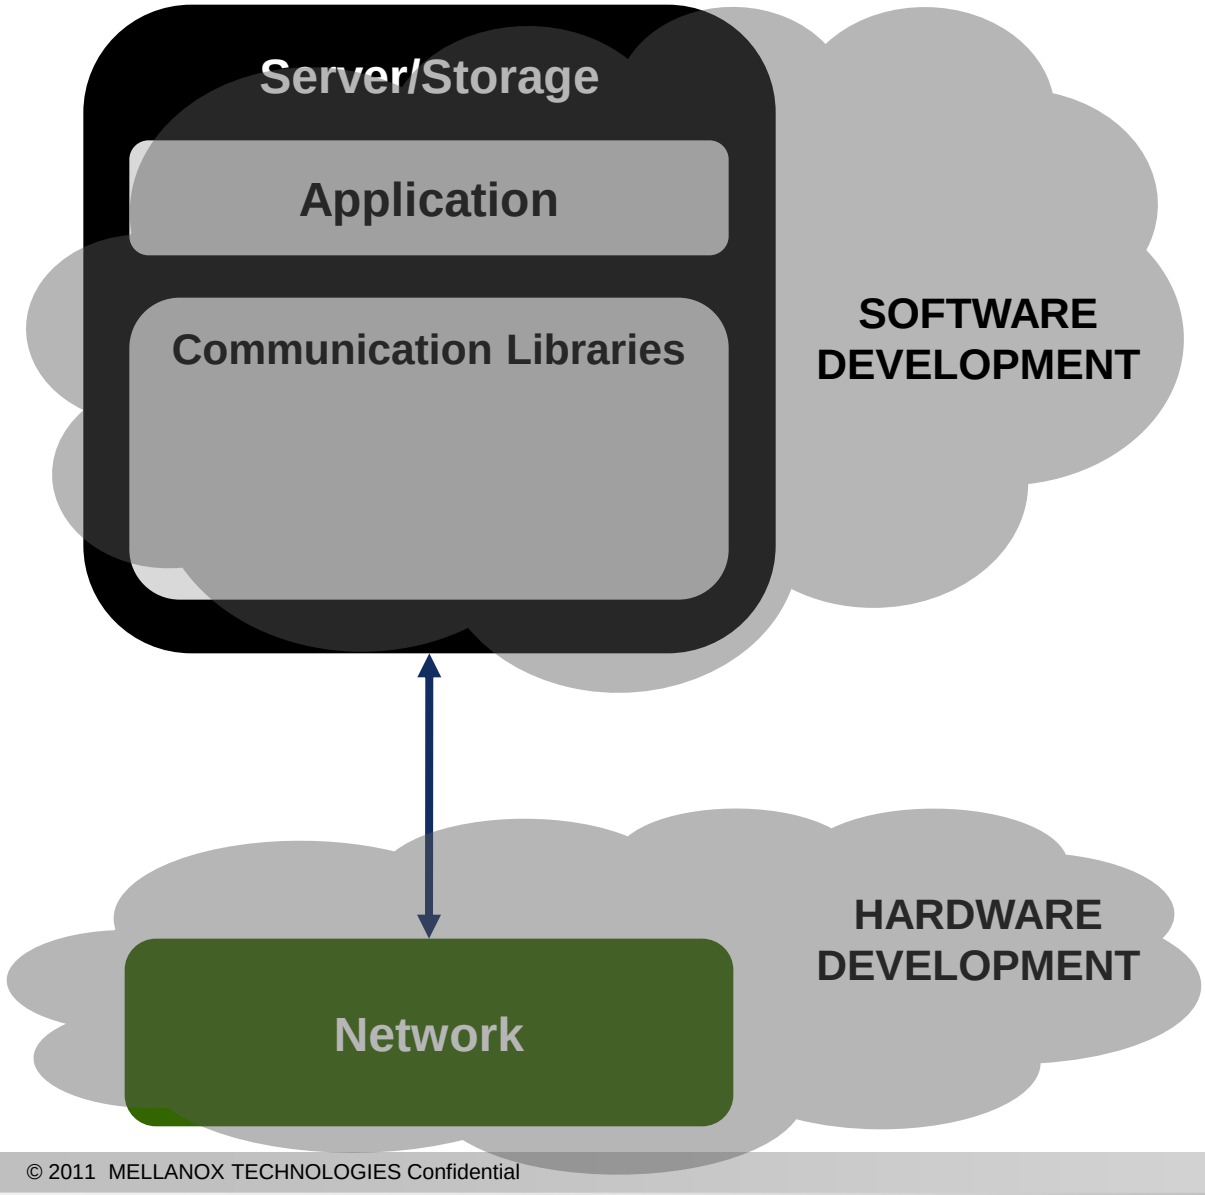

# Co-Design Strategy

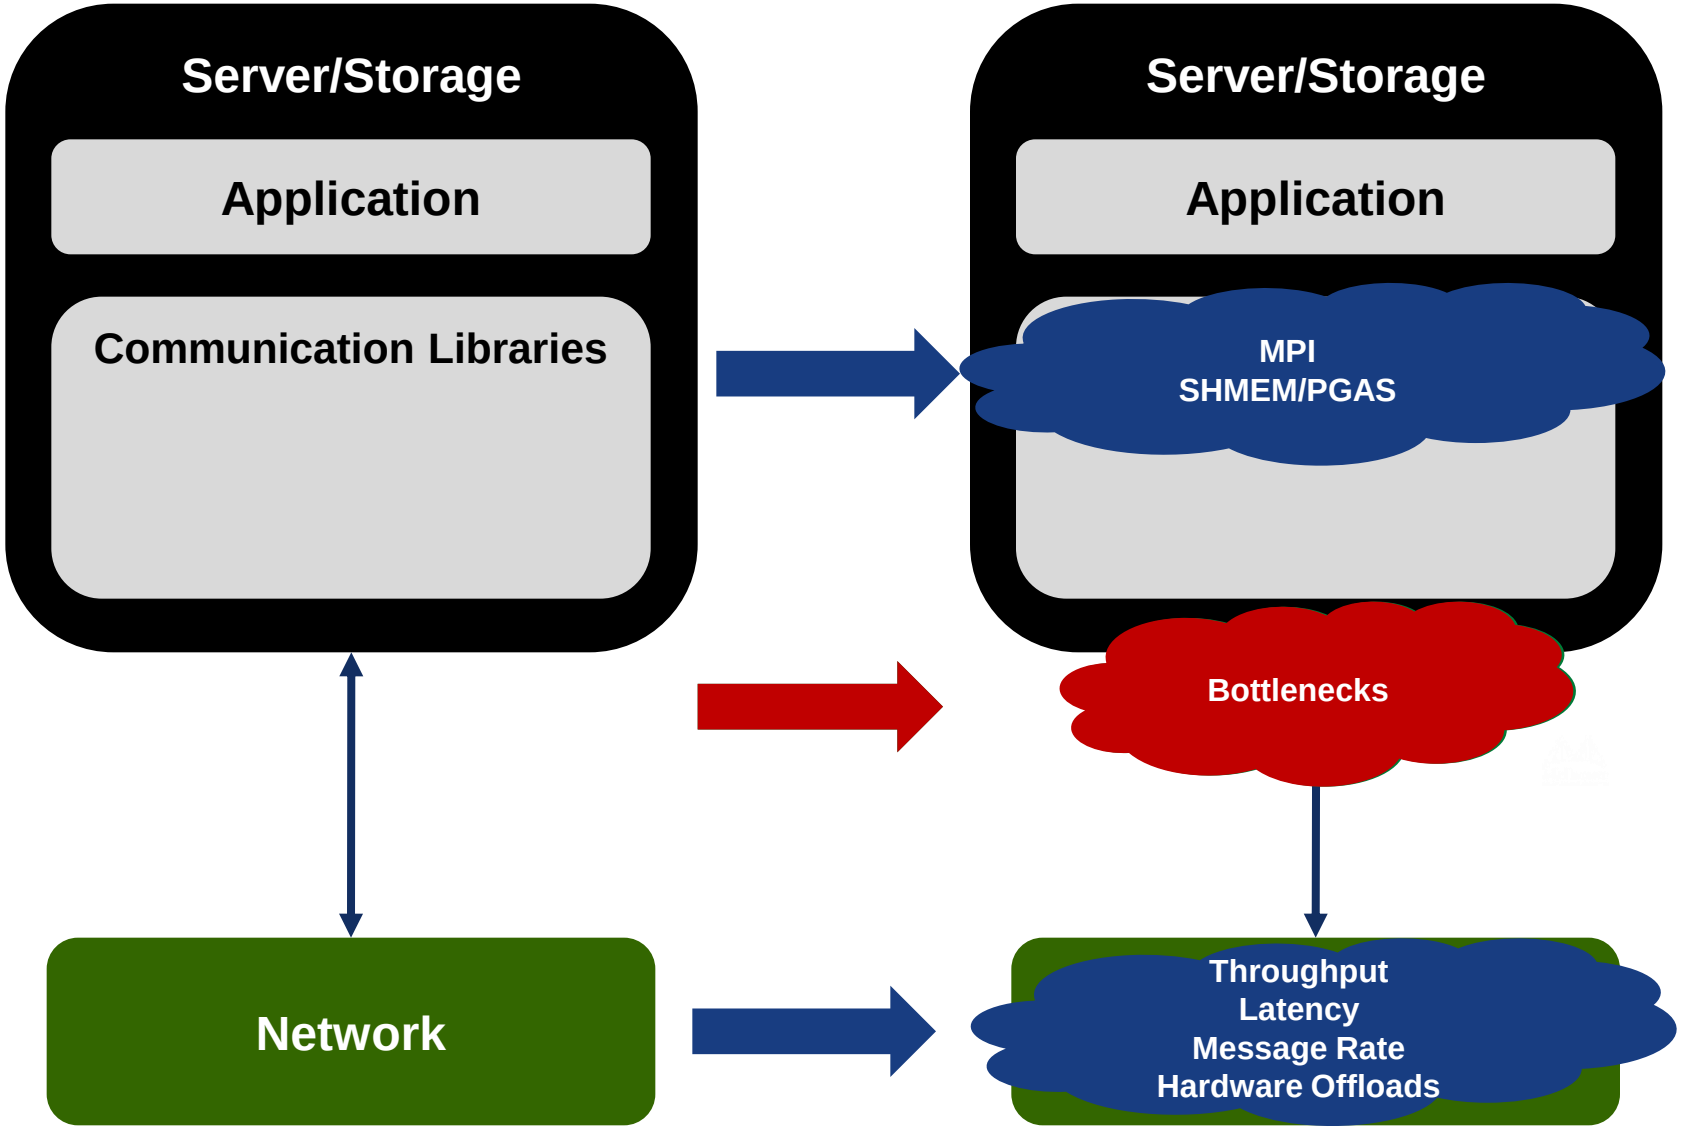

# Co-Design Strategy

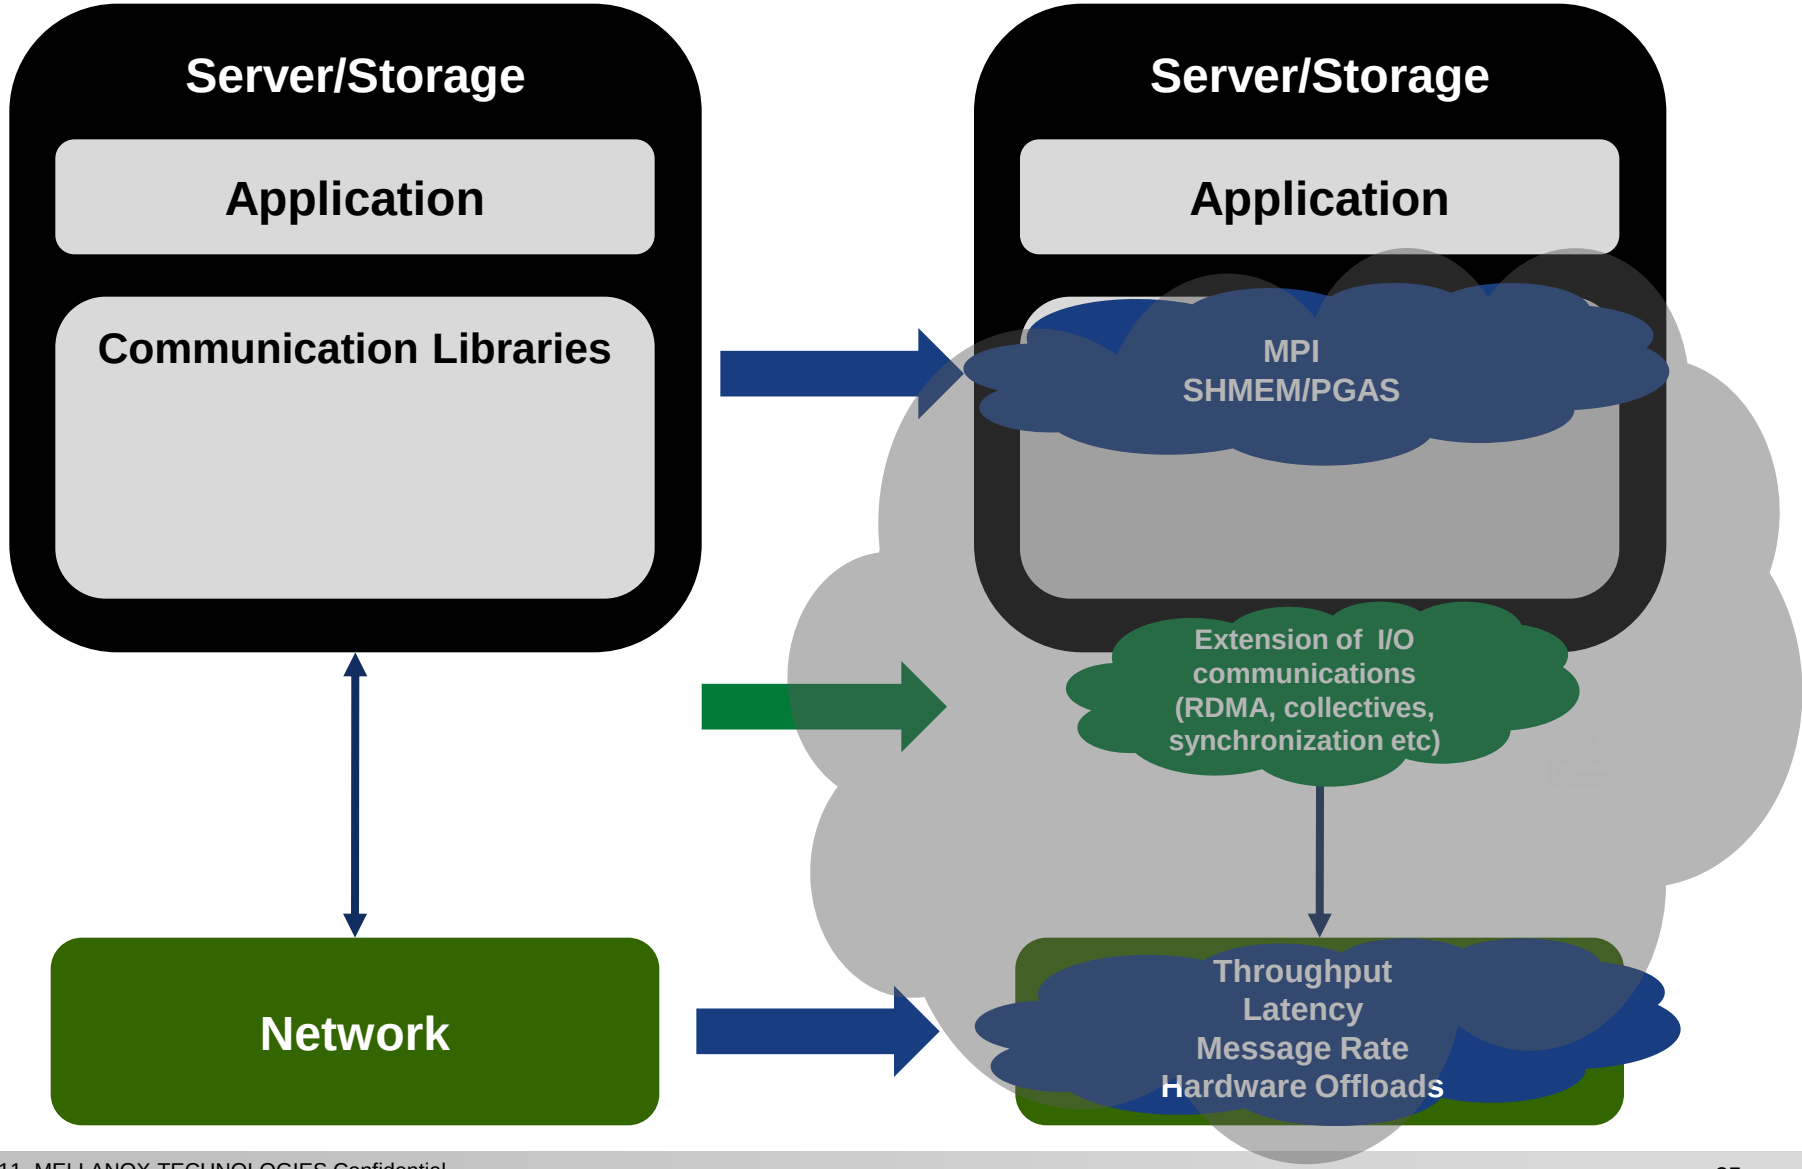

# Mellanox ScalableHPC Accelerate Parallel Applications

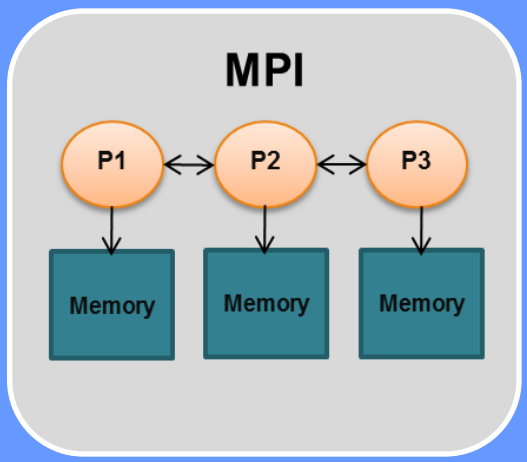


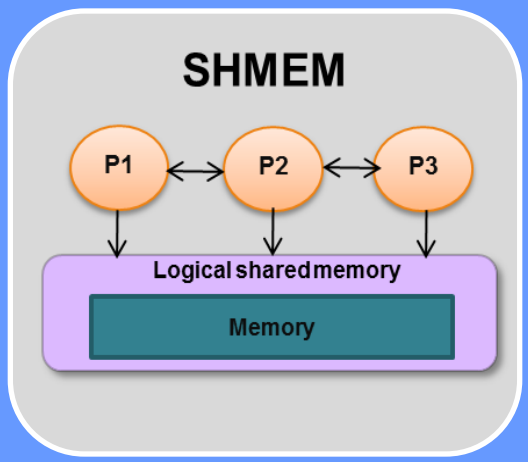


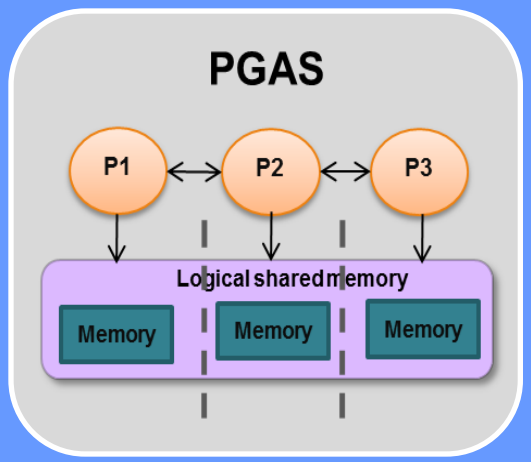


## MXM

- 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
- CoreDIrect Hardware Offload

## InfiniBand Verbs API

# MXM – Scalability

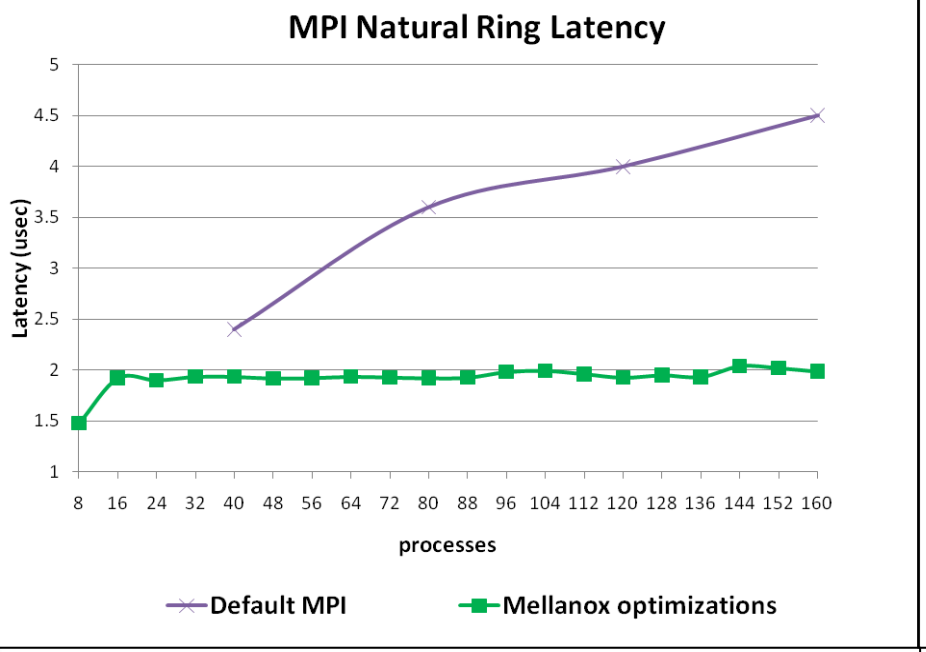


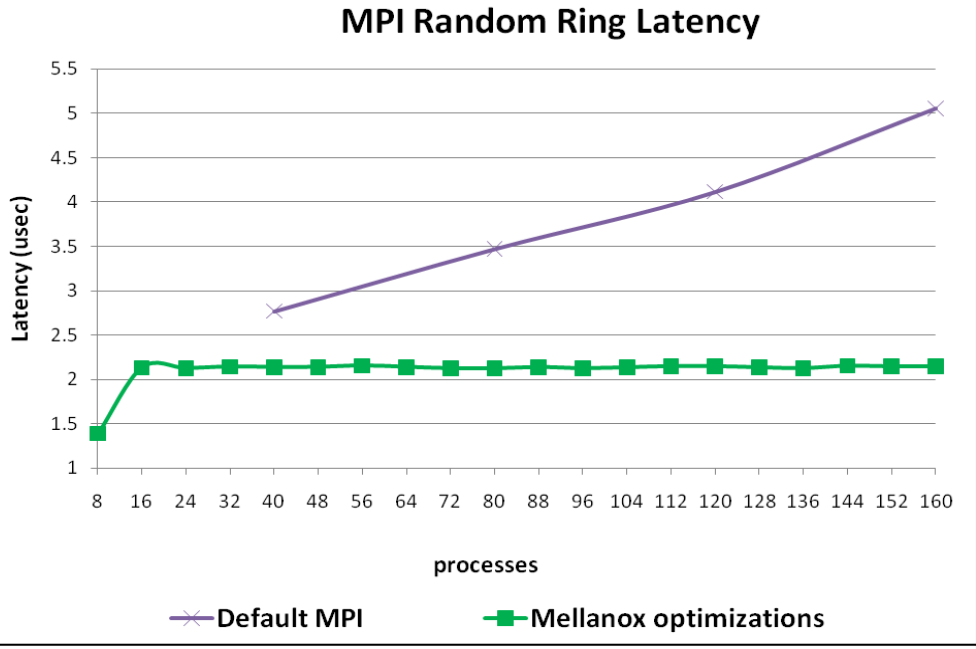


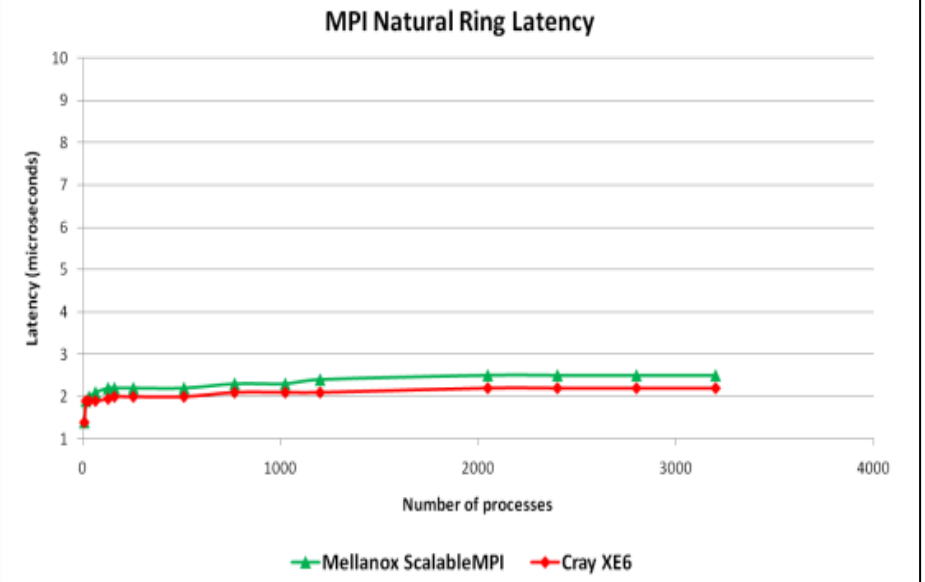


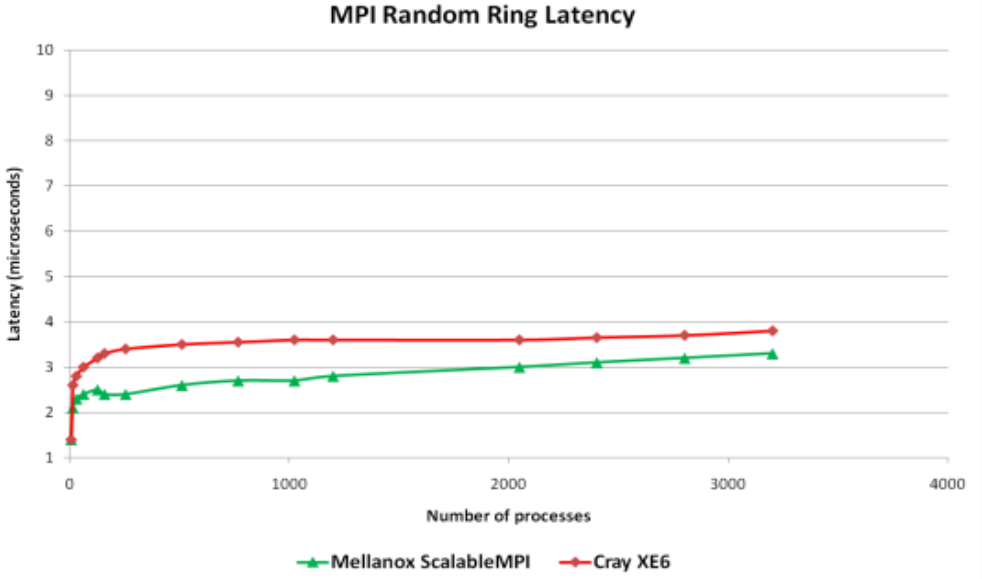

# Mellanox MPI Optimizations – Highest Scalability at LLNL

Mon, Jun 20, 2011

## Mellanox and Lawrence Livermore National Laboratory Demonstrate Leading Performance and Scalability for HPC Applications

*Continuous Collaboration Between Organizations Helps to Maximize HPC Clustering Efficiencies and Boost Application and User Workload Performance*

**ISC'11, Hamburg, Germany – June 20, 2011** – Mellanox® Technologies, Ltd. (NASDAQ: MLNX; TASE: MLNX), a leading supplier of high-performance, end-to-end connectivity solutions for data center servers and storage systems, and Lawrence Livermore National Laboratory (LLNL), today announced world-leading scalability achieved with LLNL supercomputers and Mellanox InfiniBand interconnect solutions. The results are based on collaboration between the two organizations to enhance high-performance computing (HPC) software drivers and MPI libraries on top of Mellanox's scalable interconnect solutions. The joint effort has delivered new levels of workload performance and maximized the return on investment for LLNL users.

