Mellanox Connectivity Solutions for Efficient Cloud Infrastructure

Vishal Bharat

Country Manager

-

Providing superior server and storage performance, reliability, and scalability

Mellanox TECHNOLOGIES

Company Overview

Ticker: MLNX

- Leading connectivity solutions provider for servers and storage systems
  - Foundation for the world's most powerful and energy-efficient systems
- Company headquarters:
  - Yokneam, Israel; Sunnyvale, California
  - ~800 employees; worldwide sales & support
- Solid financial position
  - Record Revenue in FY'10; $154.6M
  - 1H'11 revenue = $118.4; up 55% Y-o-Y
  - Q3'11 guidance ~$67.5M to $68.0M
- Completed acquisition of Voltaire in Feb. 2011
  - End-to-End solutions for InfiniBand and Ethernet with enhanced software capabilities

Fortune 100 Penetration

10 of top 10 Stock Exchanges

10 of top 10 Automotive Manufacturers

5 of top 10 Pharmaceutical Companies

9 of top 10 Oil and Gas Companies

Leading Supplier of End-to-End Connectivity Solutions for Servers and Storage

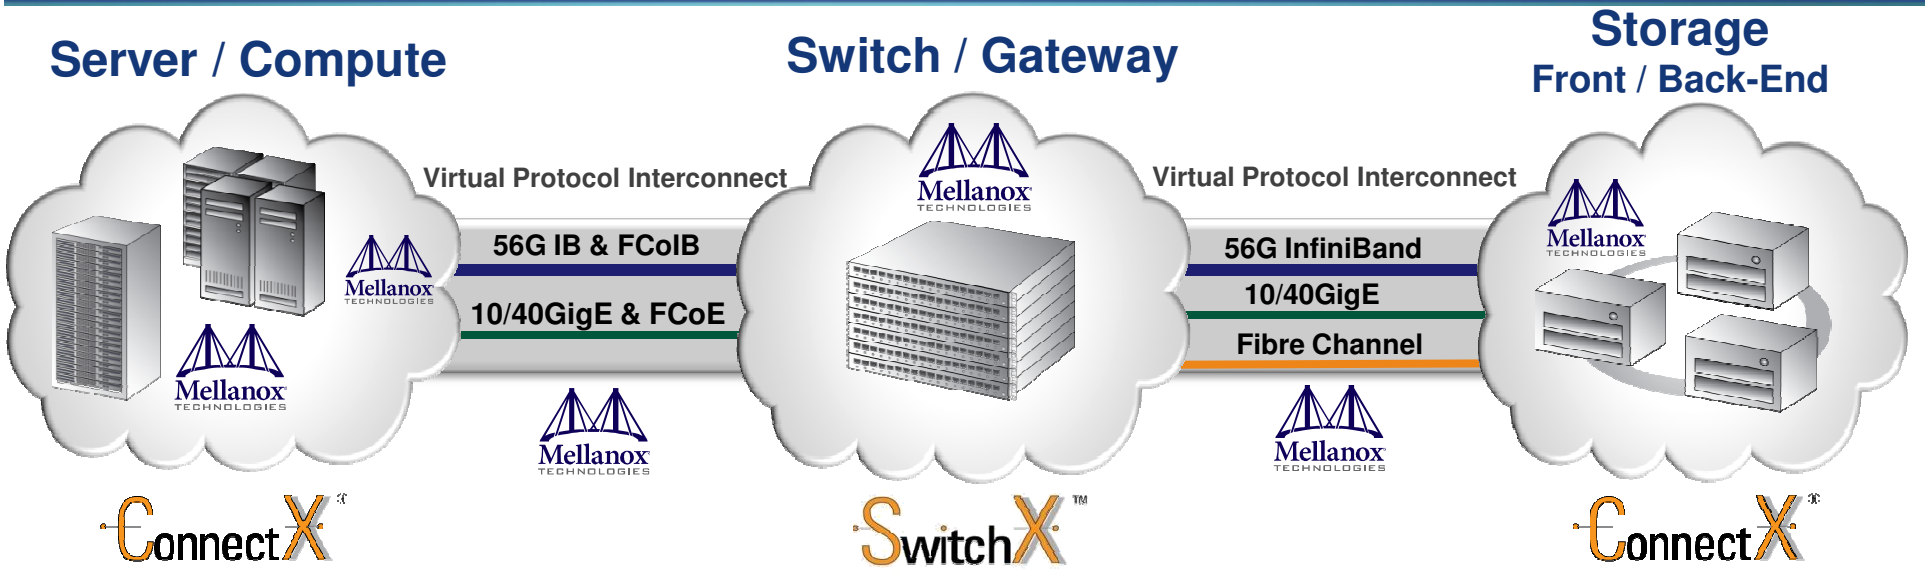

Comprehensive End-to-End InfiniBand and Ethernet Portfolio

ICs	Adapter Cards	Switches/Gateways	Host/Fabric Software	Cables
	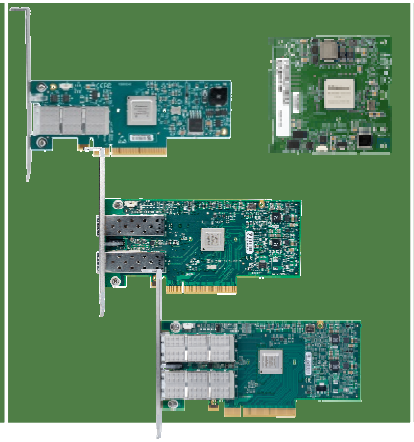	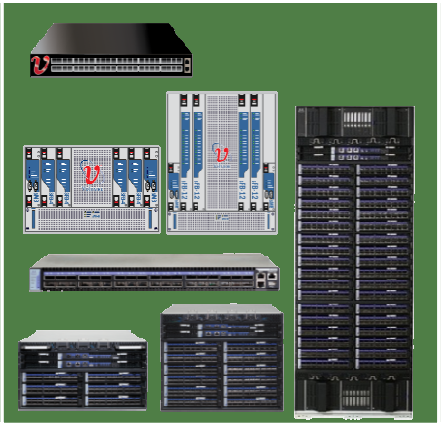		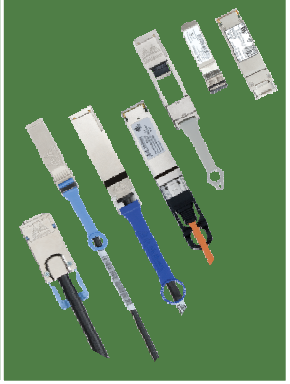

The Cloud - Service-Oriented View

- Hardware-centric view
  - Assign HW to consumer

- Application-centric view
  - Deliver service to consumer

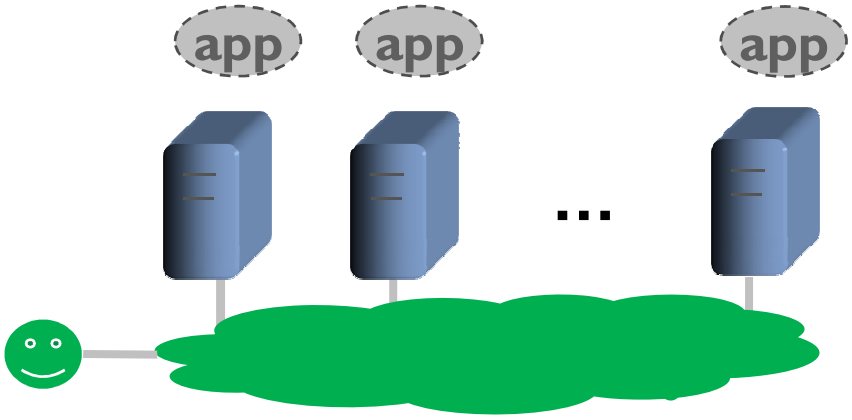

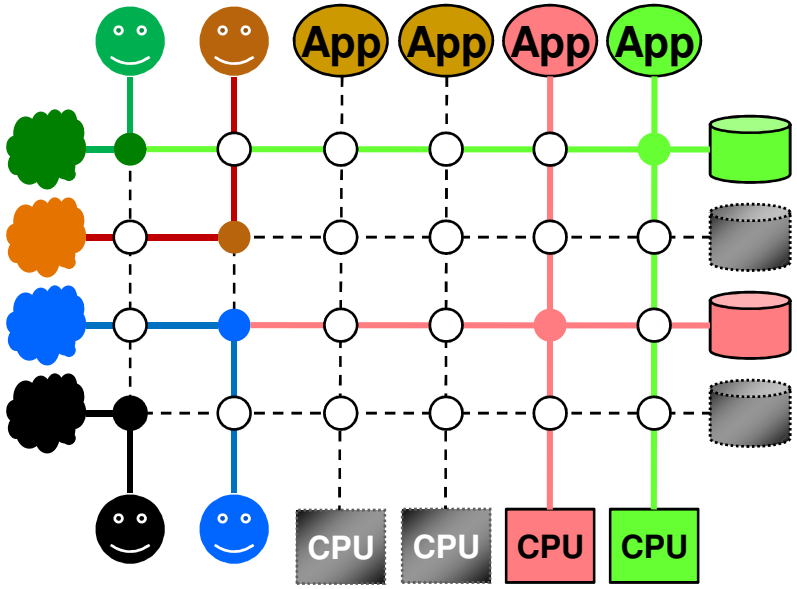

Apps Need Different I/O Service Levels