"LLNL is tasked to take on some of the world's most difficult and complex problems – all of which are very compute-intensive and place huge performance demands on the network," said Matt Leininger, deputy for advanced technology projects at LLNL. "The collaborative work with Mellanox has helped us to recently advance our scaling capabilities in our large scale Hyperion testbed cluster, delivering significant benefits on tested workloads. Deploying these enhancements on our production cluster will result in faster application runtime which increases research productivity."

The Mellanox and LLNL teams modified Message Passing Interface (MPI) libraries to optimize Mellanox InfiniBand capabilities and enhance application performance and efficiencies. The tested workload performance at scale increased up to 700 percent and resulted in significant progress to meet LLNL's petascale and exascale computing goals.

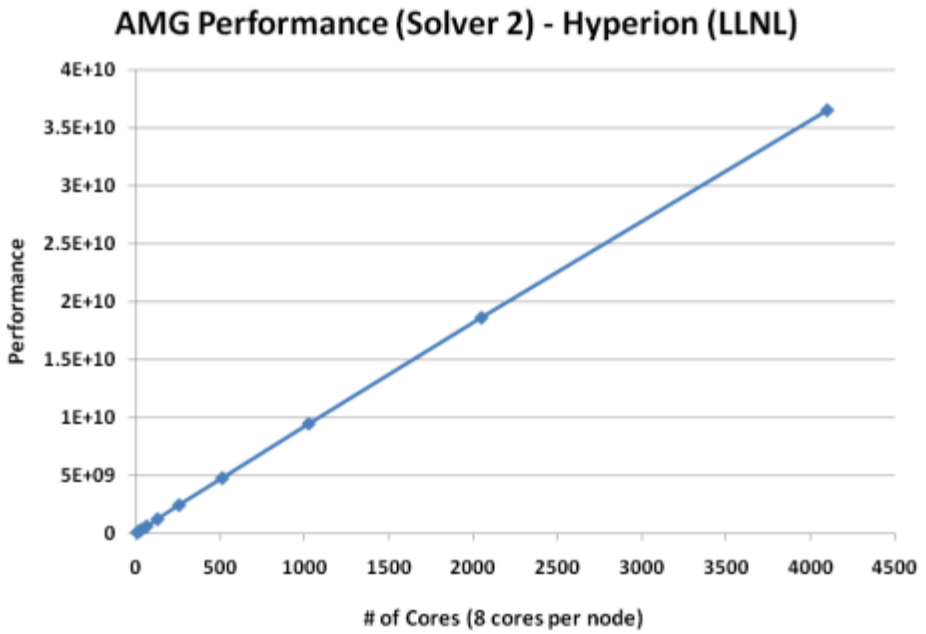


workload performance at scale increased up to 700 percent

Lawrence Livermore National Laboratory

- Mellanox MPI optimization enable linear strong scaling for LLNL application

**World Leading Performance and Scalability**

# FCA – Collective Performance with OpenMPI

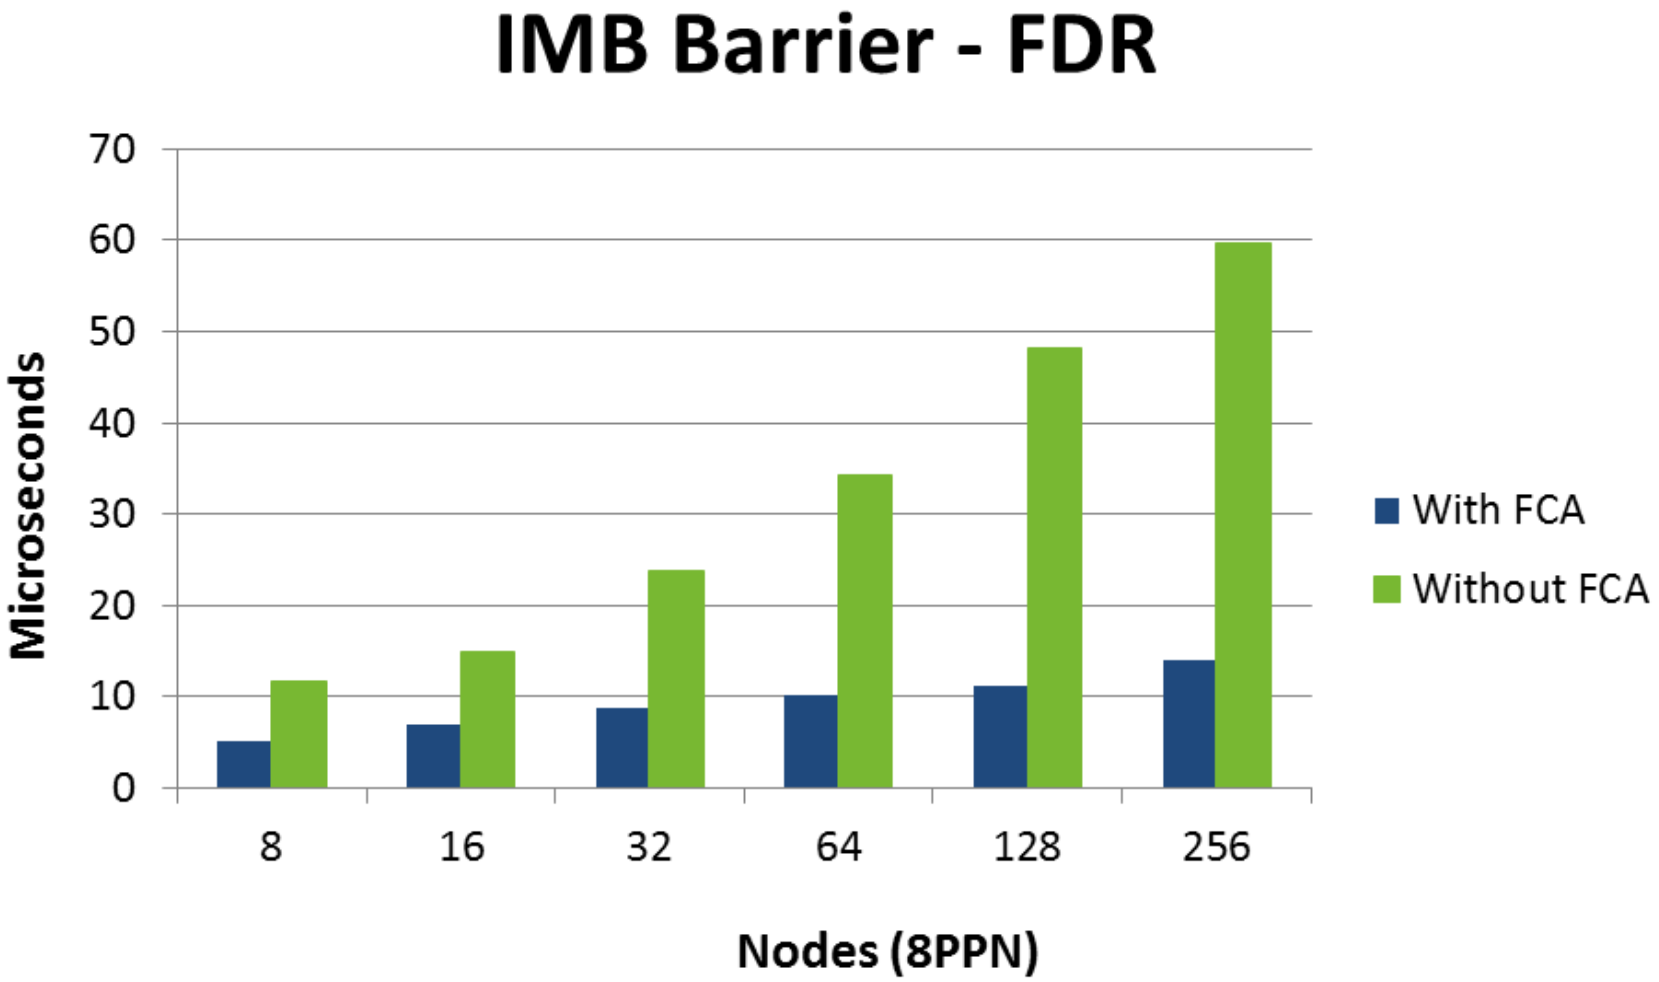

# FCA – Collective Acceleration Results for SHMEM

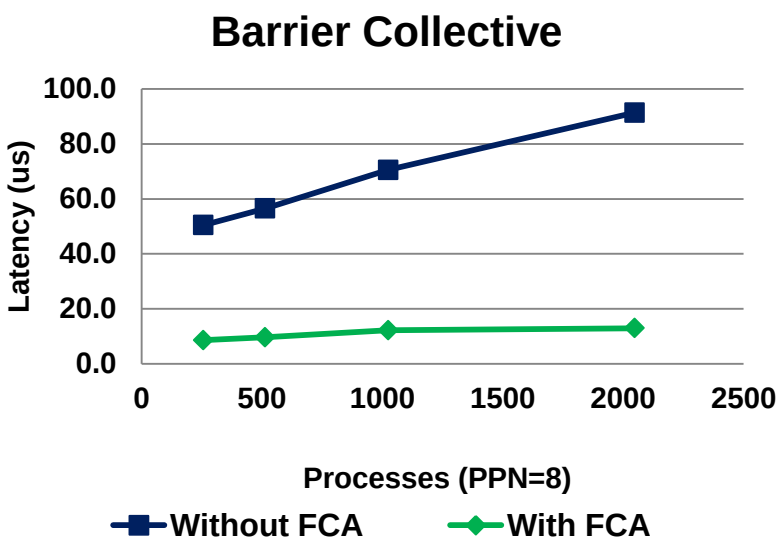


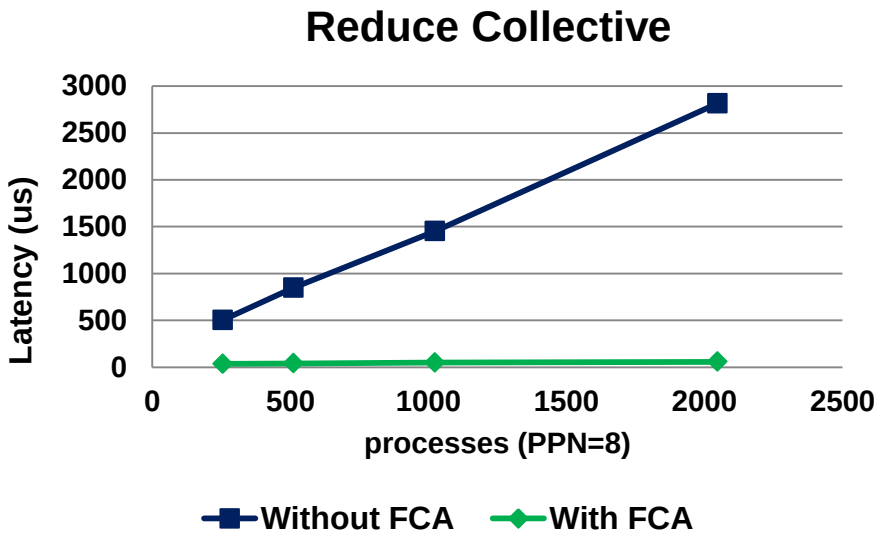


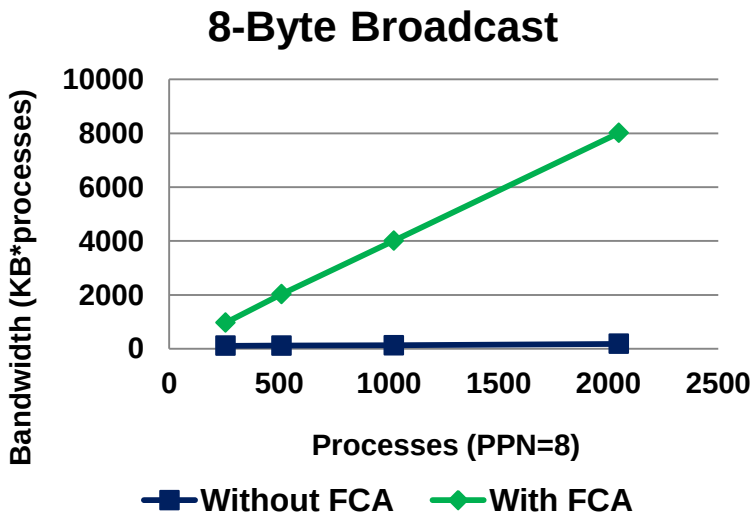

# FCA – Application Acceleration

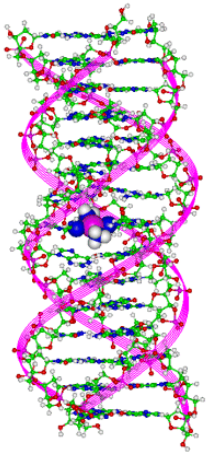

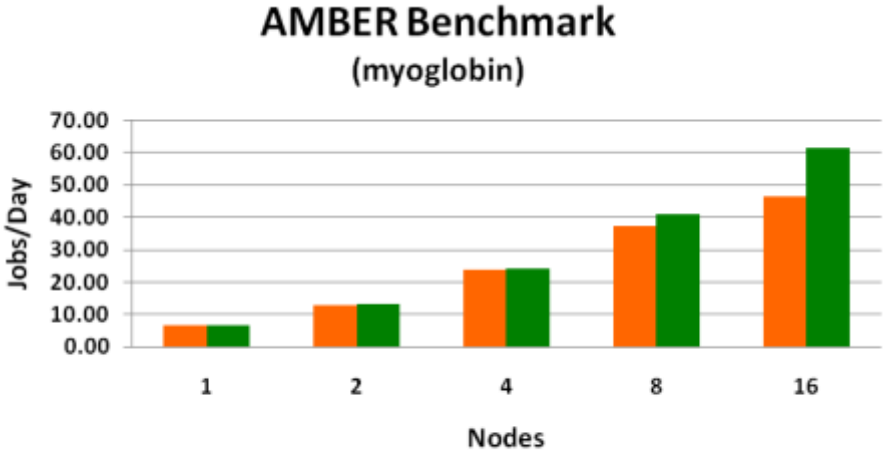


CPMD

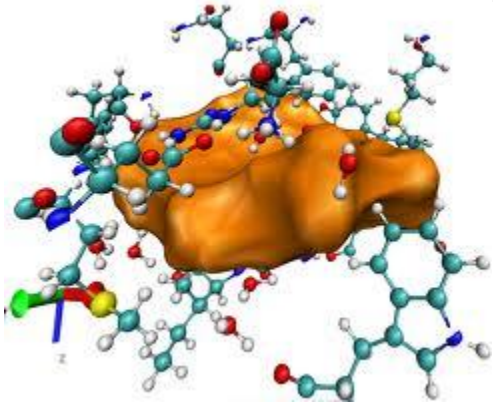

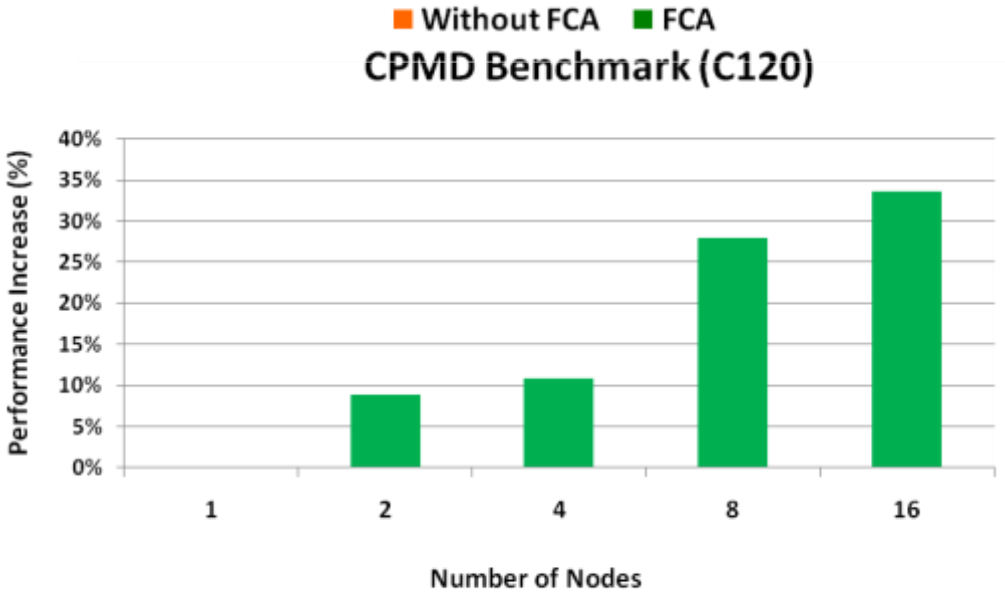


***Acknowledgment: HPC Advisory Council for providing the performance results**

# Network Congestion

- Desired behavior
  - A-G…F-G – 1/6 link BW
  - X-Y – 5/6 link BW
- Congestion effect – lossless:
  - A-G…E-G – 1/6 link BW
  - F-G – 1/12 link BW
  - X-Y – 1/12 link BW
- Congestion effect – lossy:
  - Massive packets' drop
  - Application-visible impact

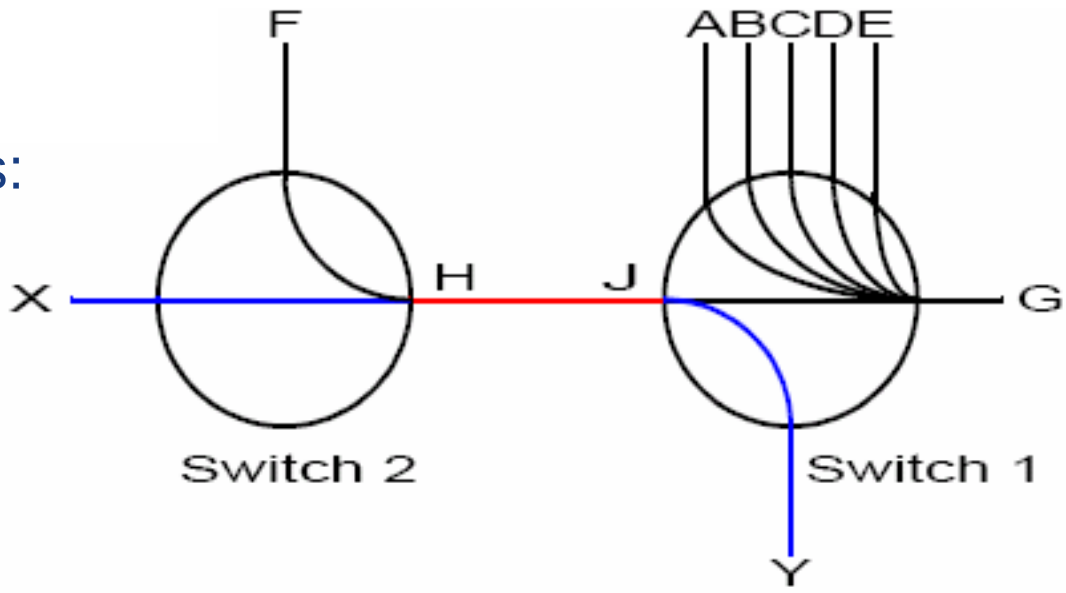


***Solution – Slow Down Injection Rate***

# HW-based Congestion Handling – InfiniBand

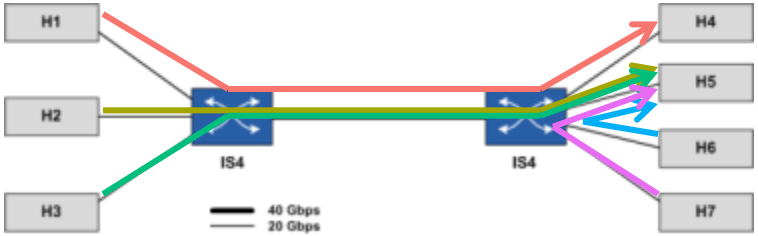


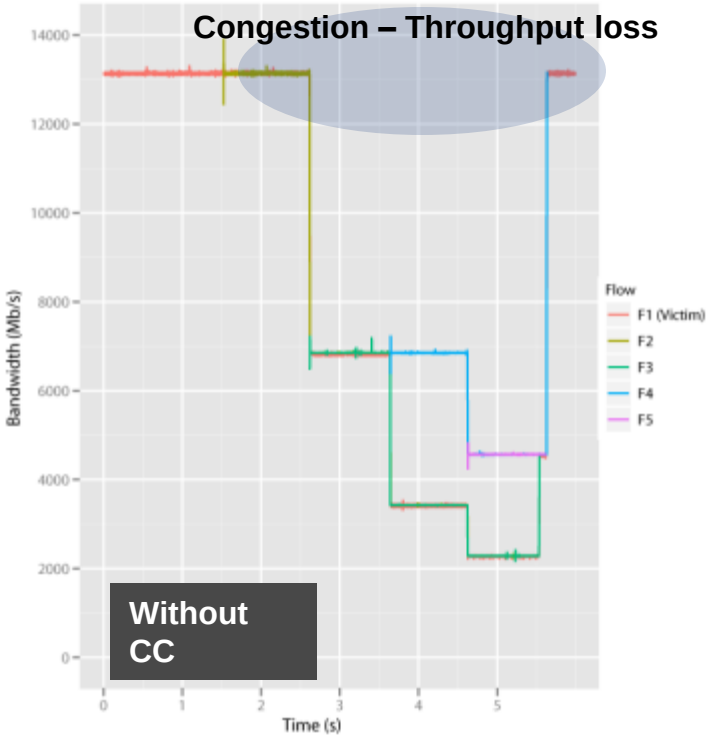


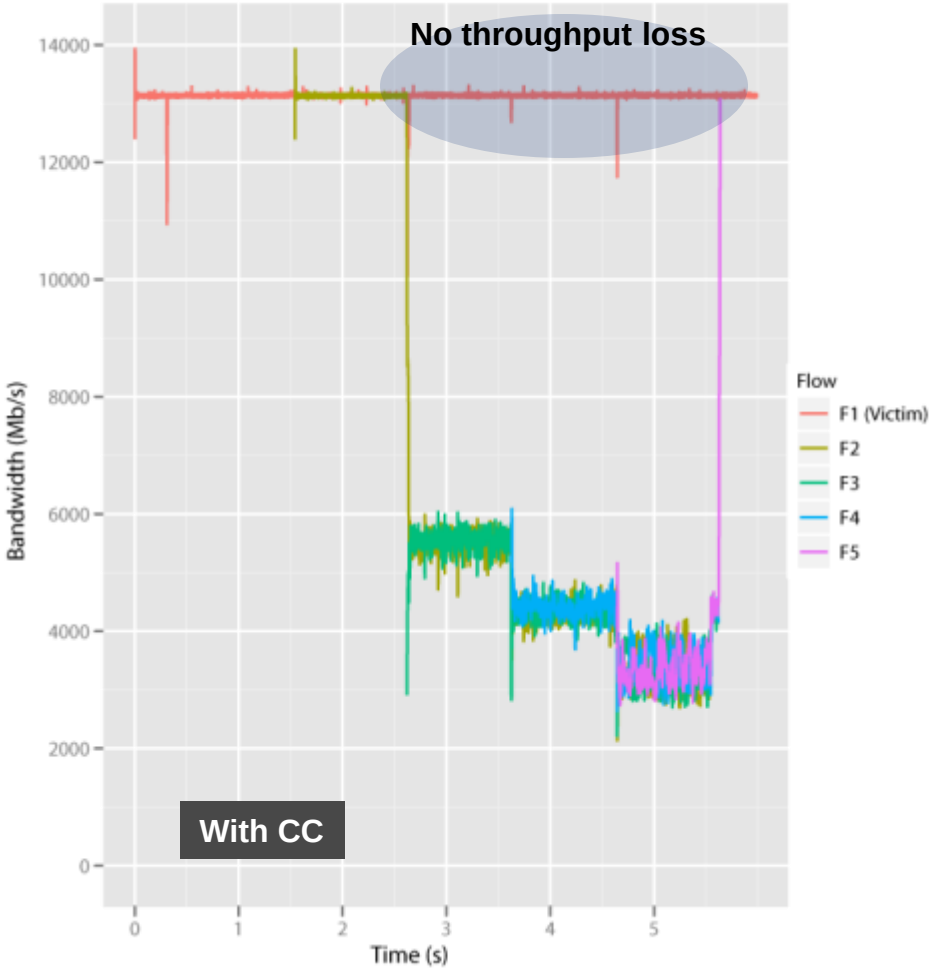

# Application Performance

| Network Latency | % Improvement |
|---|---|
| Ping Pong Latency | 88% |
| Natural Ring Latency | 81.6% |
| Random Ring Latency | 81.3% |
| Ping Pong bandwidth | 85.5% |
| **Applications (HPCC)** | **% Improvement** |
| PTRANS | 76% |
| FFT | 40% |

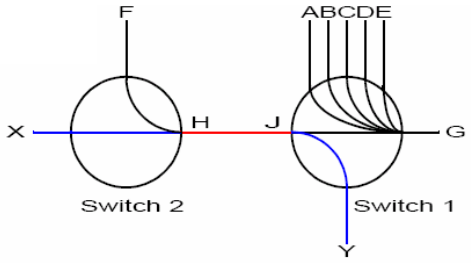


**For more performance examples: "First Experiences with Congestion Control in InfiniBand Hardware"; Ernst Gunnar Gran, Magne Eimot, Sven-Arne Reinemo, Tor Skeie, Olav Lysne, Lars Paul Huse, Gilad Shainer; IPDPS 2010**

[ simula . research laboratory ] - by thinking constantly about it

**Congestion Free Network For Highest Efficiency**

# In-Network Congestion

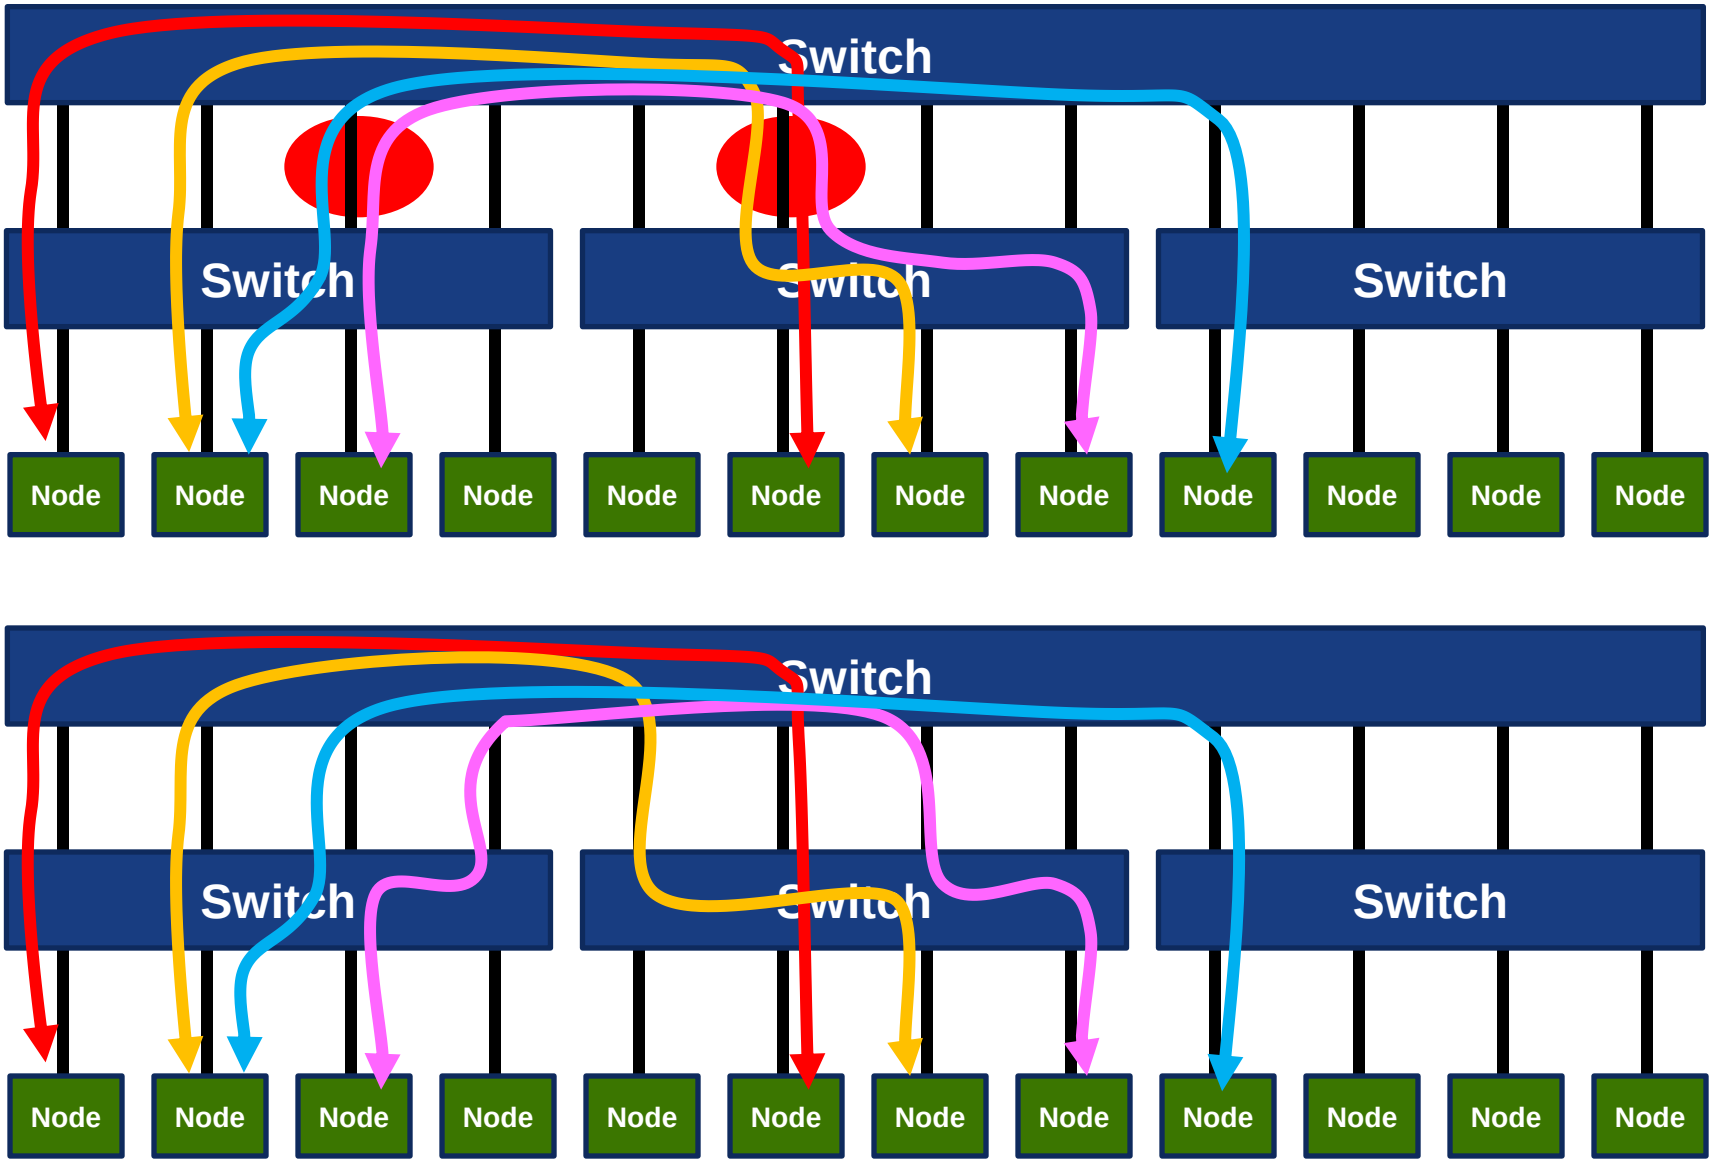

# Adaptive Routing Operation

- Static Network Configuration
  - Identify equivalent paths
  - Program Routing Tables with alternative paths, set decision thresholds
- Run-Time

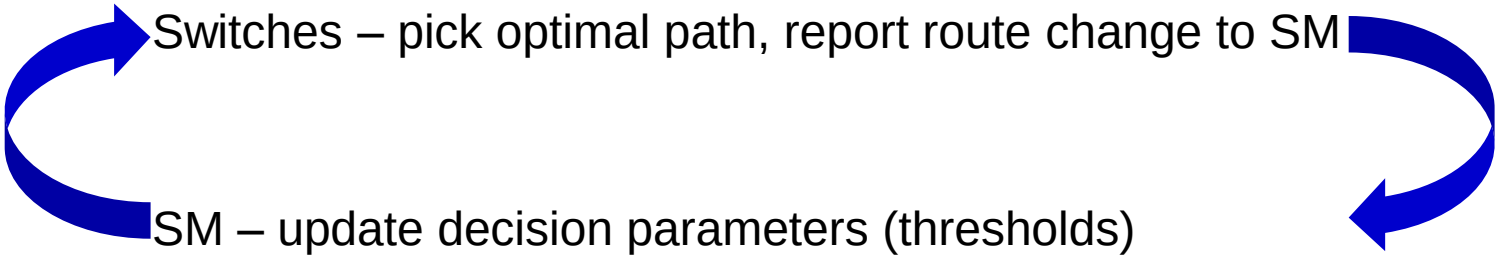


**220 server node system, Mellanox InfiniBand HCAs and switches**

| | No Adaptive Routing | | With Adaptive Routing | |
|---|---|---|---|---|
| Hot Spots Configuration | Average Bandwidth | Minimum Bandwidth | Average Bandwidth | Minimum Bandwidth |
| None | 100.0% | | NA | NA |
| 2:41,3:4,4:1 | 79.4% | 37.3% | 99.7% | 99.6% |
| 2:35,3:5,4:1 | 80.3% | 37.1% | 99.7% | 99.0% |
| 2:30,3:10,4:1 | 75.9% | 36.5% | 99.7% | 99.5% |
| 2:41,3:5,4:1 | 82.4% | 37.7% | 99.7% | 99.4% |

# Scale-out Network Topologies

- CLOS (fat-tree)
  - Can be fully non-blocking (1:1) or blocking (x:1)
  - Typically enables best performance

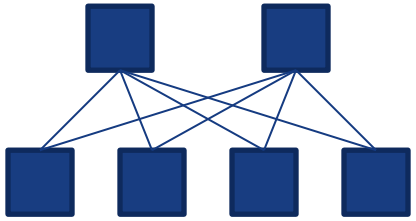

- Mesh or 3D Torus
  - Great performance solutions for applications with locality
  - Simple expansion for future growth

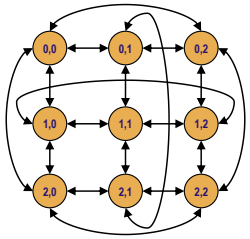


- Dragonfly
  - Hierarchical scale-out
  - Cabling-optimized topology

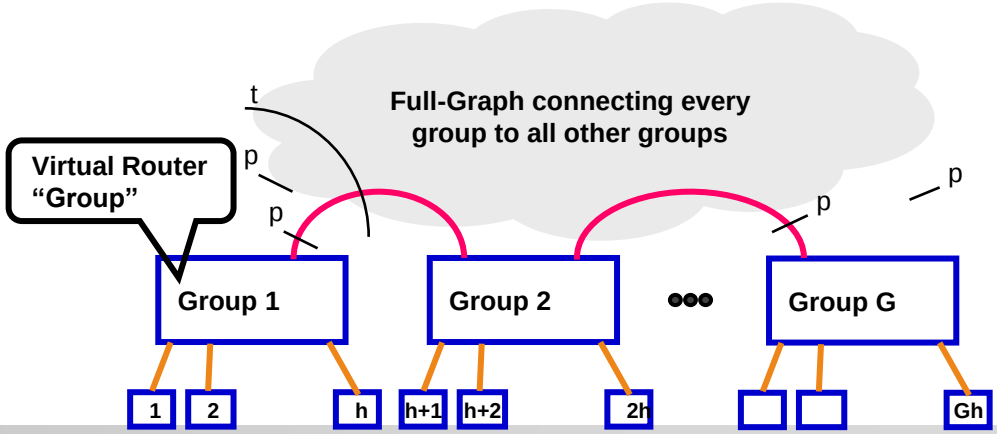

# Biggest cluster

# On the Way to Exascale – Scalability Challenge

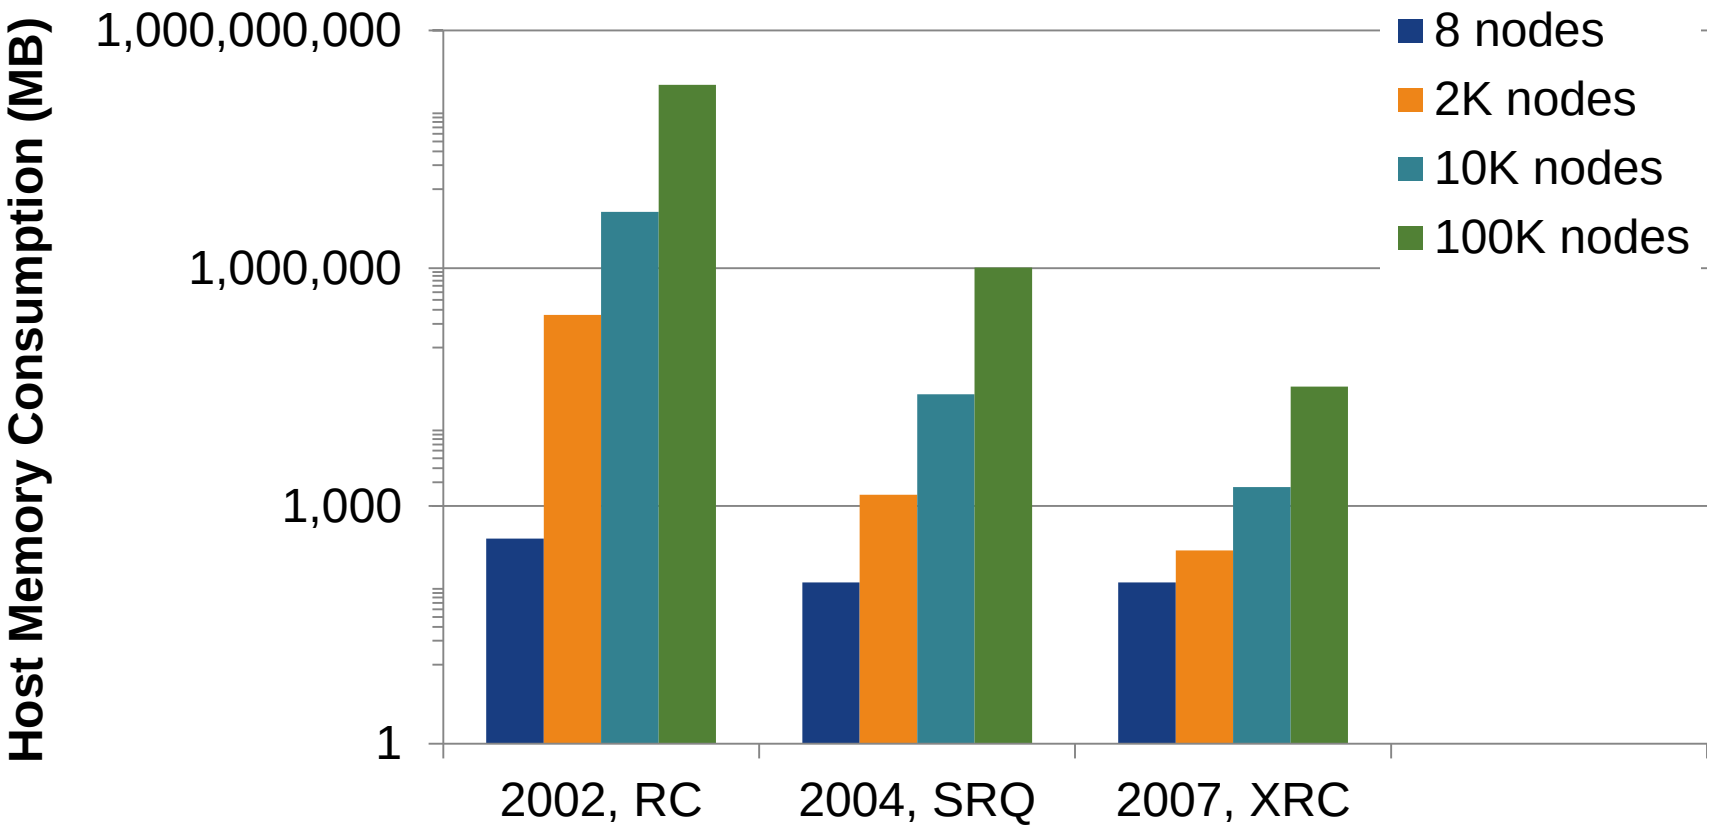

# The New Kid in Town

## InfiniHost

**World's first InfiniBand HCA**

**10Gb/s InfiniBand**
**PCI-X host interface**
**1 million msg/sec**


**2002**

## InfiniHost III

**World's first PCIe InfiniBand HCA**

**20Gb/s InfiniBand**
**PCIe 1.0**
**2 million msg/sec**


**2005**

## ConnectX (1,2,3)

**World's first Virtual Protocol Interconnect (VPI) Adapter**

**40/56Gb/s InfiniBand**
**PCIe 2.0, 3.0 x8**
**33 million msg/sec**


**2008-11**

## Connect-IB

**The Exascale Foundation**

Connect IB™

**June 2012**

# Announcing Connect-IB: The Exascale Foundation

- Salability Solution at Transport Layer
  - Unlimited inter-node Scalability
  - Intra-node scalability above 256 cores
- Data Integrity
  - Internal datapath ECC
  - Unbreakable link technology
  - Application-level data integrity
- World's fastest adapter
  - 100Gbit/Sec bandwidth, <0.7uSec latency
  - 130M messages/sec, 20M IOPs/sec
- HW virtualization
  - 512 Virtual Machines at HW performance
  - VM-transparent migration

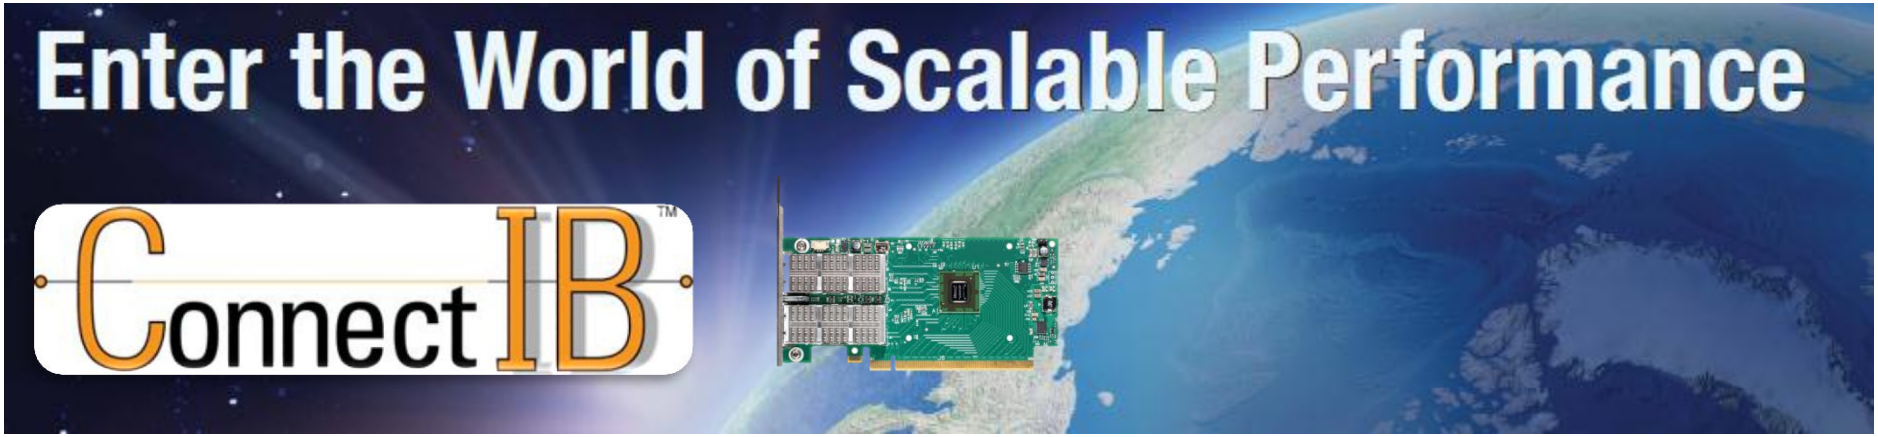

# ConnectIB – Exascale Scalability

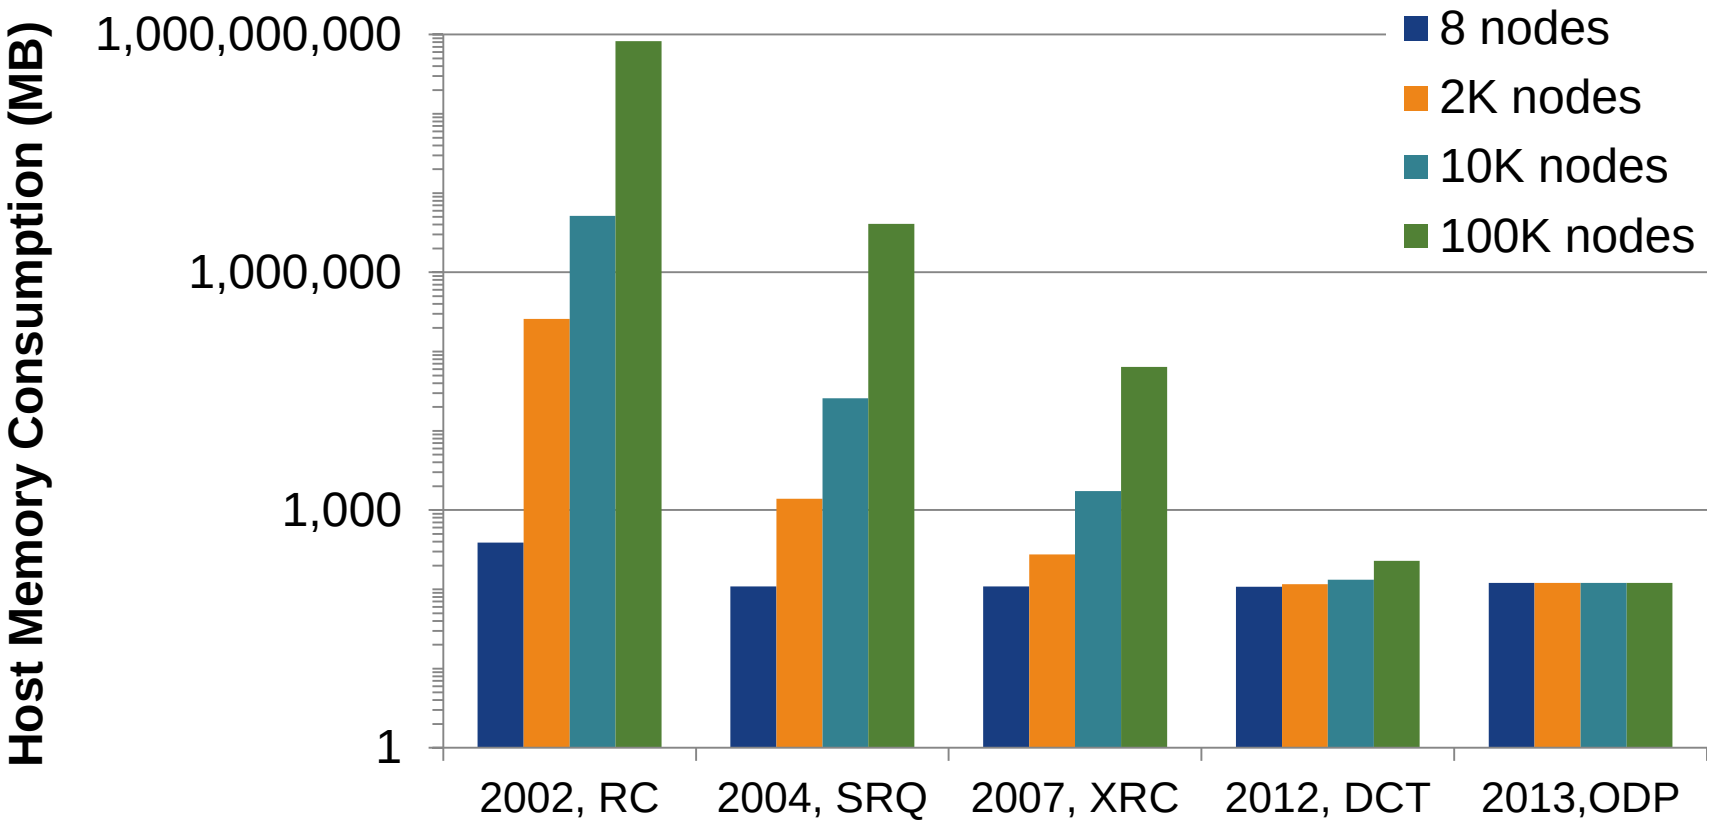

# Storage

# Solving the Storage Bottleneck with RDMA

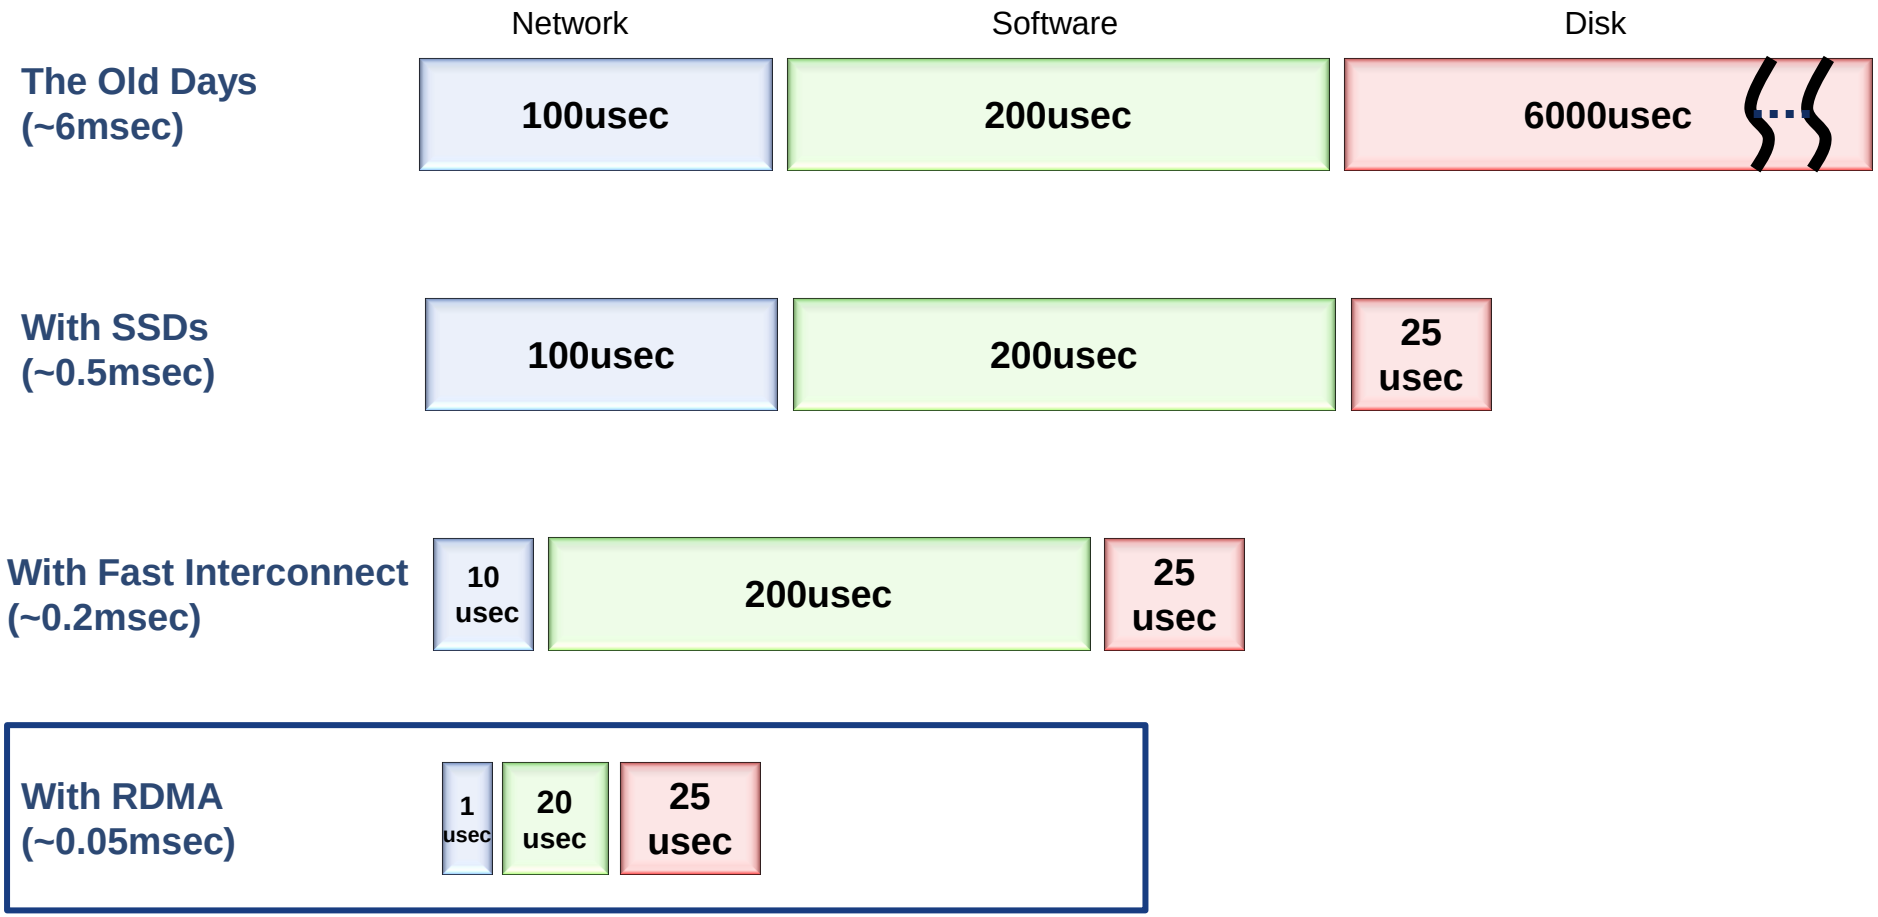

# Mellanox Storage Accelerator – RDMA Block Storage

- High-performance target driver for block storage
- iSER (RDMA iSCSI) for data path
- Supports InfiniBand and Ethernet
- Key Features
  - Native IB Storage
  - Scale-out solution (non-disruptive)
  - Central storage management
  - Integrated Flash caching
  - Optional FC Storage Gateway


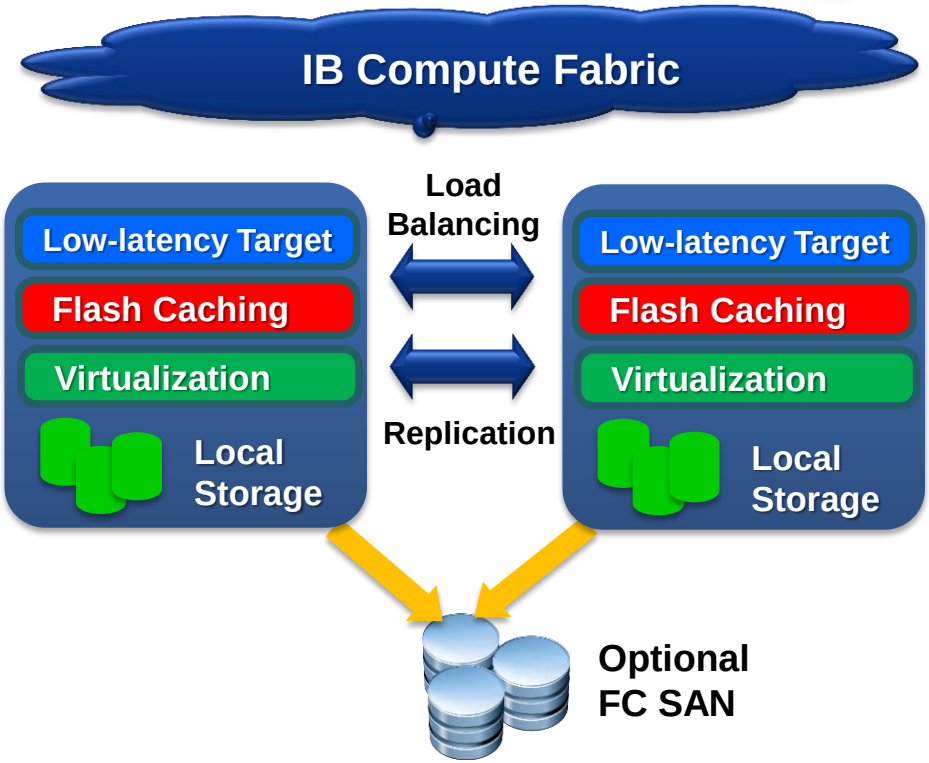

# Evaluation Platform

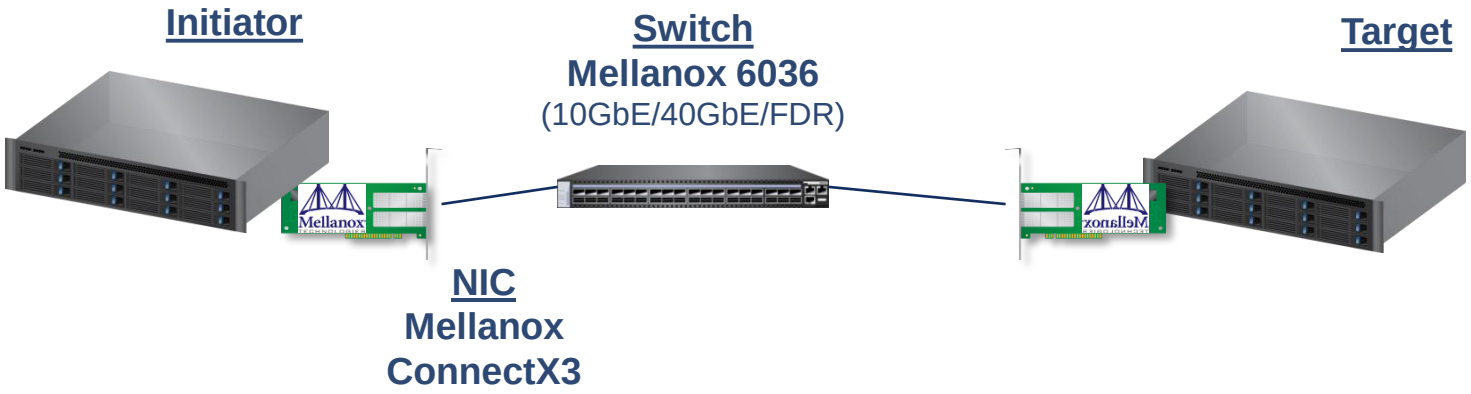


## Hardware/OS configuration

- SuperMicro X9DRW
- 32GB RAM, 2 x E5-2650 CPUs
- PCIe Gen3
- RHEL 6.2 x64
- ConnextX3 VPI (FW: 2.10.800)
- Drivers: MOFED-1.5.3-3.1.0

## Storage configuration

- VSA ver: 2.1.1-1.el6
- 3 target processes
- 6 RAM LUNs (in total)
- Client use 8 threads and IO depth of 16-64 IOs

# Bandwidth Tests

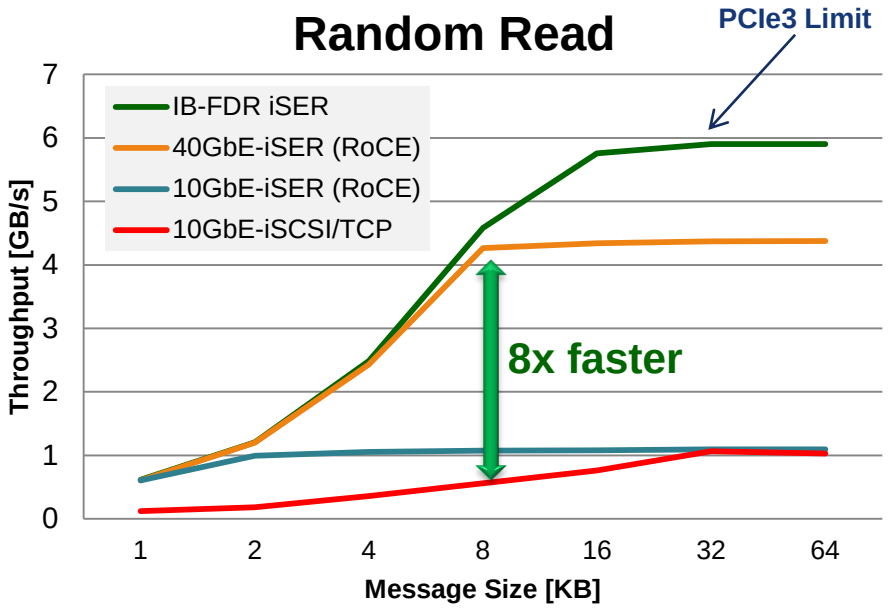


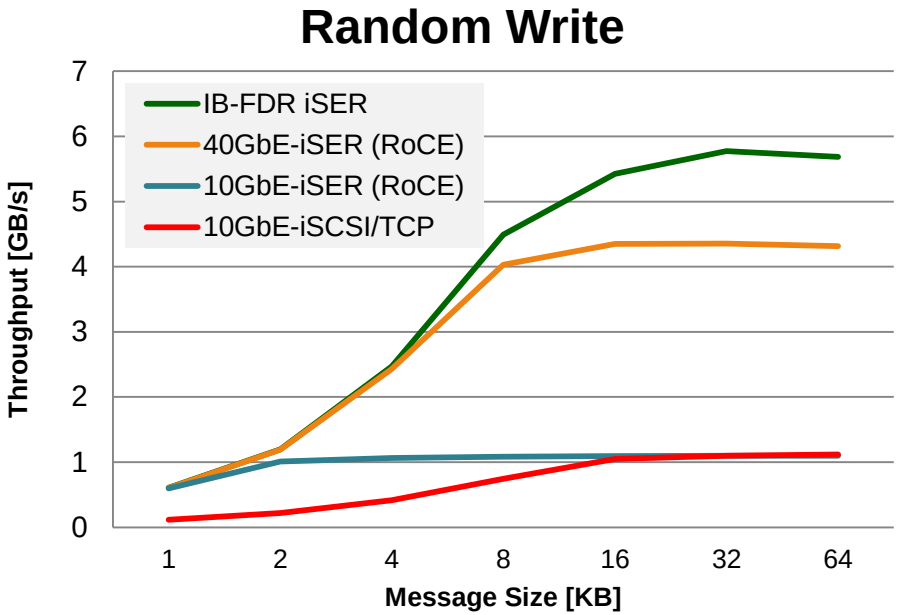


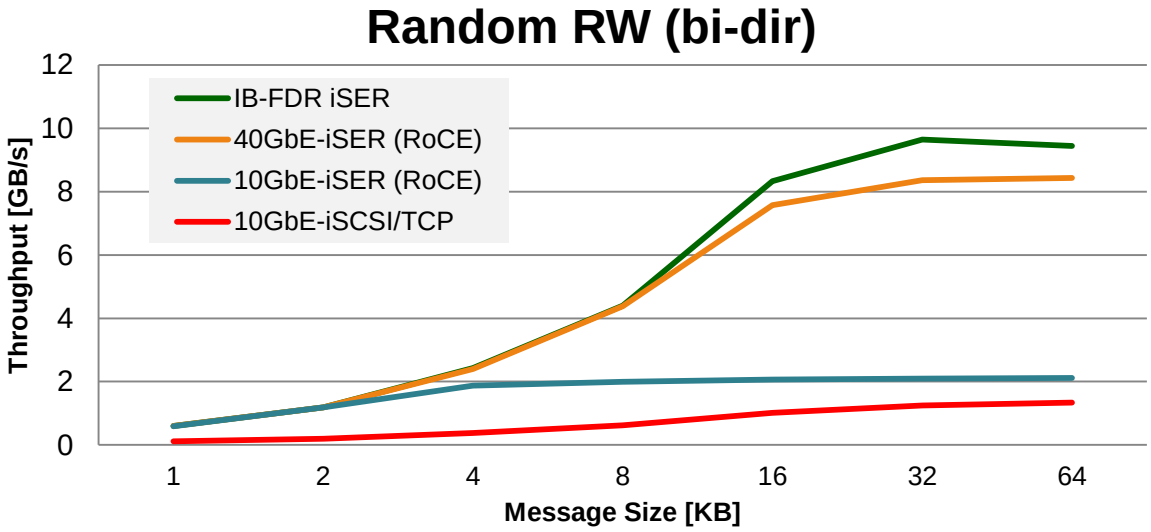

# Single I/O Latency (Synchronous I/O)

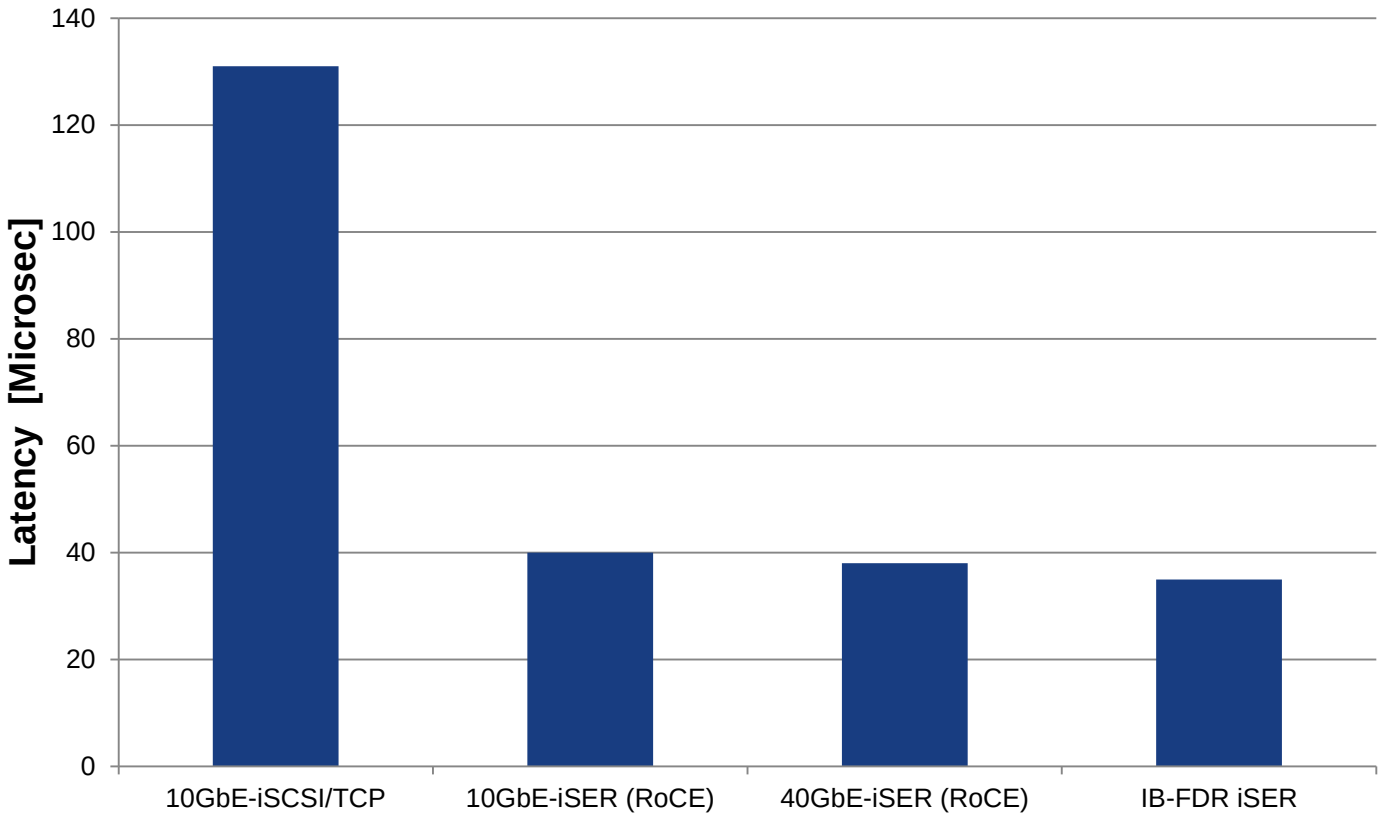


**3-4x lower latency (on Synchronous load)**

# IOPS and CPU Utilization

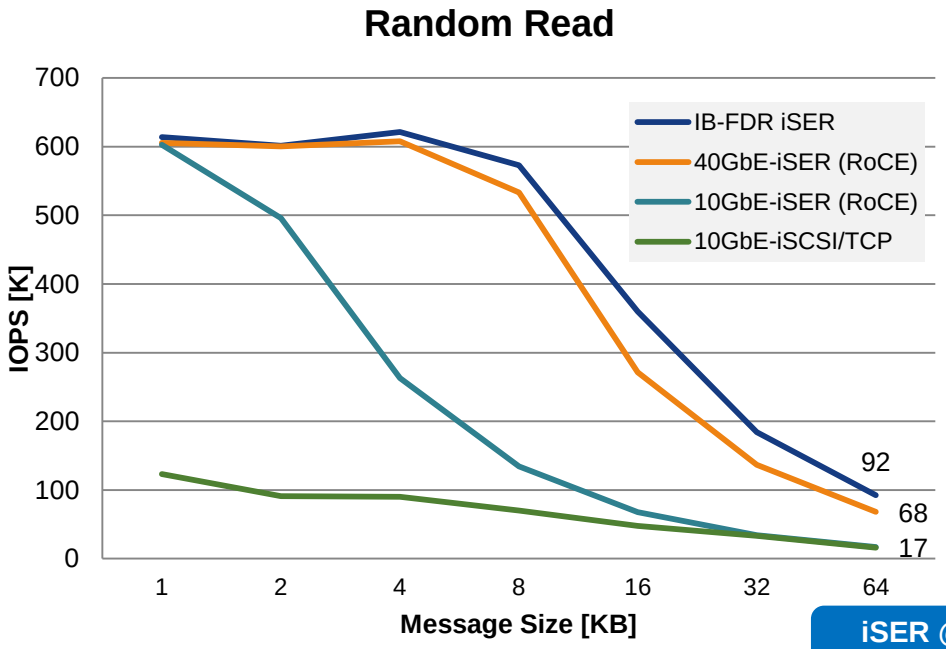


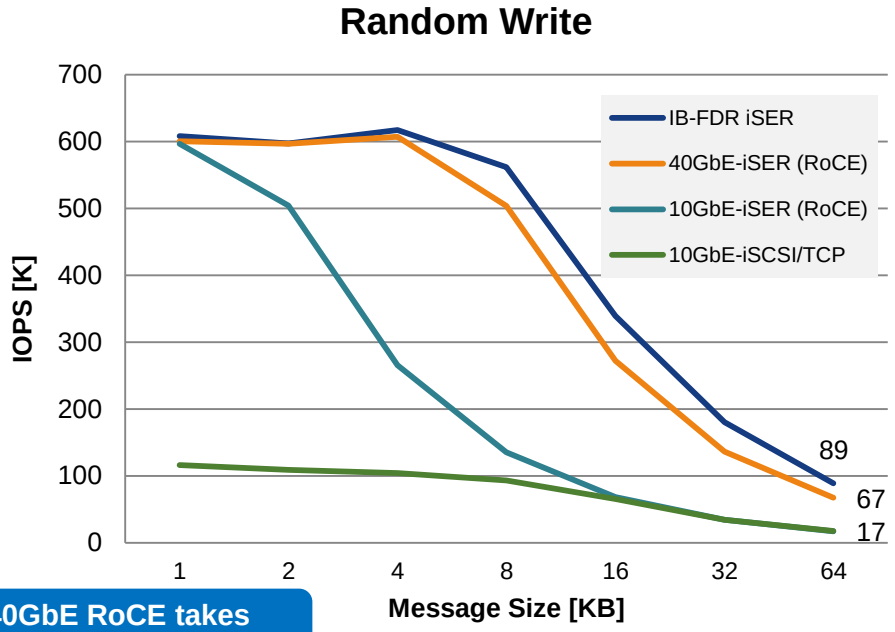


**iSER @ 40GbE RoCE takes 5-7x less CPU per IOP**

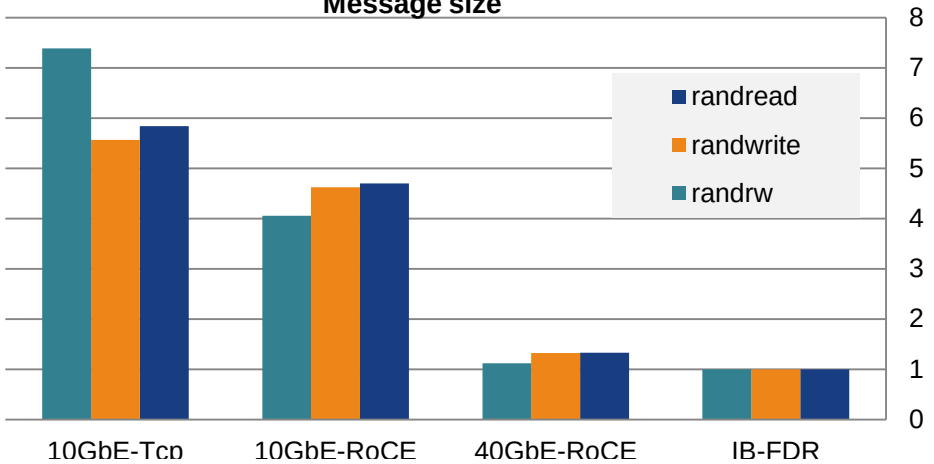


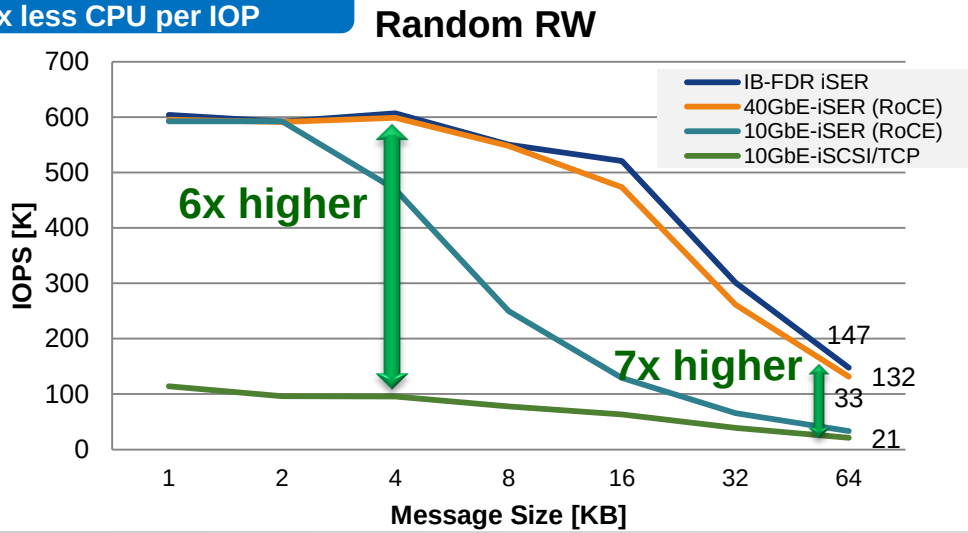

# Summary

- iSER at 10GbE RoCE shows significant performance improvements over 10GbE iSCSI/TCP

| | 10GbE iSCSI (TCP) | 10GbE iSER (RoCE) | 40GbE iSER (RoCE) |
|---|---|---|---|
| Bandwidth @ 8KB | 1x | 1.9x | 7.6x |
| Single IO Sync Latency | 1x | 0.3x | 0.3x |
| Full load IO Async Latency @ 4KB Read | 1x | 0.32x | 0.13x |
| IOPs @ 4KB Read | 1x | 2.9x | 6.8x |
| CPU Utilization per IOP @ 4KB Read | 1x | 0.29x | 0.17x |

- iSER at 40GbE RoCE captures most of the benefits of FDR
  - BW, latency, IOPs and CPU utilization

# FDR InfiniBand Meets the Needs of Changing Storage World

SSDs, the storage hierarchy, In-Memory Computing…..

Remote I/O access needs to be equal to local I/O access

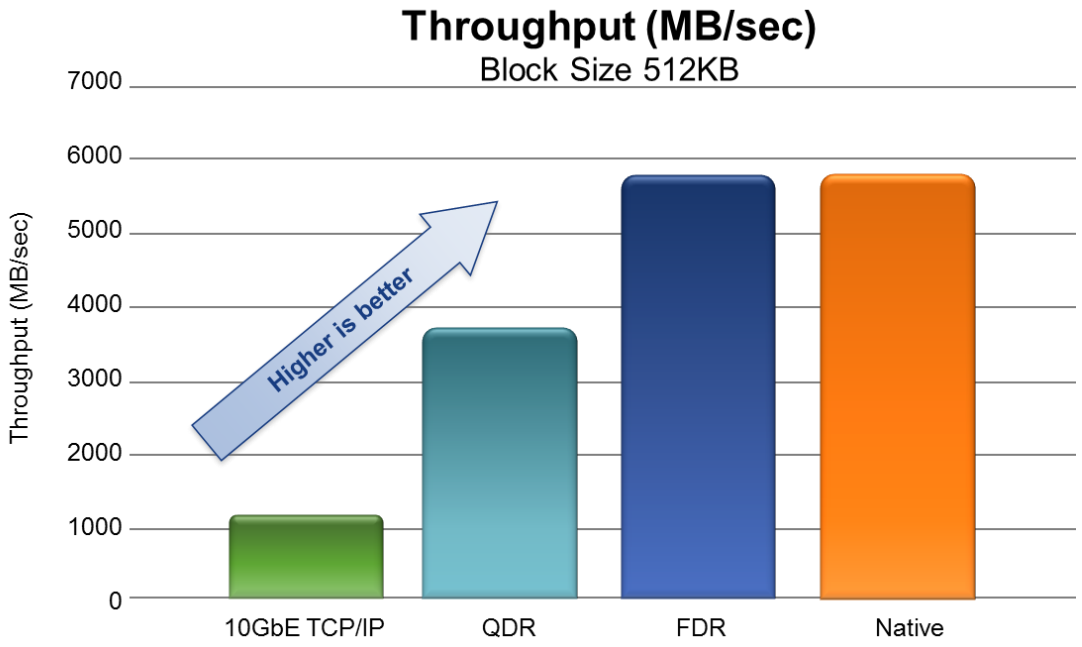


**SMB 3.0 + RDMA (InfiniBand FDR)**

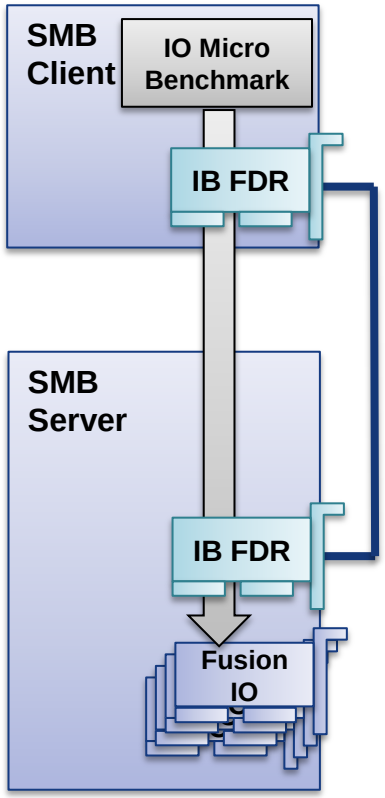


**Native Throughput Performance over InfiniBand FDR**

# Mellanox Interconnect: Standard for Storage and Database

## Oracle ZFS

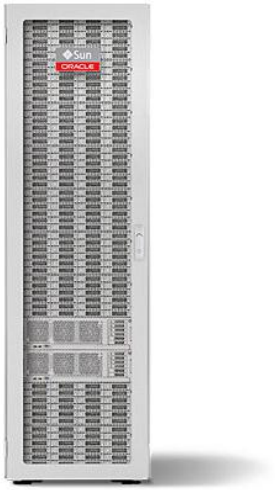

## EMC Isilon

ISILON SYSTEMS

EMC² where information lives®

**Scale out Network Attached Storage**

## IBM XIV Storage System

**Power to serve even more applications from a single system**

## IBM DB2 pureScale

**95% Efficiency and Scalability**

# Server Message Block Direct – SMB over RDMA

- Increased throughput
- Low Latency
- Low CPU overhead
- Fault Tolerance

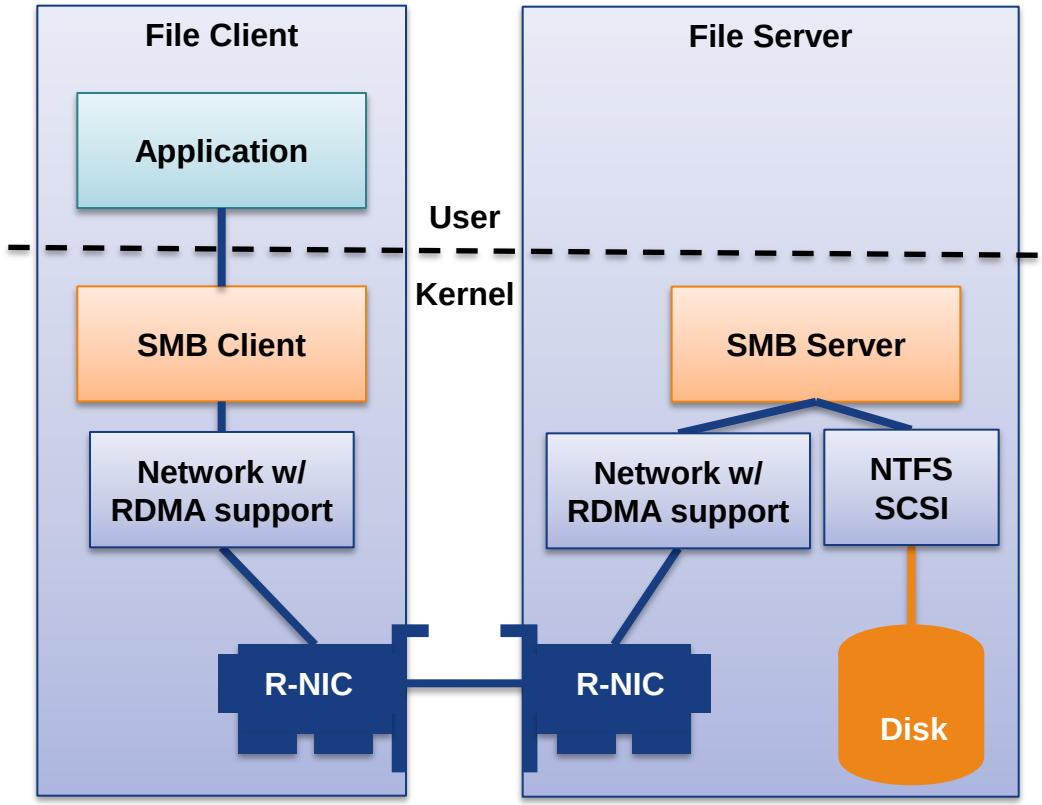


**New class of SMB file storage for the Enterprise**

# SMB Direct over InfiniBand and RoCE

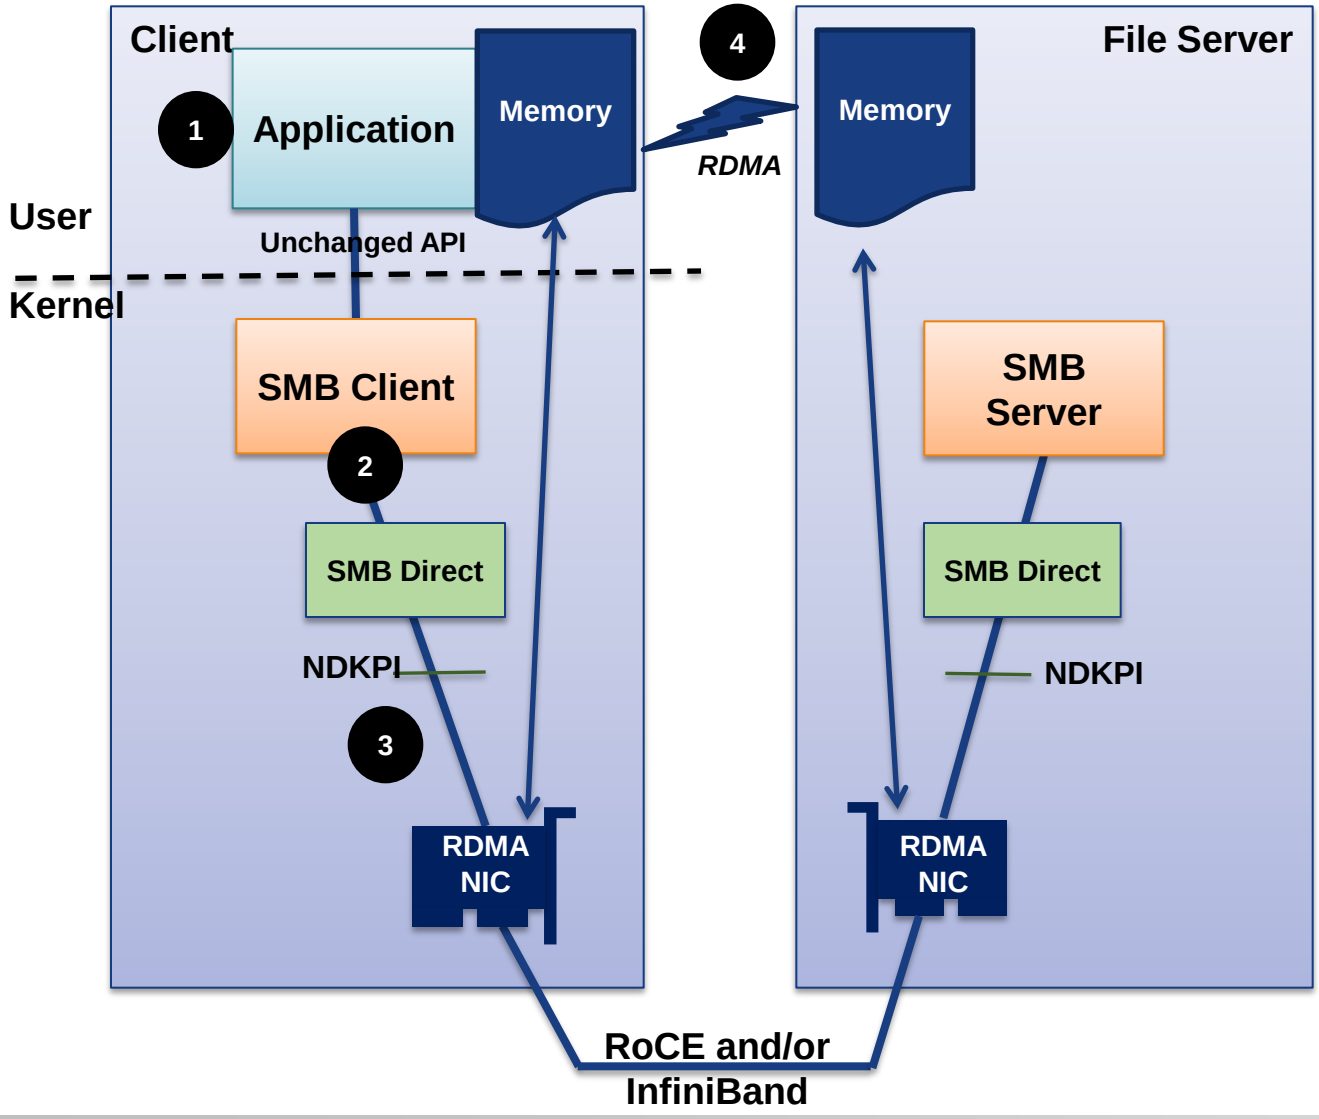


1. Application (Hyper-V, SQL Server) does not need to change.
2. SMB client makes the decision to use SMB Direct at run time
3. NDKPI provides a much thinner layer than TCP/IP
4. Remote Direct Memory Access performed by the network interfaces.

# Measuring SMB Direct Performance

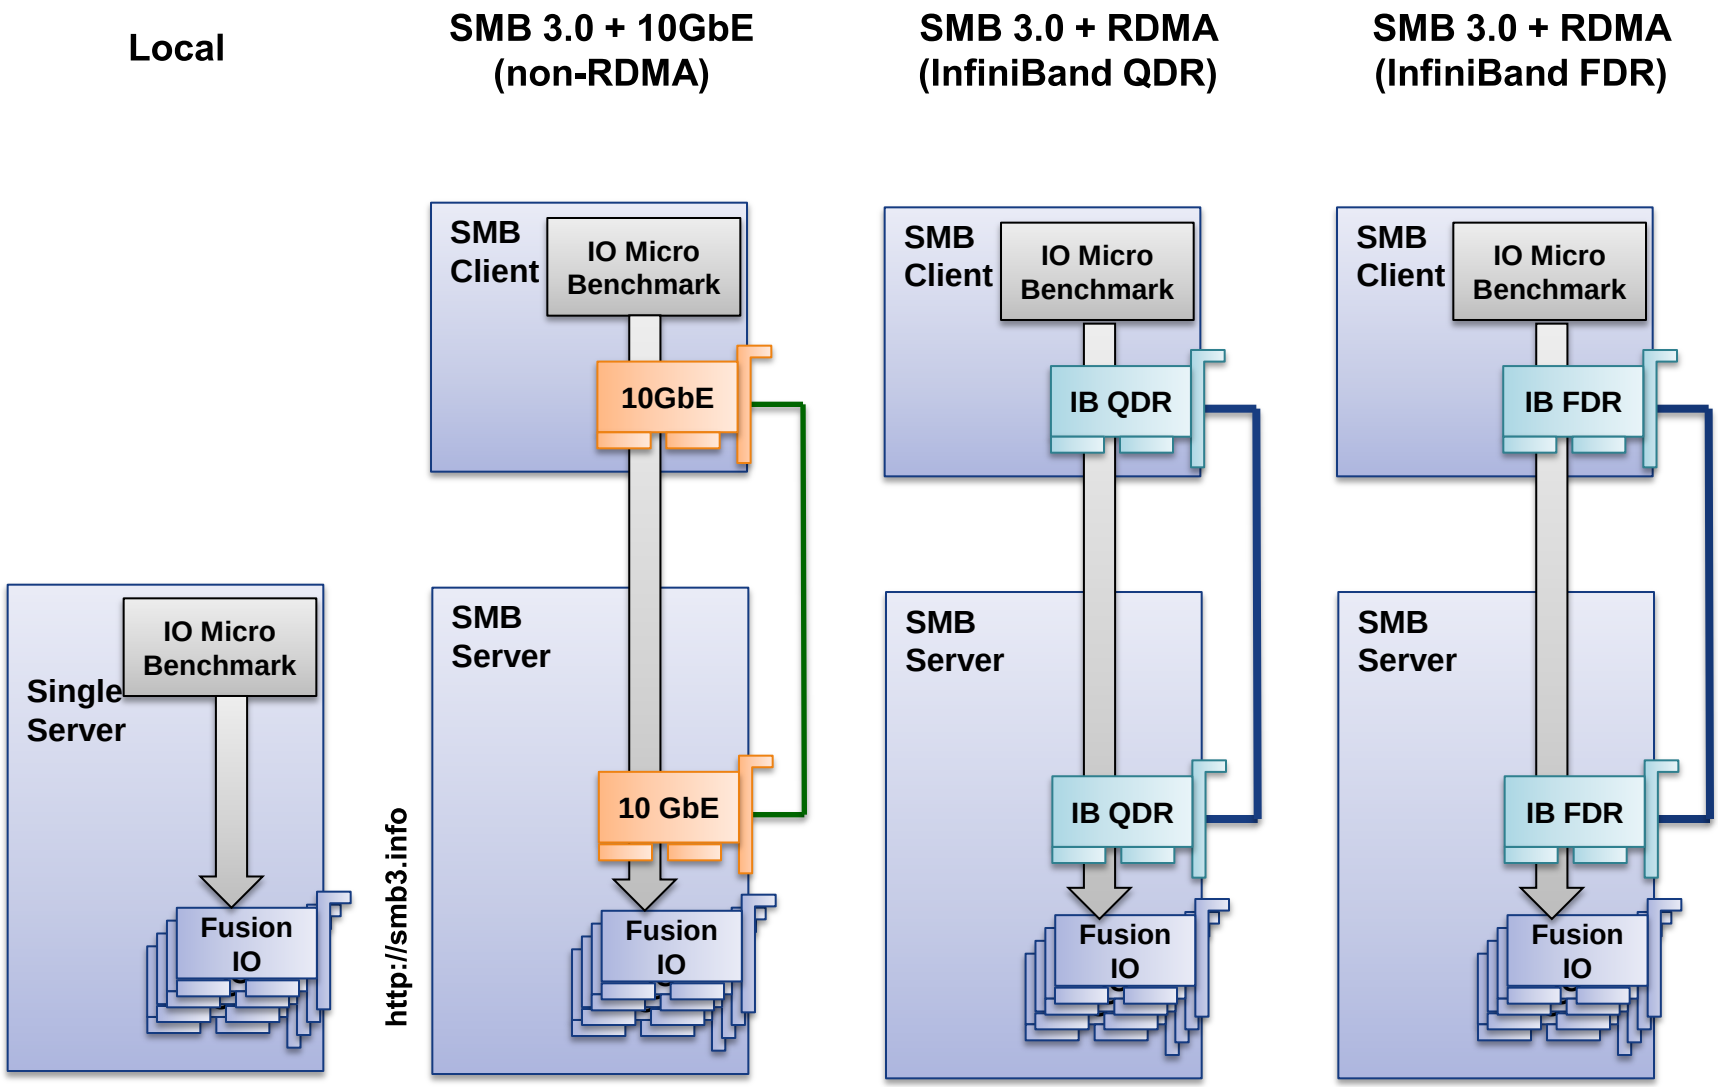

# Microsoft Delivers Low-Cost Replacement to High-End Storage

## FDR 56Gb/s InfiniBand delivers 5X higher throughput with 50% less CPU overhead vs. 10GbE

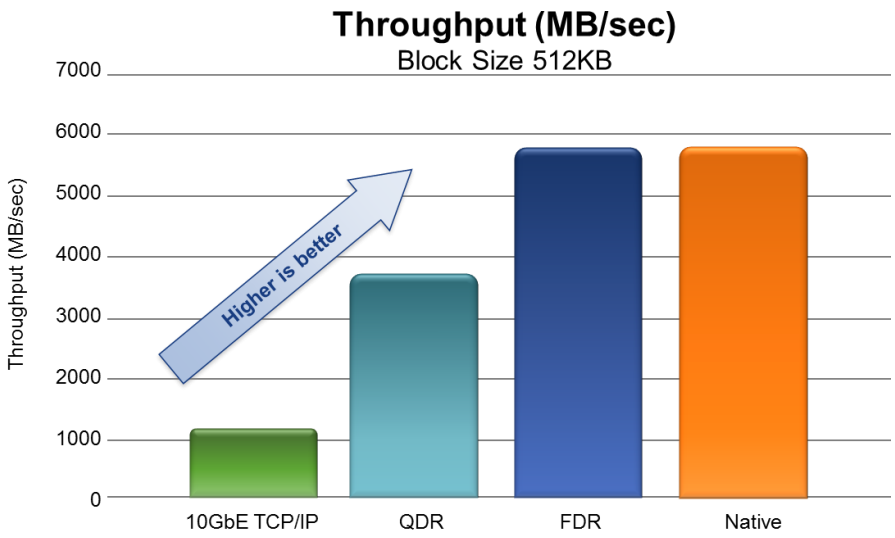


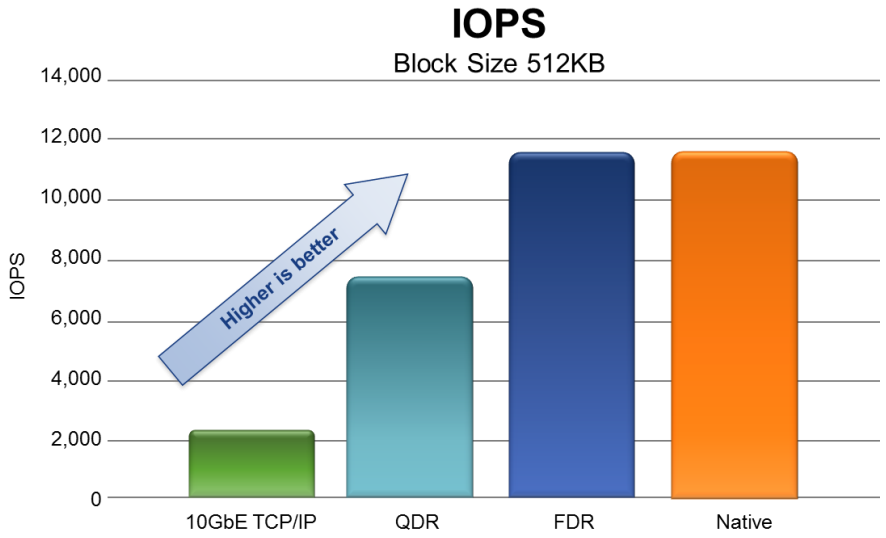


**Native Throughput Performance over FDR InfiniBand**

# Leading Interconnect, Leading Performance

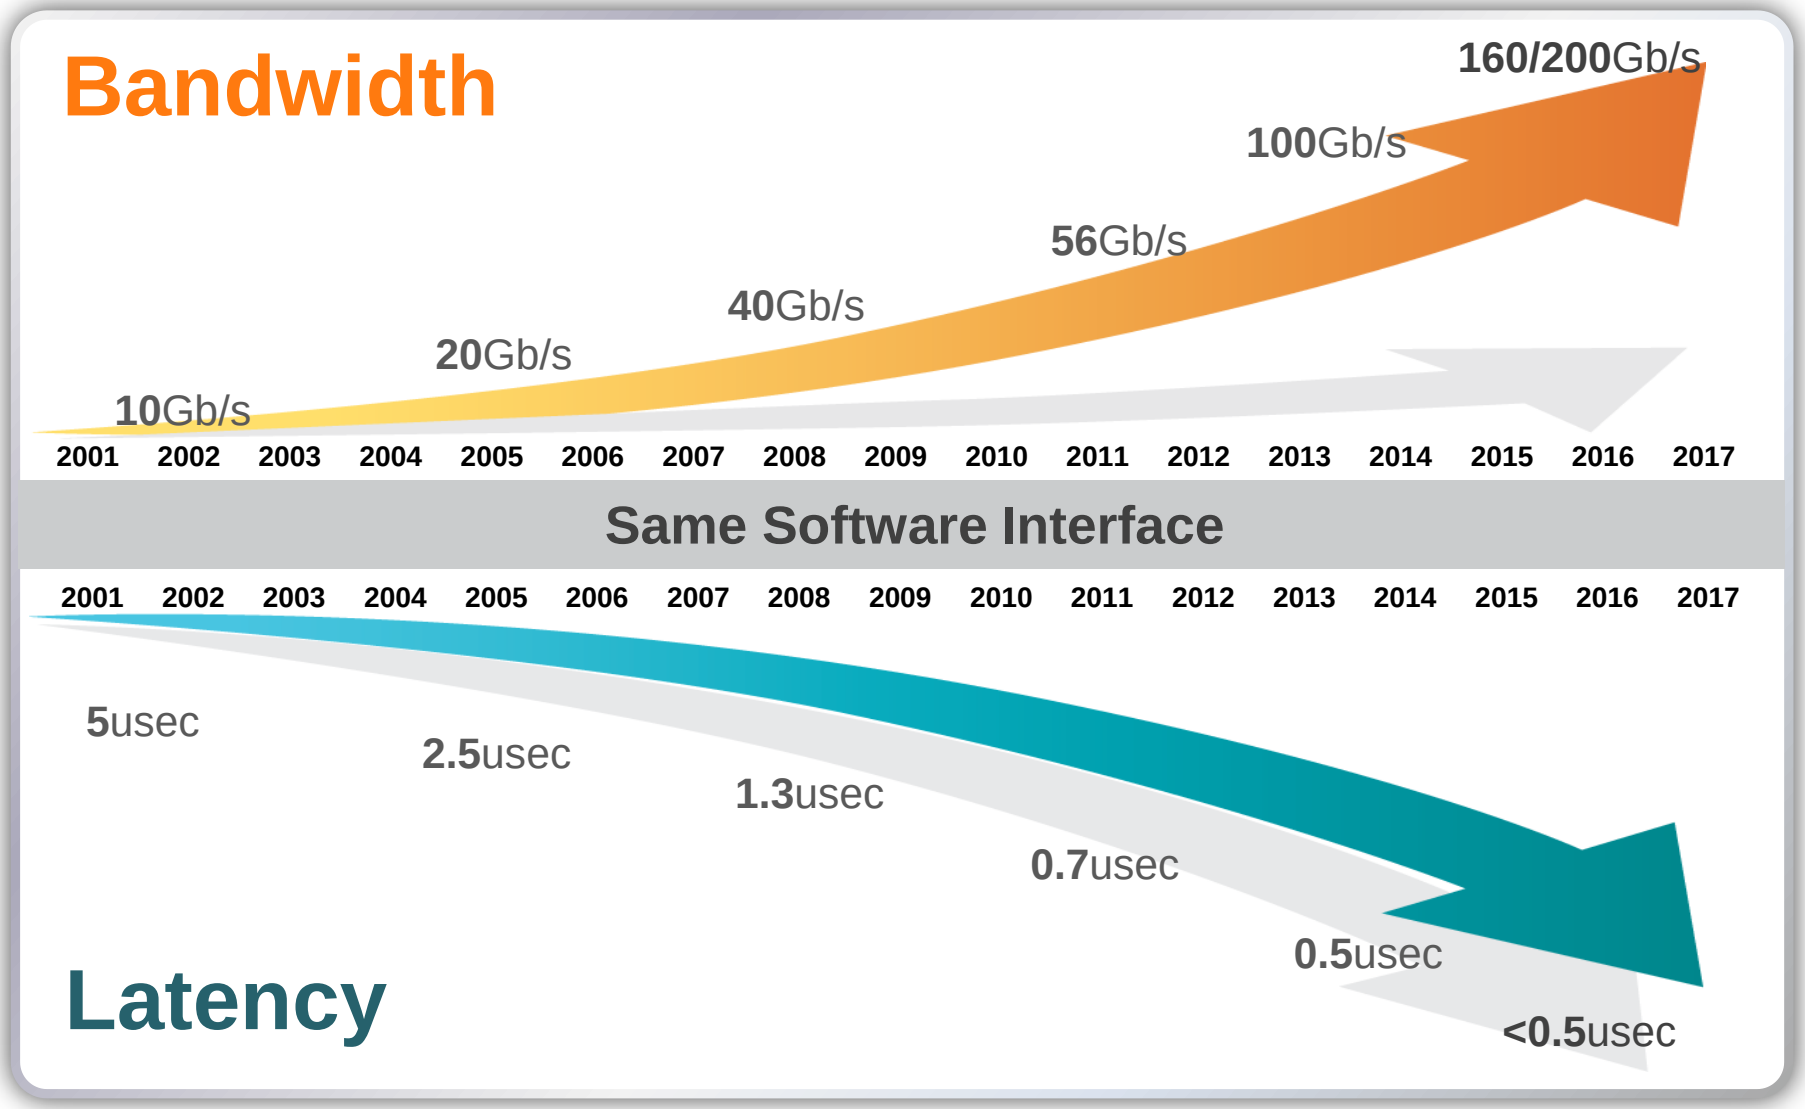

# Advanced Networks Unleashes Opportunities

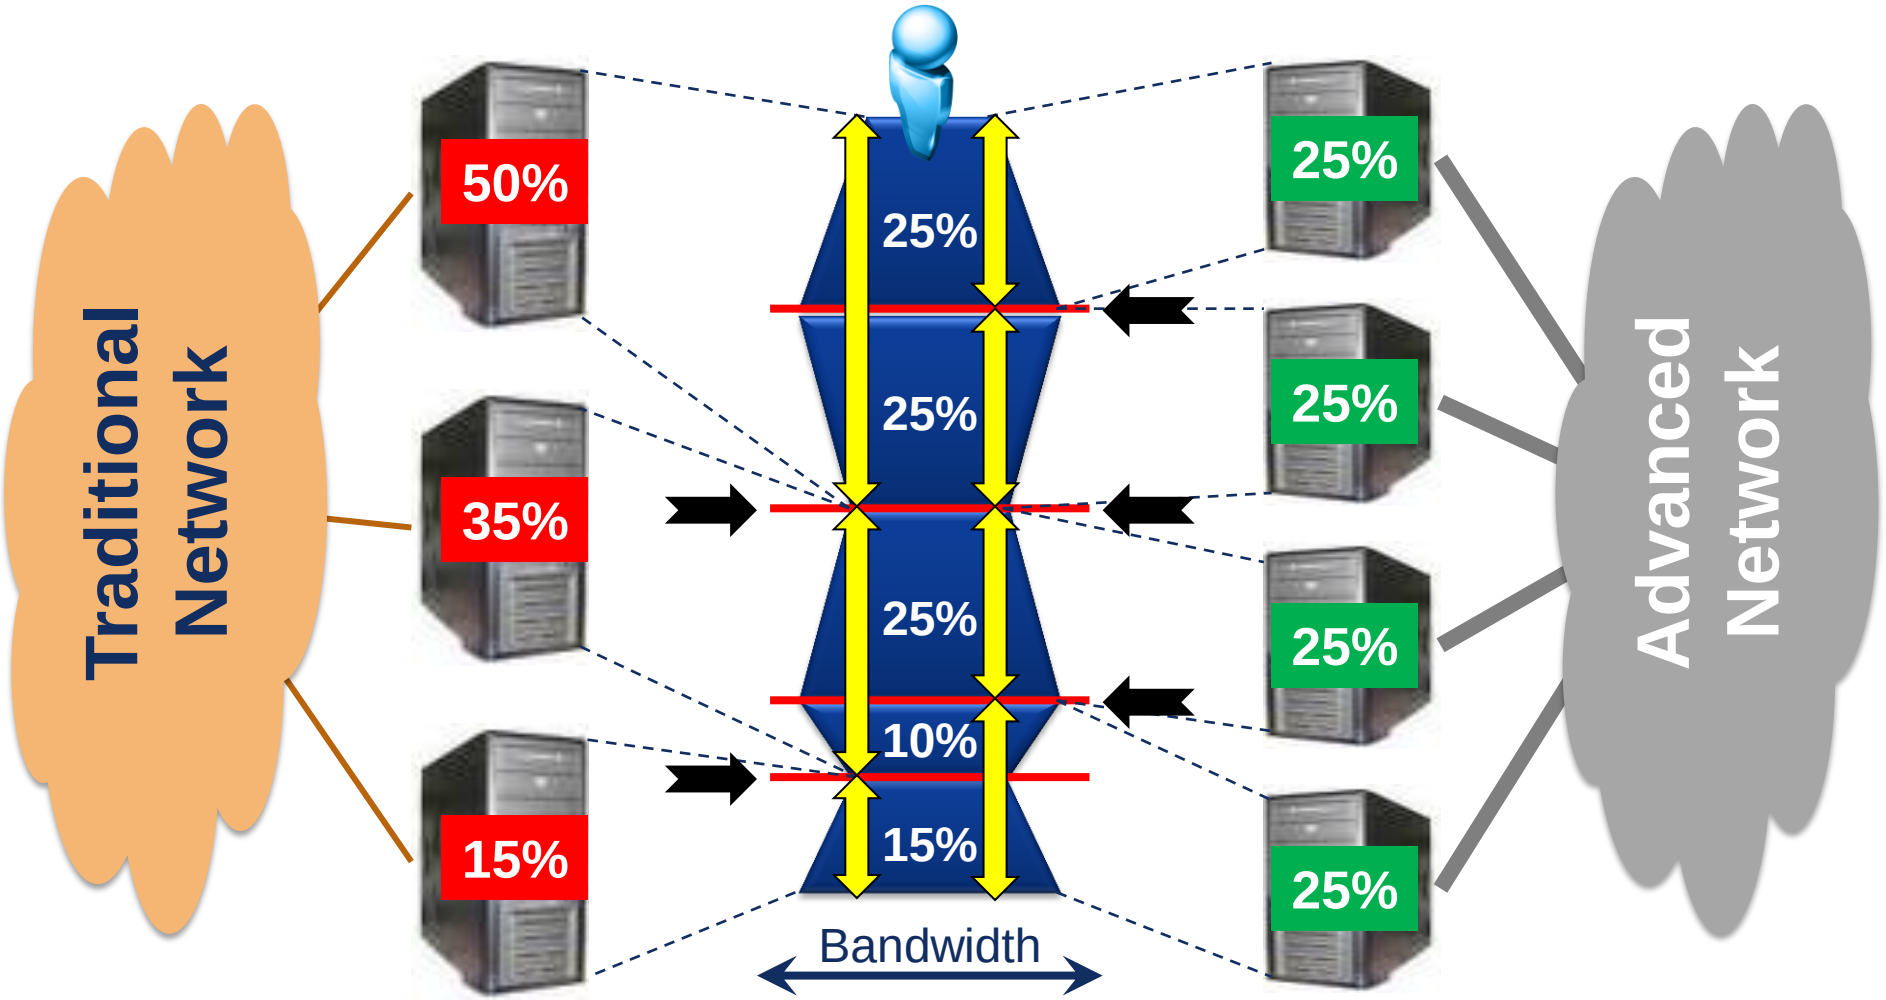


***Network-Constrained***

***Resource-Optimized***

# High-performance Network + RDMA

***"…everybody is using Ethernet, we are using InfiniBand, 40Gbps InfiniBand"***

***Larry Ellison Keynote at Oracle Openworld***

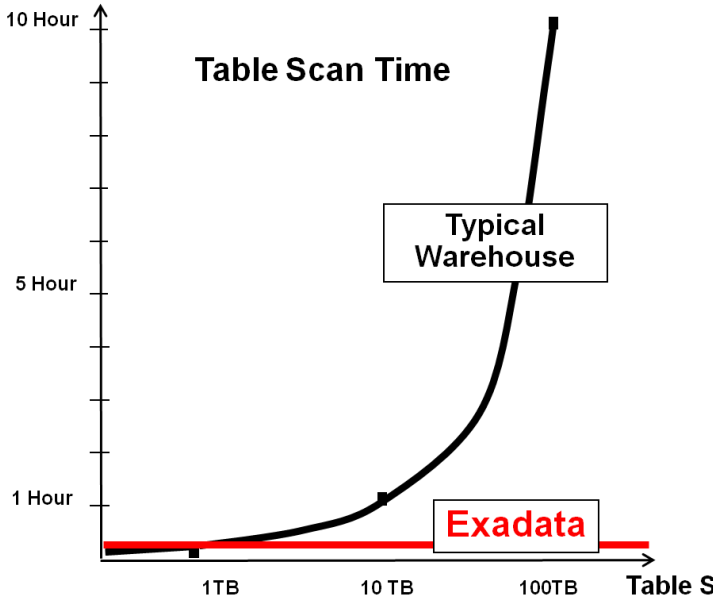


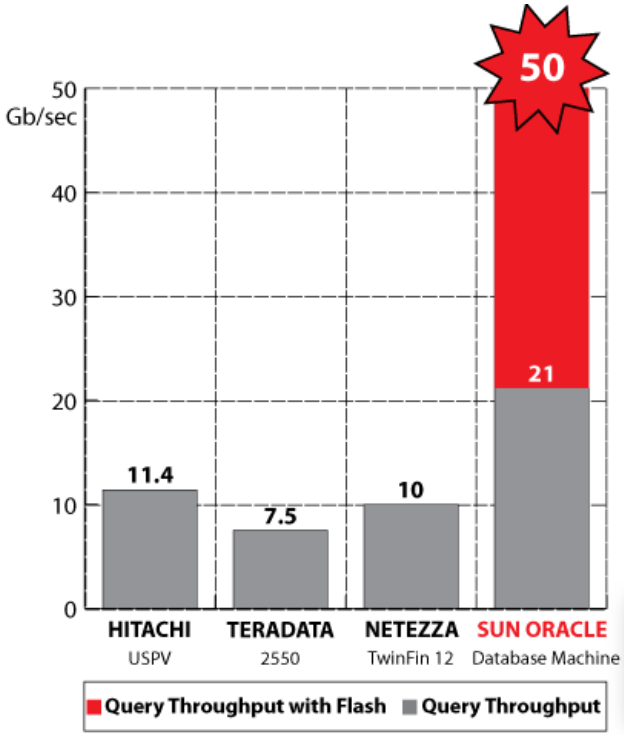



ORACLE EXADATA

# RDMA-accelerated VM migration

## VM Migration with VMware vMotion over RDMA*

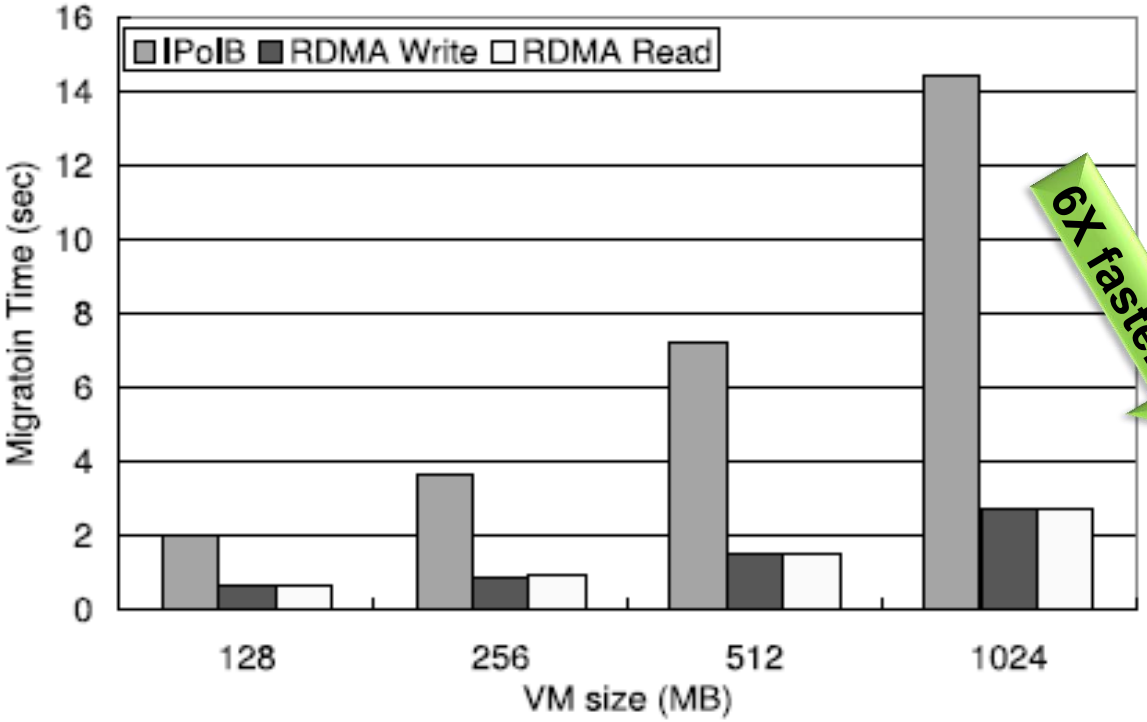


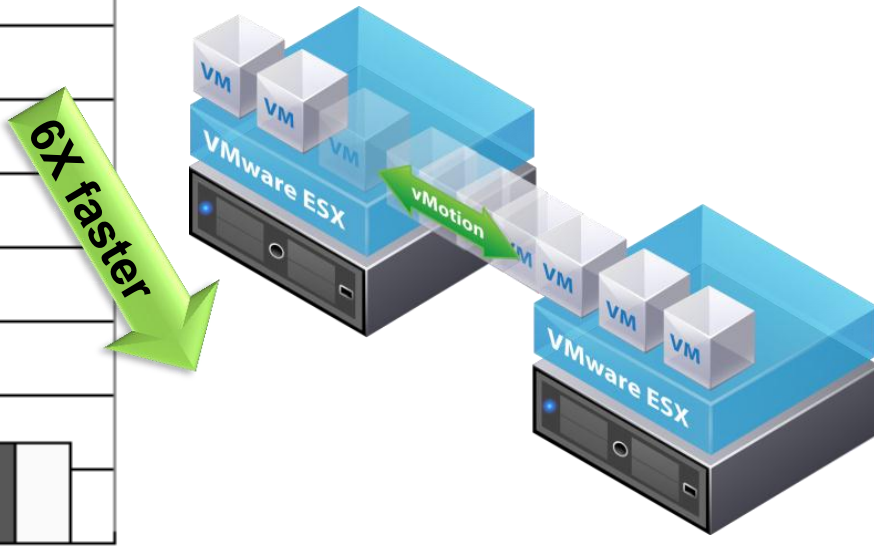


*
**High Performance Virtual Machine Migration with RDMA over Modern Interconnects,**
**Wei Huang, Qi Gao, Jiuxing Liu, Dhabaleswar K. Panda,**
**Computer Science and Engineering, The Ohio State University and**
**IBM T. J. Watson Research Center**

**RDMA Delivers the Value**



62

# Hadoop

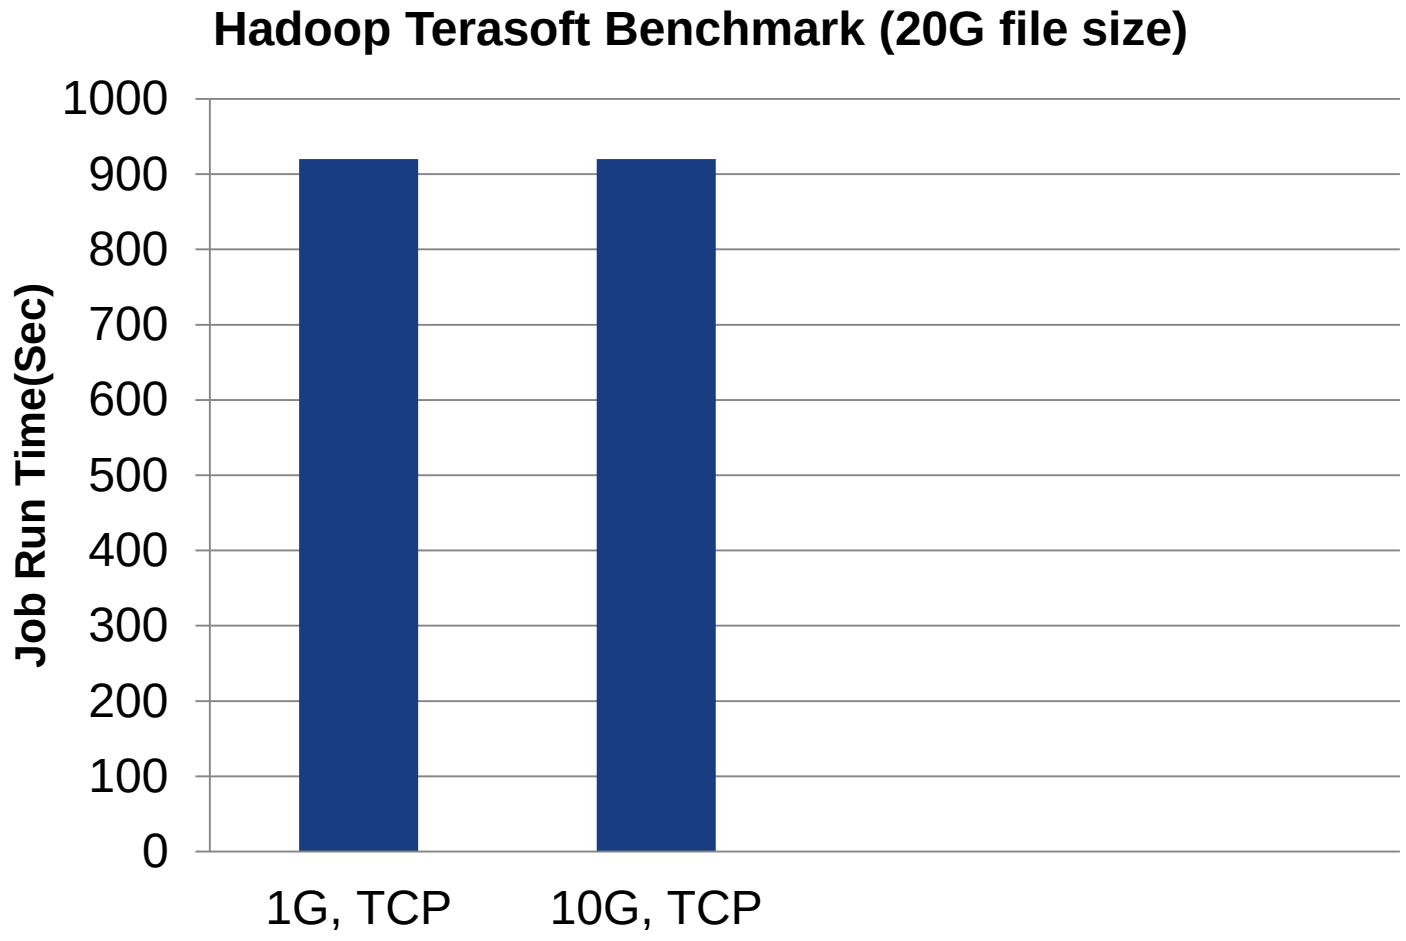

# UDA Plug-in

- Plug-in architecture
  - Hadoop applications are unmodified
  - Open source
  - Enabled via xml configuration
- Efficient MOF transfer over RDMA
  - Supports InfiniBand and Ethernet
  - Zero copy, transport offload, kernel bypass
- Accelerated merge sort

```
        <value>false</value>
    </property>
    <property>
        <!-- mapred.roce.setting represents how
        0: disable roce.
        1: use roce with roce-merger. -->
        <name>mapred.roce.setting</name>
        <value>0</value>
    </property>
    <property>
        <!-- merging approach
```

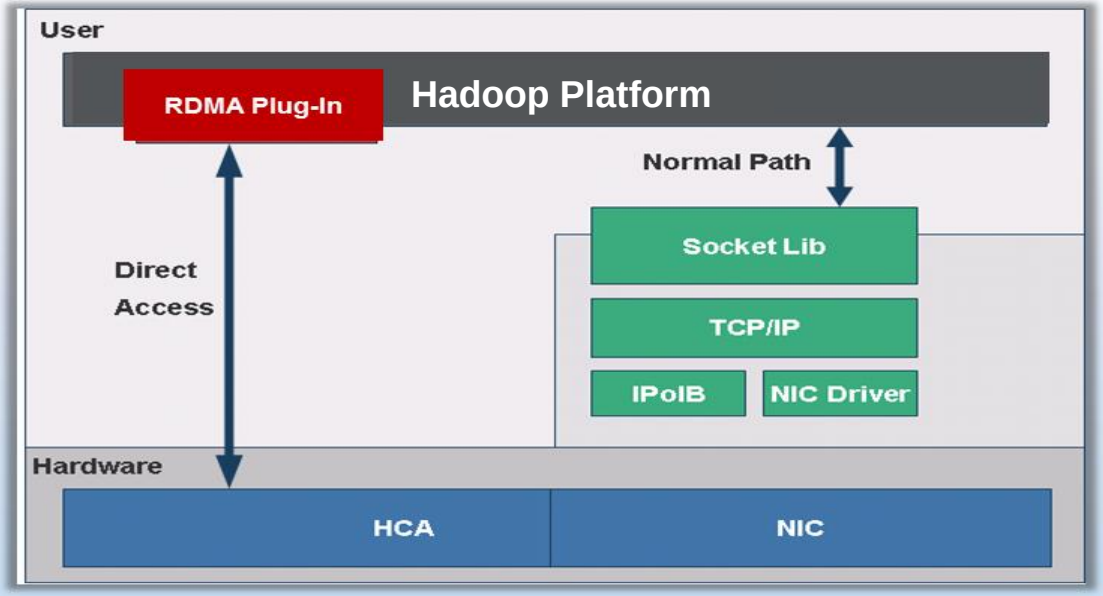

# Advanced Network Interface Value

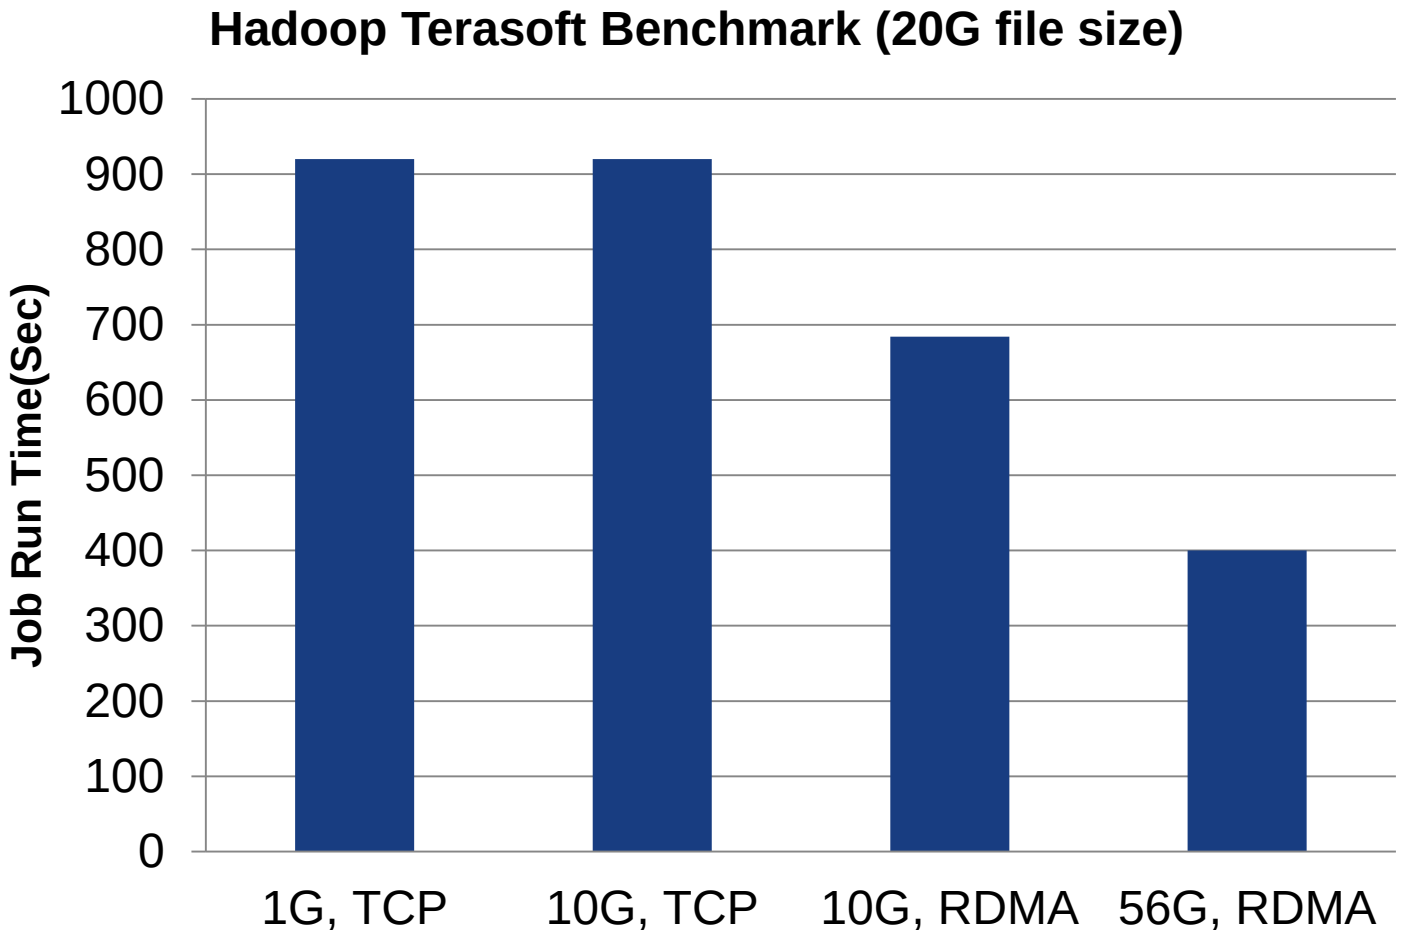

# Local vs. Remote Access – the Cost

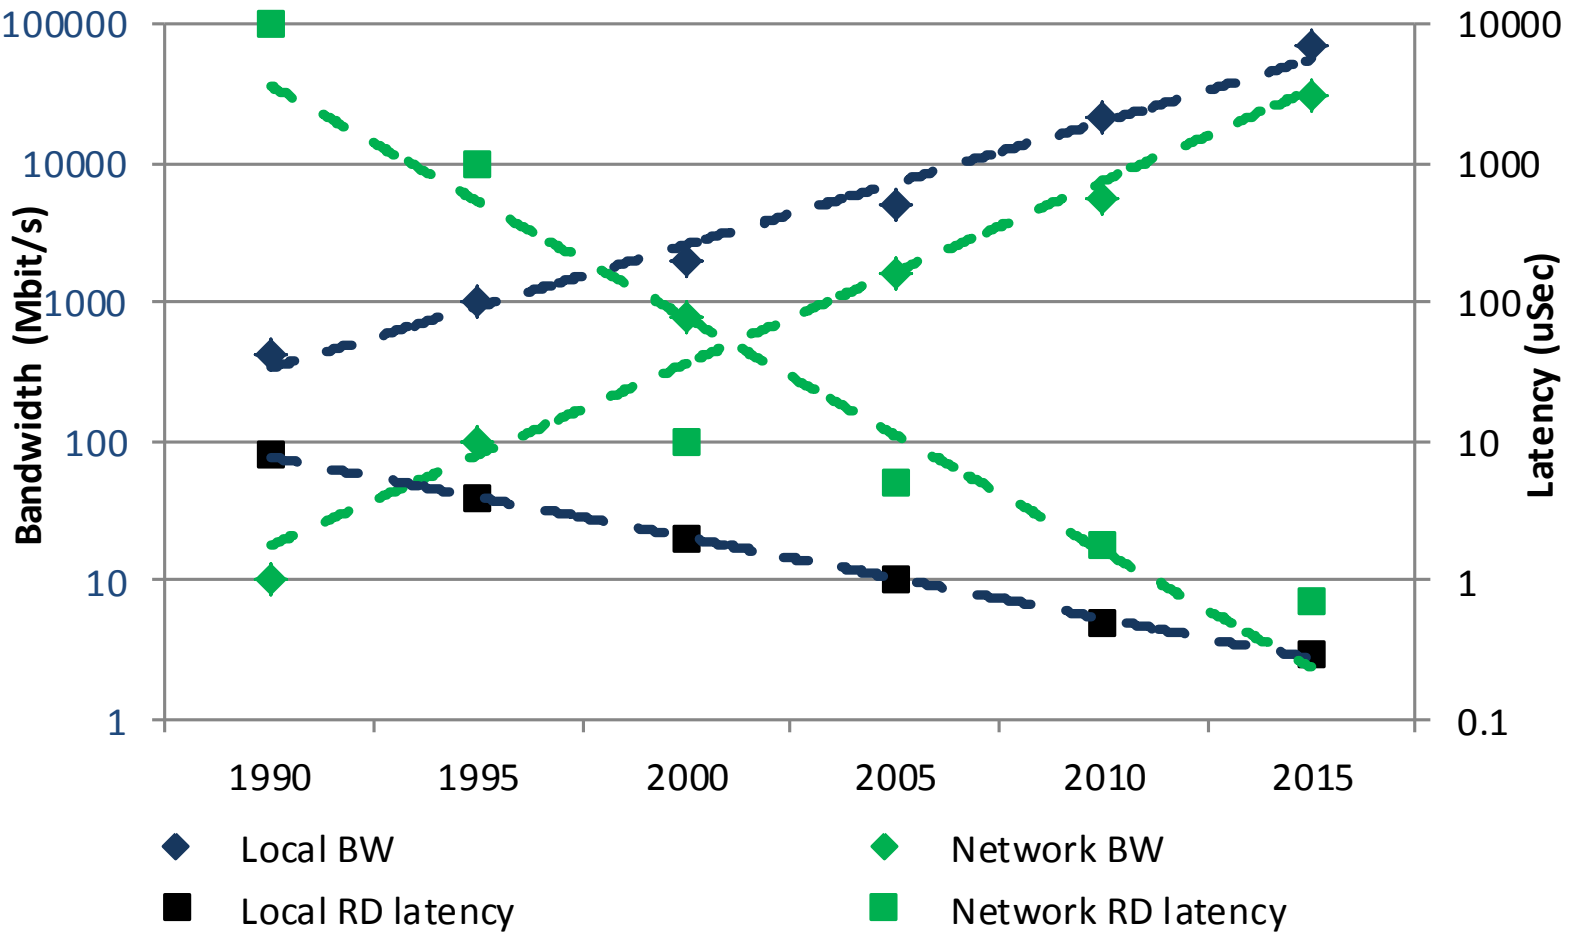

# Remote GPU access – rCUDA

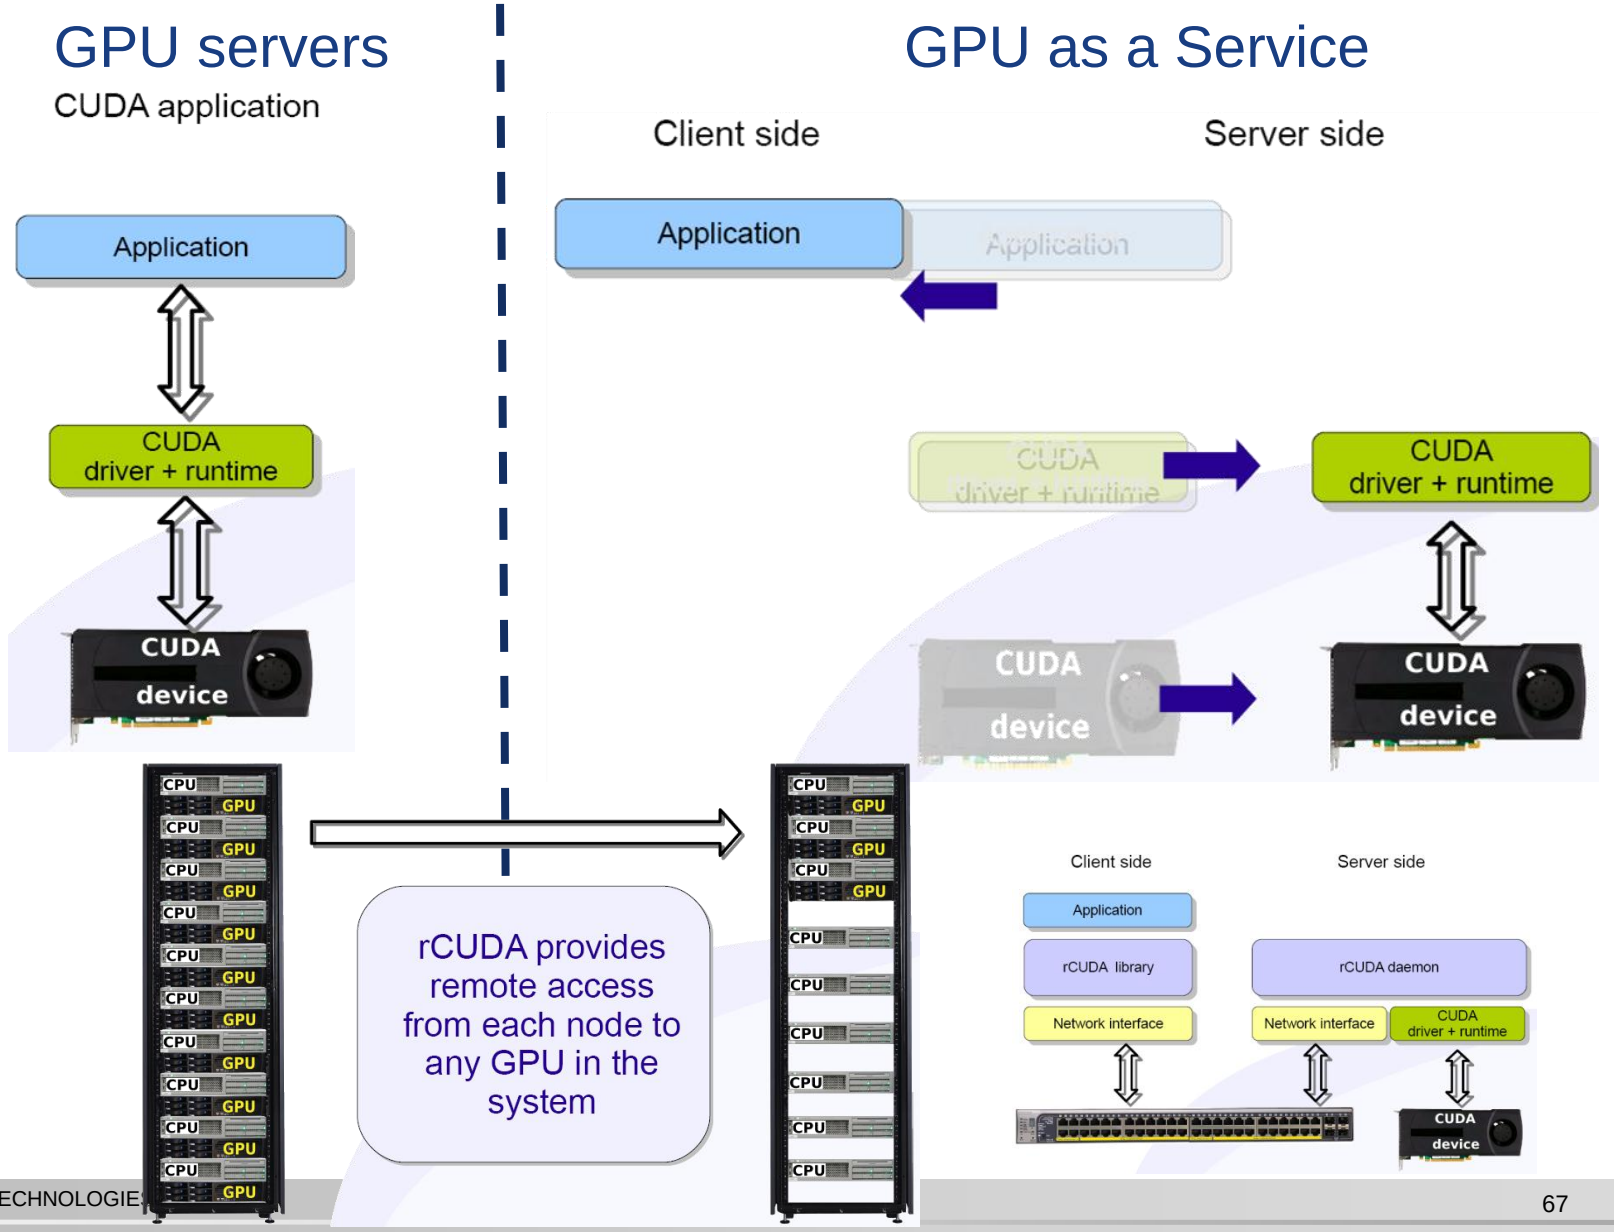

# GPU as a Service

## Servers with GPUs

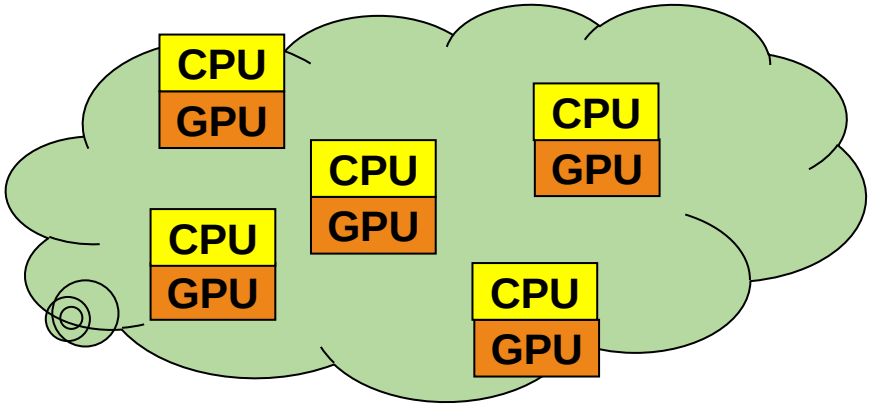


## GPU as a Service

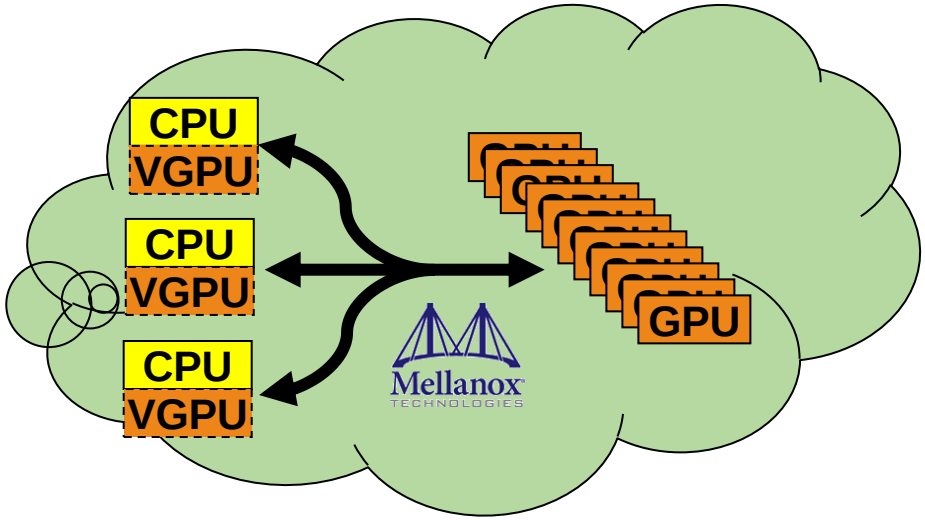


**PCIe-equivalent performance**

- **56Gb/s bandwidth**
- **0.7usec latency**

**RDMA dwarfs overhead**

- **Maintains local access model**
- **Supports memory management**

**Independent GPU management**

- **GPU as network-resident service**

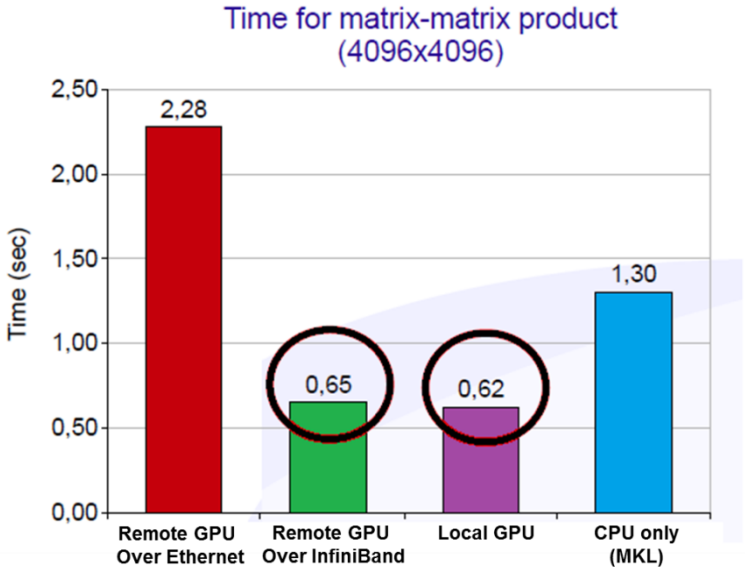

# Virtual Protocol Interconnect (VPI) Technology

## ConnectX®-3 VPI Adapter

redhat. Novell SUSE Microsoft vmware OPENFABRICS SOFTWARE

Applications

Networking | Storage | Clustering | Management

Acceleration Engines

PCI EXPRESS™ 3.0

Ethernet: 10/40 Gb/s

InfiniBand:10/20/40/ 56 Gb/s

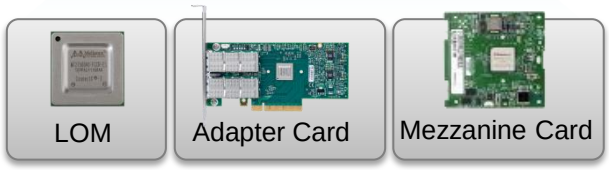


## SwitchX™ VPI Switch

Unified Fabric Manager

Switch OS Layer


64 ports 10GbE
36 ports 40GbE
48 10GbE + 12 40GbE
36 ports IB up to 56Gb/s
8 VPI subnets

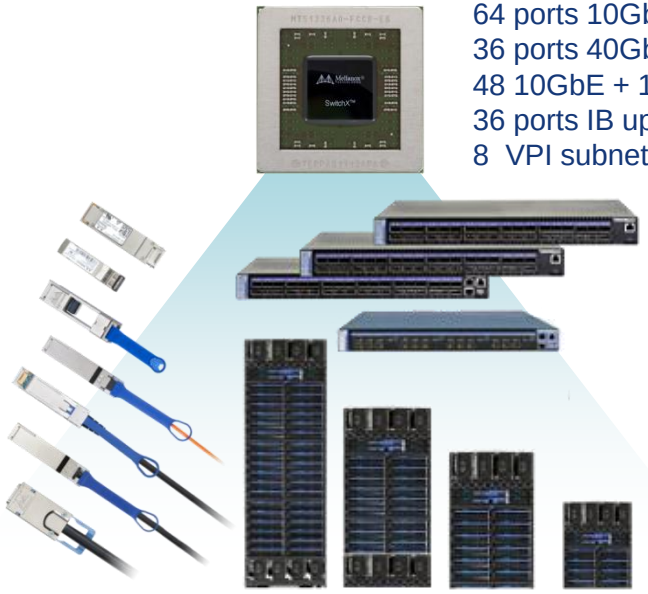

# Modular Building Blocks – Single Chip

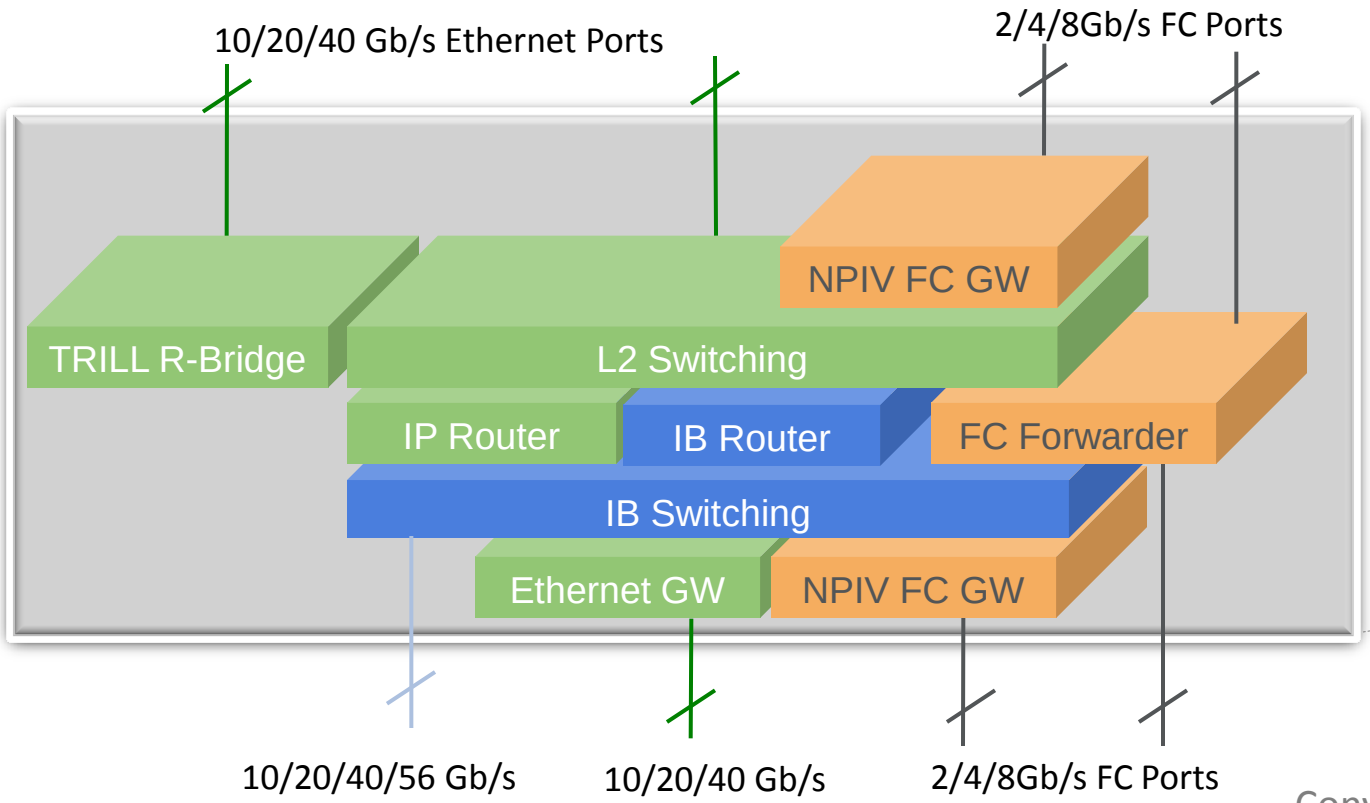



Converged I/O with Ethernet

Repurposing

Converged I/O with InfiniBand

# Industry's Only End-to-End FDR 56Gb/s InfiniBand Provider

- Highest performing interconnect to accelerate applications
  - Highest available bandwidth – 56Gb/s per port
  - Lowest end-to-end latency – <1us
  - End-to-end congestion control enables highest fabric efficiencies
- Superior application performance
  - Big Data, Web 2.0, cloud, database, financial service, storage, etc.
  - Automotive, oil & gas, database mining, financial analysis, etc.
- Best-in-class ROI
  - High server utilization and productivity
  - Ultra-low power

# Industry's Only End-to-End 10/40GbE Provider

- Superior Ethernet performance with unprecedented scalability
  - Fastest end-to-end Ethernet available
  - Lowest end-to-end latency
  - Backwards compatibility with 10GbE for infrastructure flexibility
- Required by multi-core CPUs, virtualization
  - Removes I/O bottlenecks that idles CPUs
  - Enables more full-performance VMs per server
- Best-in-class ROI
  - High server utilization and productivity
  - Leading low-power and density

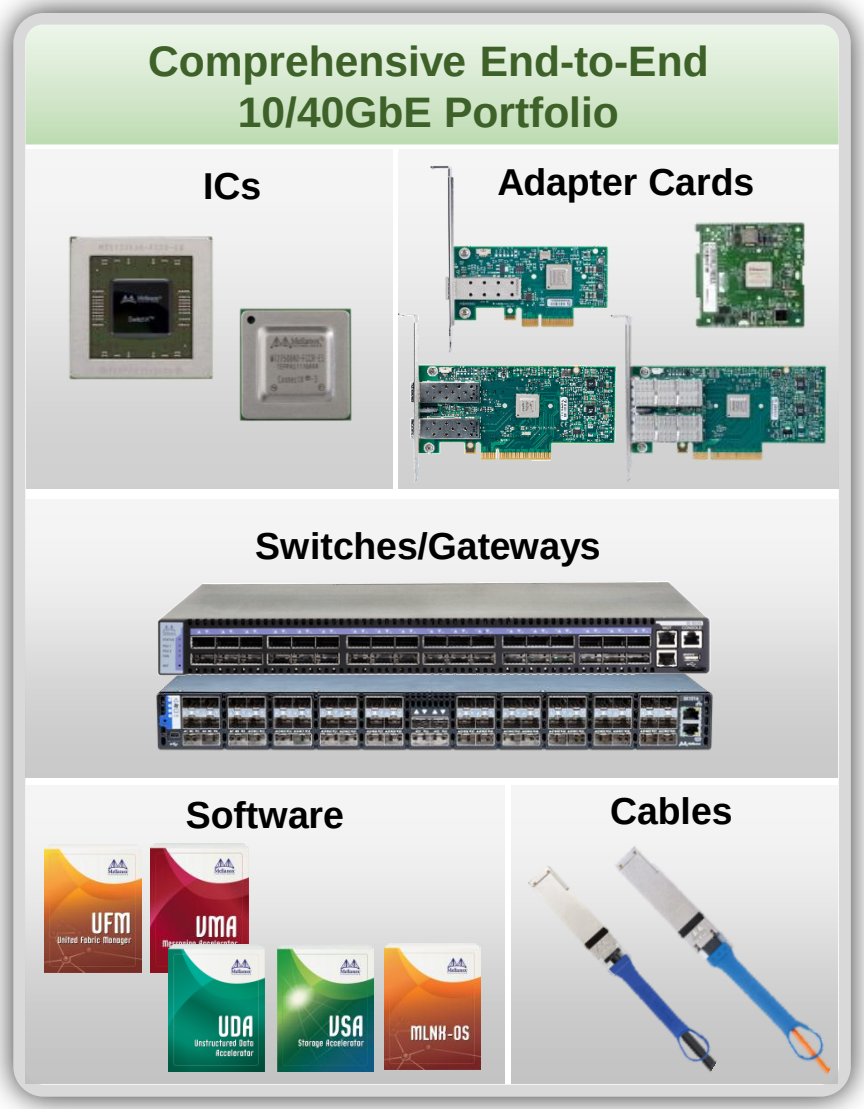

# Mellanox 10GbE Accelerated Penetration

- 10GbE NIC Market Leader
  - ~7X growth increase
  - 24.6% market share with 4% share over closest competitor
- Complete End-to-End
  - From silicon to systems
  - NICs, switches (new), cables, software
- Strengths in areas of increasing importance:
  - Low latency
  - CPU efficiency
  - High density
  - Low power

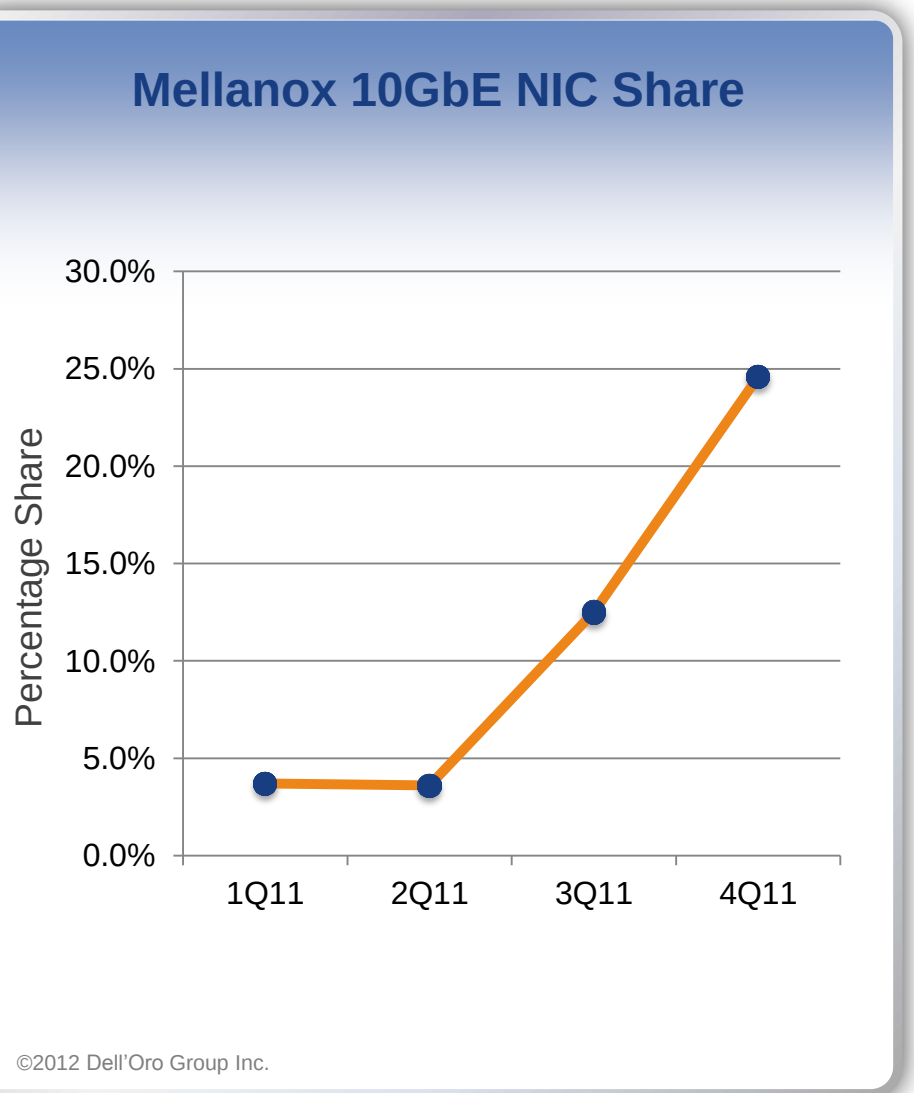

# Unmatched Efficiency & Scalability for Web 2.0

## *Big Data Analytics (e.g. Hadoop)*

- Run more jobs in parallel
- Faster rebalancing after node/link failure
- Faster data load into HDFS
- Ready for major increase in disks per node

**2X**


## *In-Memory Caching (e.g. Memcached)*

- More transactions per second
- Faster response/retrieval time
- More users/clients per server

**13X**

## *Price-Performance*

- More bandwidth at the same price
- Lower price per Gb/s
- Lower power per Gb/s

**4X**

Source: Company testing for order of magnitude improvement data

# Mellanox End-to-End Advantage

## Standard

- 3rd party cluster management
- Standard debuggers
- MPI collectives
- InfiniBand, TCP, UDP
- Eth, IB and FC L2 protocols
- Congestion control
- Link speed/width, FEC

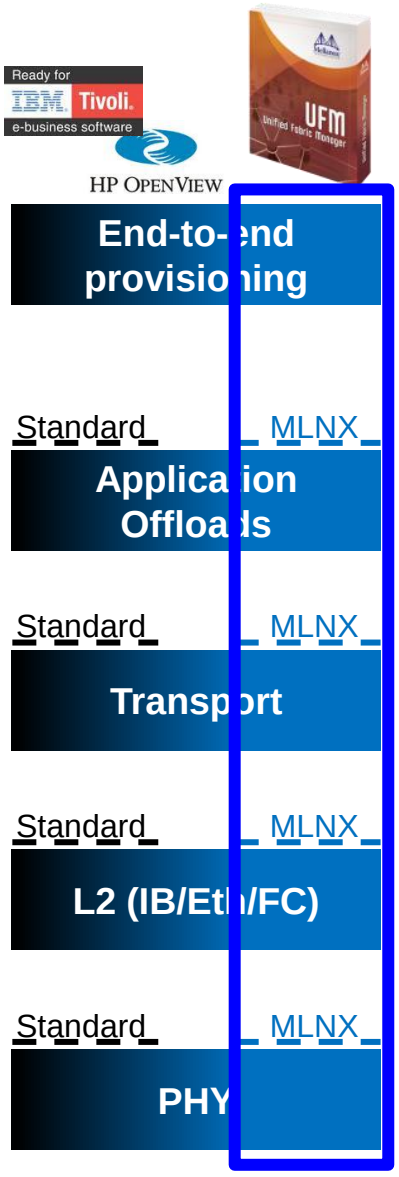


## Mellanox value-add

- End-to-end performance debug
- Paths' monitoring
- Datacenter-wide provisioning
- IOPs time stamping
- On-demand paging
- Ultra-scalable reliable transport
- QoS, network debuggers
- Adaptive L2 protocol choice
- Adaptive routing
- Unbreakable link technology

# The Cloud

# Cloud – the Service Provider

## Server-centric

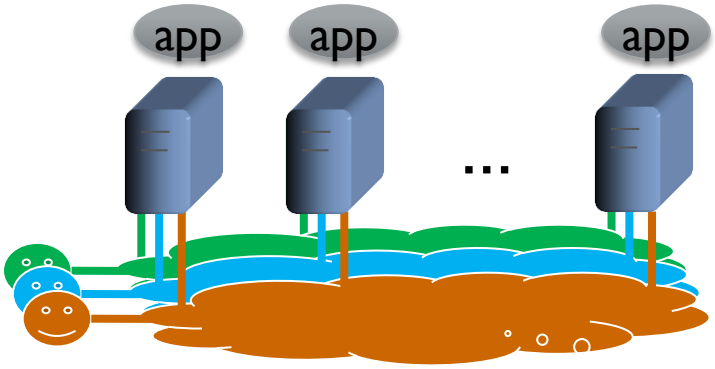


## Service-centric

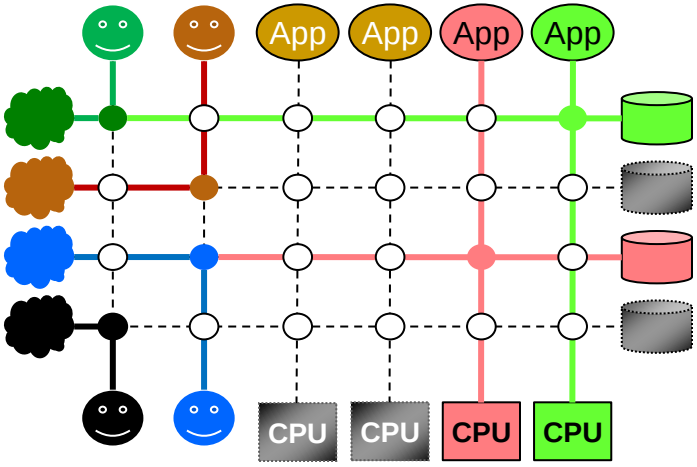

# Computing and Network as a Service

## Computing

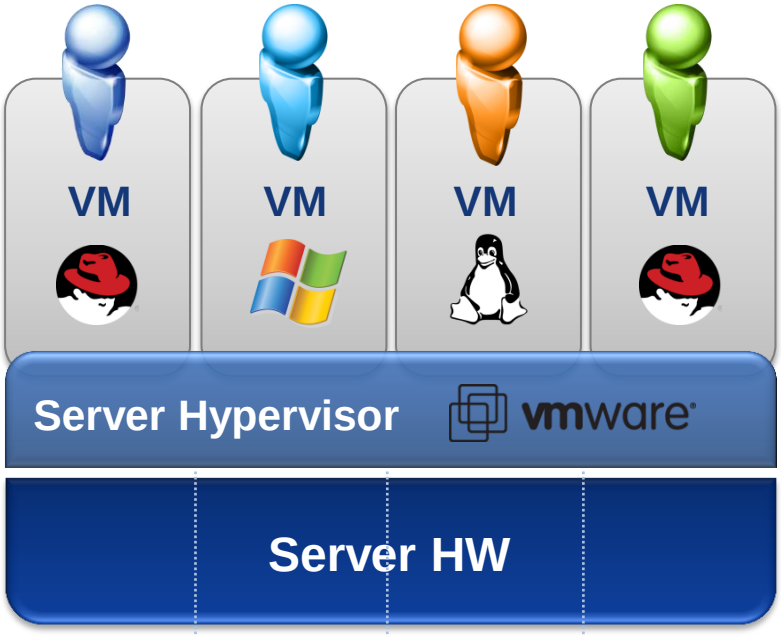


## Network

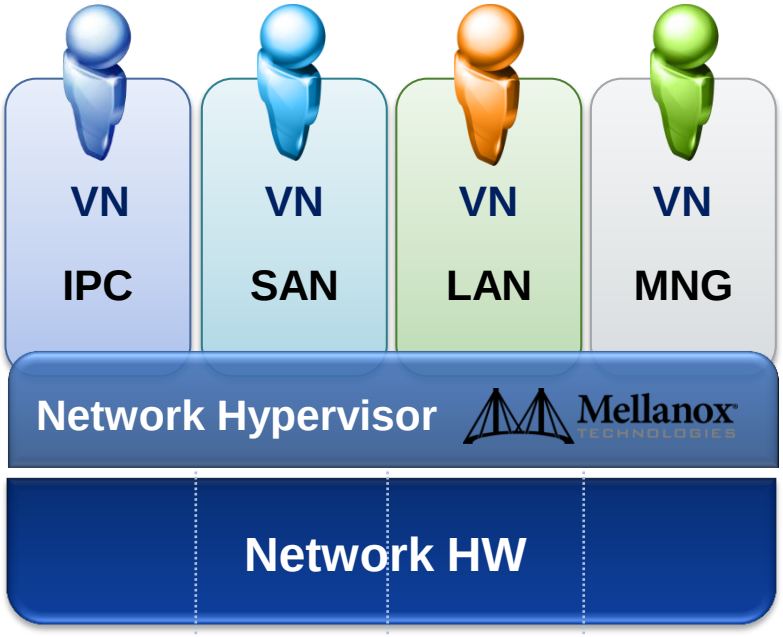


**Computing as a Service**

**Network as a Service**

# I/O as a Service

- **I/O adapter**
  - Best performance, cost, power
  - Virtualization

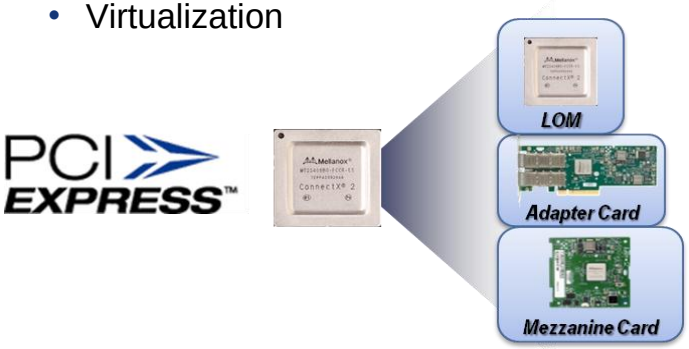


- **Network**
  - Best performance, cost power
  - Virtualization, Connectivity


- **Management**
  - End-to-end provisioning


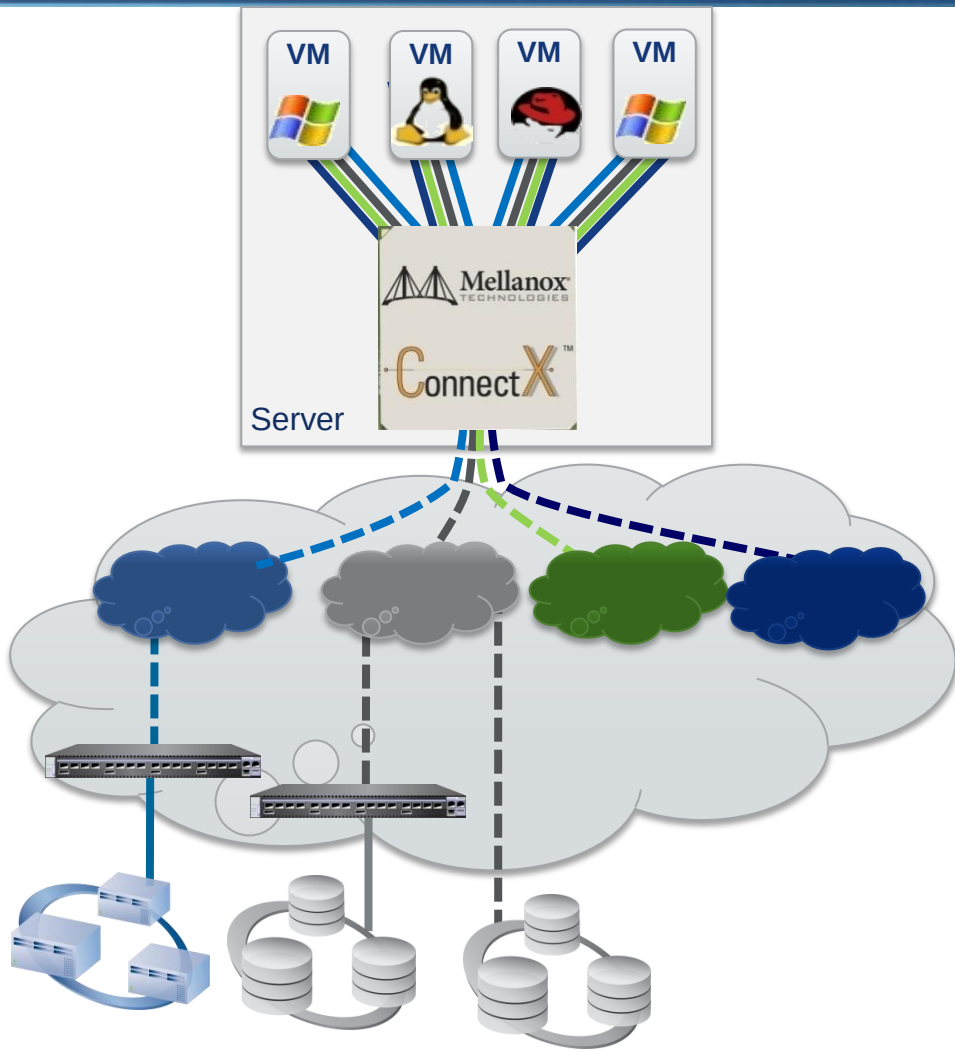

# Data Center Management

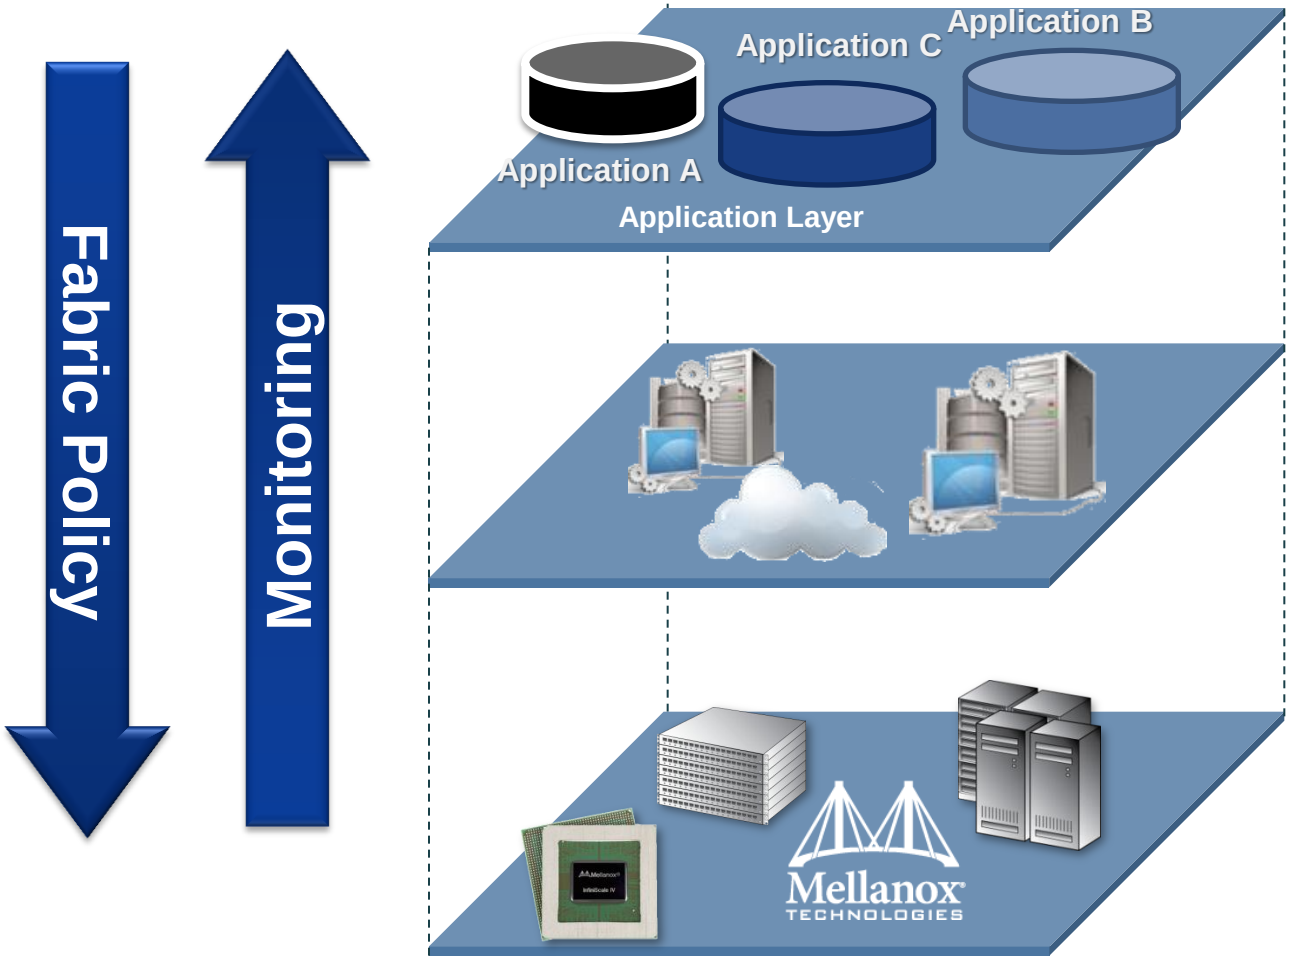


**Application Layer**

**Logical Layer**

**Physical & topology Layer**

# InfiniBand Cloud Appliance


**Mellanox 40Gb/s InfiniBand Accelerated**

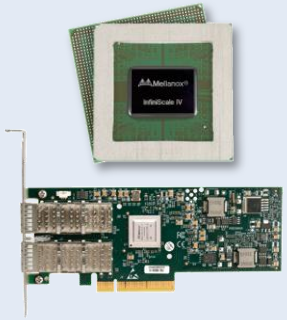

**Runs Oracle 10X Faster**

**The World's Fastest Database Machine**

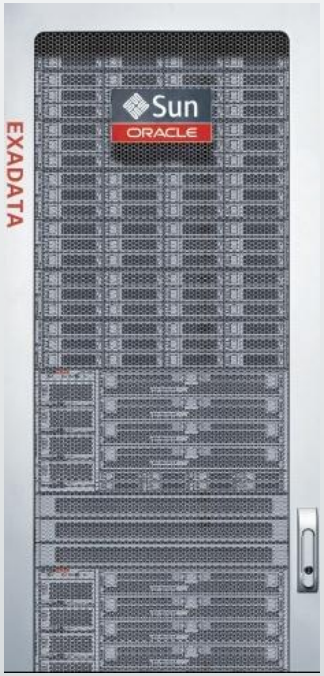

# Bing maps

Given its compute- and data-intensive mission, and the huge scale of its operations, the team requires enormous amounts of processing power, data storage capacity, and memory. It measures its storage needs in petabytes, its memory needs in terabytes, and its processing needs in tens of thousands of cores. At the same time, the team must work continuously to drive down the cost of processing this data.

they changed the server design to meet the unique storage and I/O needs of the BITs processing environment and they integrated InfiniBand networking infrastructure into the solution.

"In working with Dell, we were able to build a microsite that's cost-effective, reliable, agile, and rugged, that could be located in extreme environments within various countries around the globe," he notes.

'Better all the way around'

bing

"With the new approach to our site design, coupled with the Dell Modular Data Center solution, we project roughly an 8x reduction in cost when compared to traditional models."

*Brad Clark, group program manager, Bing Imagery Technologies*

## Technology at work

**Hardware**

- Dell Modular Data Centers
- AMD®-based Dell Nucleon servers
- InfiBand networking from Mellanox®

**Services**

- Infrastructure Consulting Services
- Custom Services and Integration
- Deployment

## Benefits

- Reduce power consumption
- Deliver fast throughput
- Accelerate IT deployment
- Create a model for future sites



View Dell case studies at: Dell.com/casestudies

# 10,000-Node Data Center Comparison

Mellanox TECHNOLOGIES

**3 Core switches**

**1,800 Cables**

**200 TOR switches**

- <5 Watts per node
- <2 Microsecond latency
- Small footprint, less cables
- 56Gb/s backbone
- 10Gb/s effective BW per Server

**Competitor**

**5 Core switches**

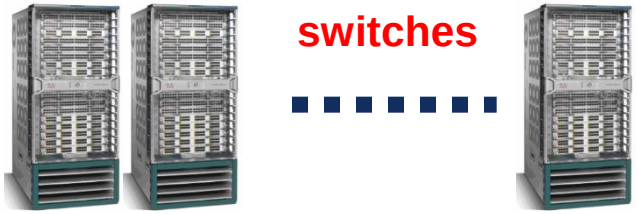

**3,350 Cables**

**209 TOR switches**

- >13 Watts per node
- >10 Microsecond latency
- Huge space, lots of cables
- 10Gb/s backbone
- 3.3Gb/s effective BW per server

**Mellanox Data Center Interconnect Solution – $8.7M Cost Saving!**

# Mellanox Fast Interconnect Drives Data Center ROI

| HPC | Web 2.0 | DB/Enterprise | Cloud | Financial Services |
|---|---|---|---|---|
|  |  |  |  |  |
| Up to 10X Performance and Simulation Runtime<br><br>33% Higher GPU Performance<br><br>Unlimited Scalability | 2X Hadoop Performance<br><br>13X Memcached Performance<br><br>4X Price/Perf | 10X Database Query Performance<br><br>4X Faster VM Migration<br><br>More VMs per Server and More Bandwidth per VM | 12X More Throughput<br><br>Support More Users at Higher Bandwidth<br><br>Improve and Guarantee SLAs | Lowest Latency<br><br>62% Better Execution Time<br><br>42% Faster Messages Per Second |

**Storage**

**Mellanox storage acceleration software provides >80% more IOPS (I/O operations per second)**

# Technology Leadership Marathon

# Run as fast as you can – and accelerate!

# Mellanox Track Record – Acceleration

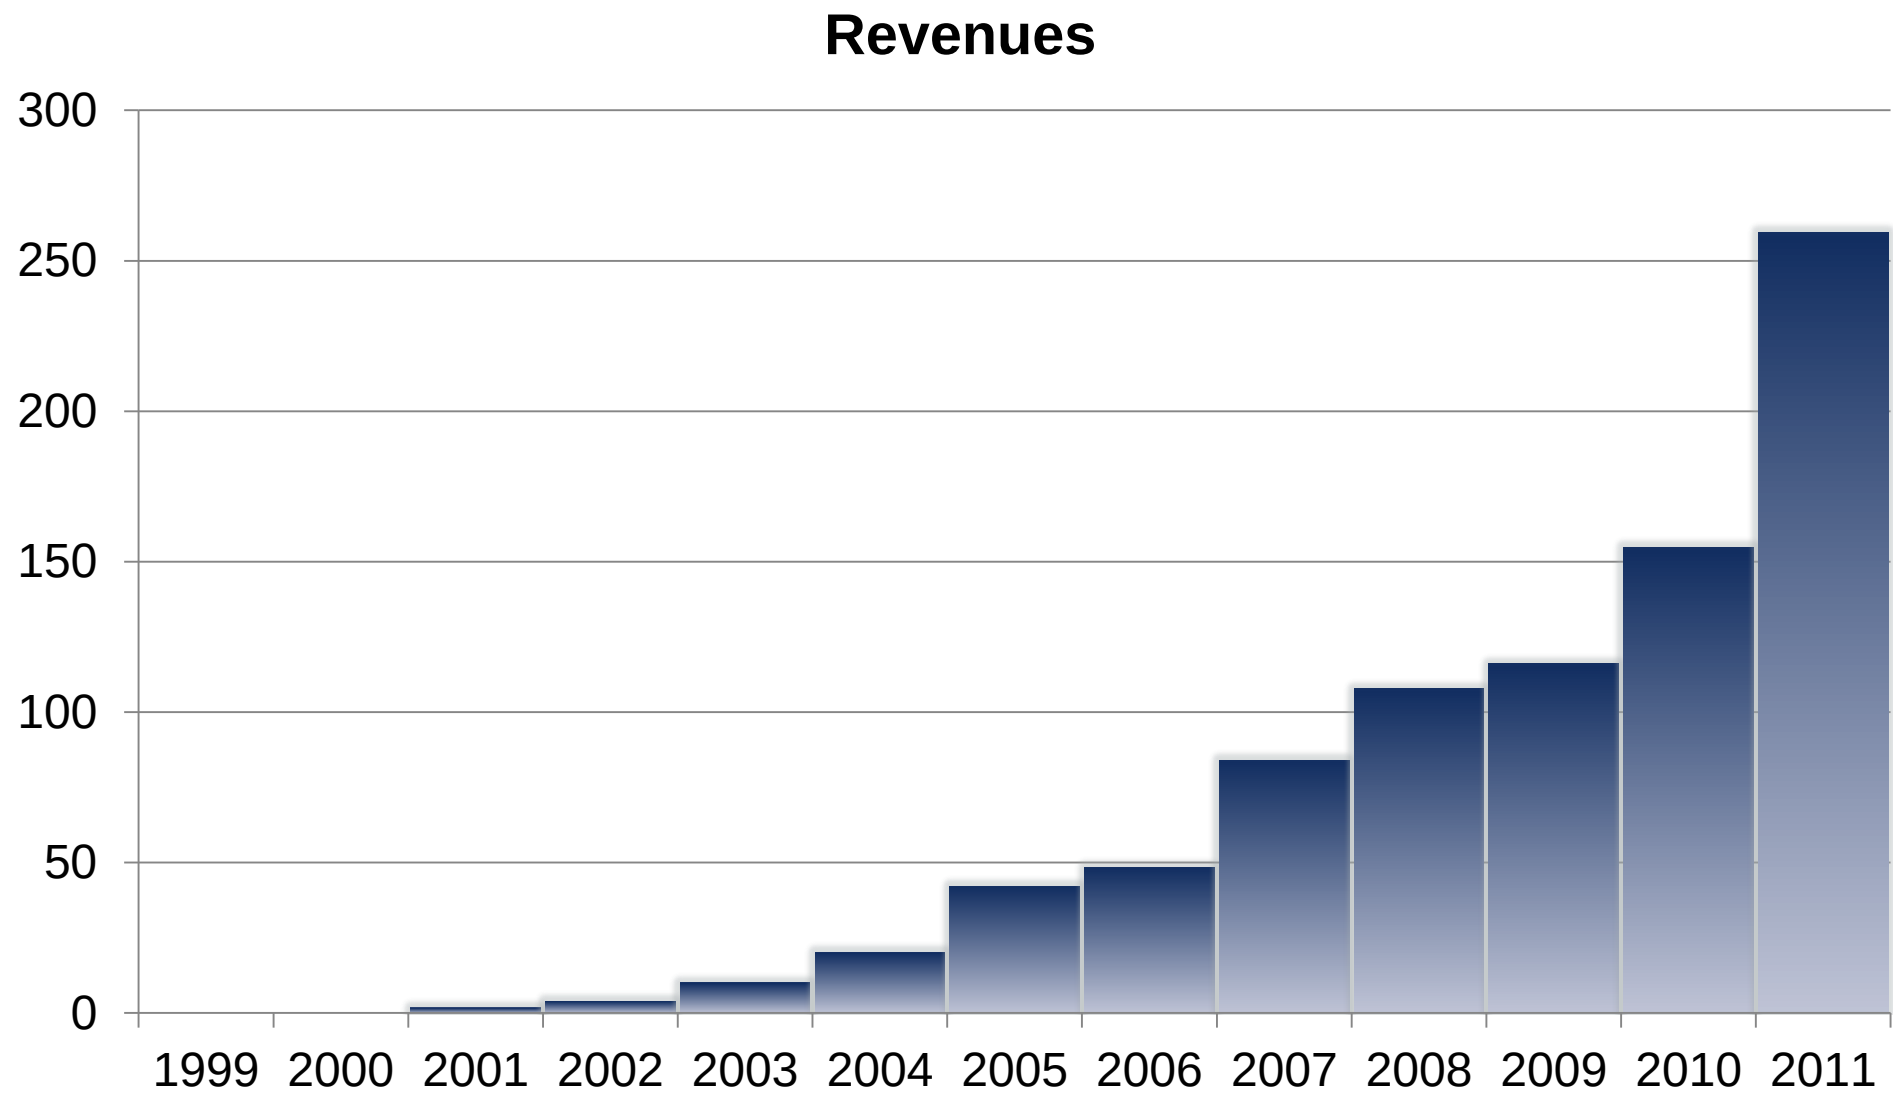

# Mission – Be the Leader

**Run as fast as you can – and accelerate!**

*Cheetah toady is a Turtle tomorrow*

# Mellanox Vision: Interconnect is Us

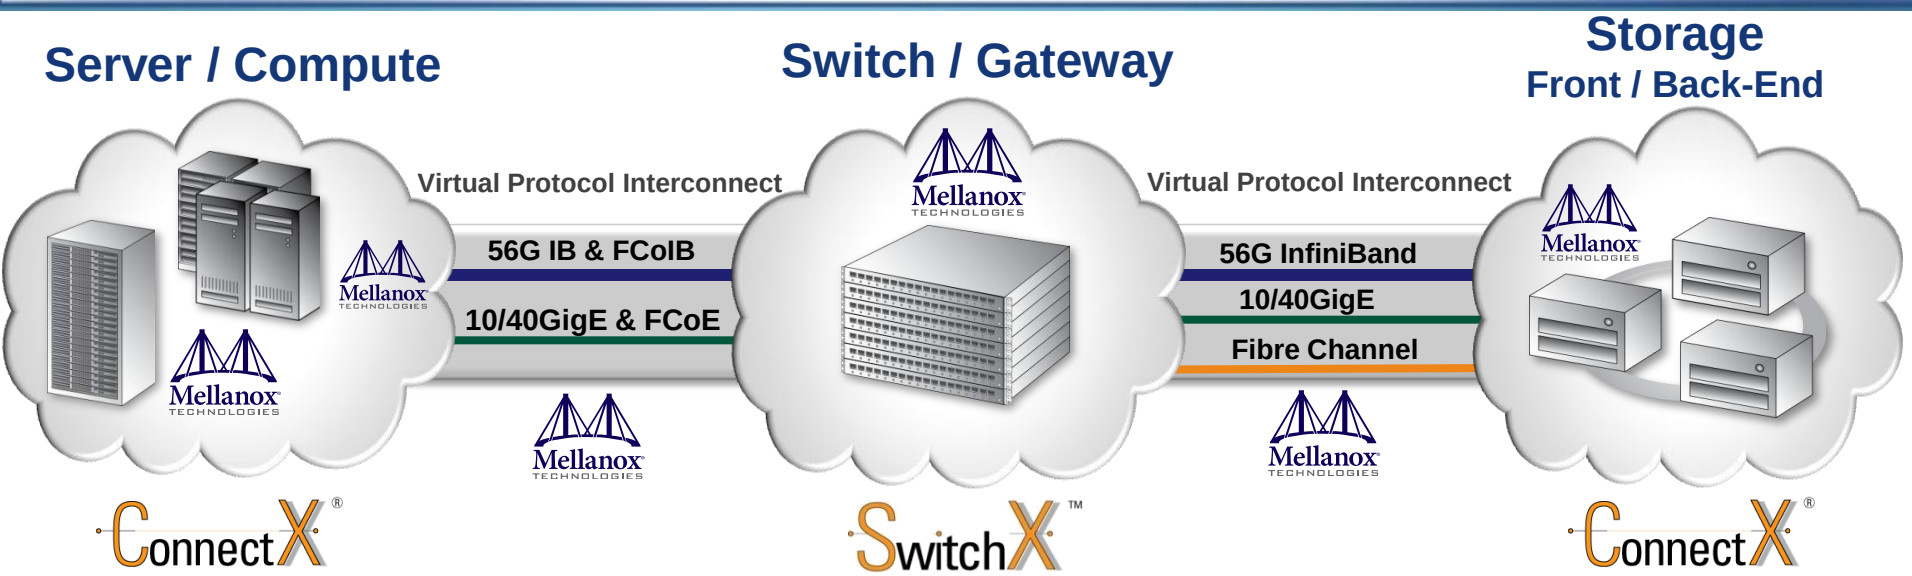


## Industries Only End-to-End InfiniBand and Ethernet Portfolio

| ICs | Adapter Cards | Host/Fabric Software | Switches/Gateways | Cables |
|---|---|---|---|---|
|  | 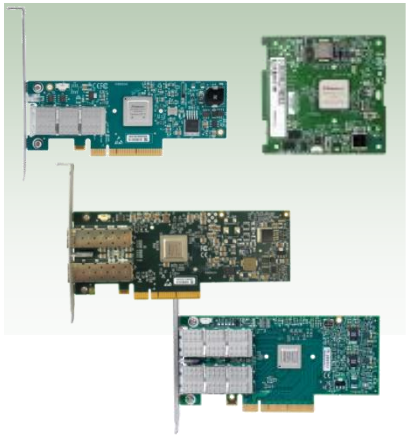 |  |  |  |

THANK YOU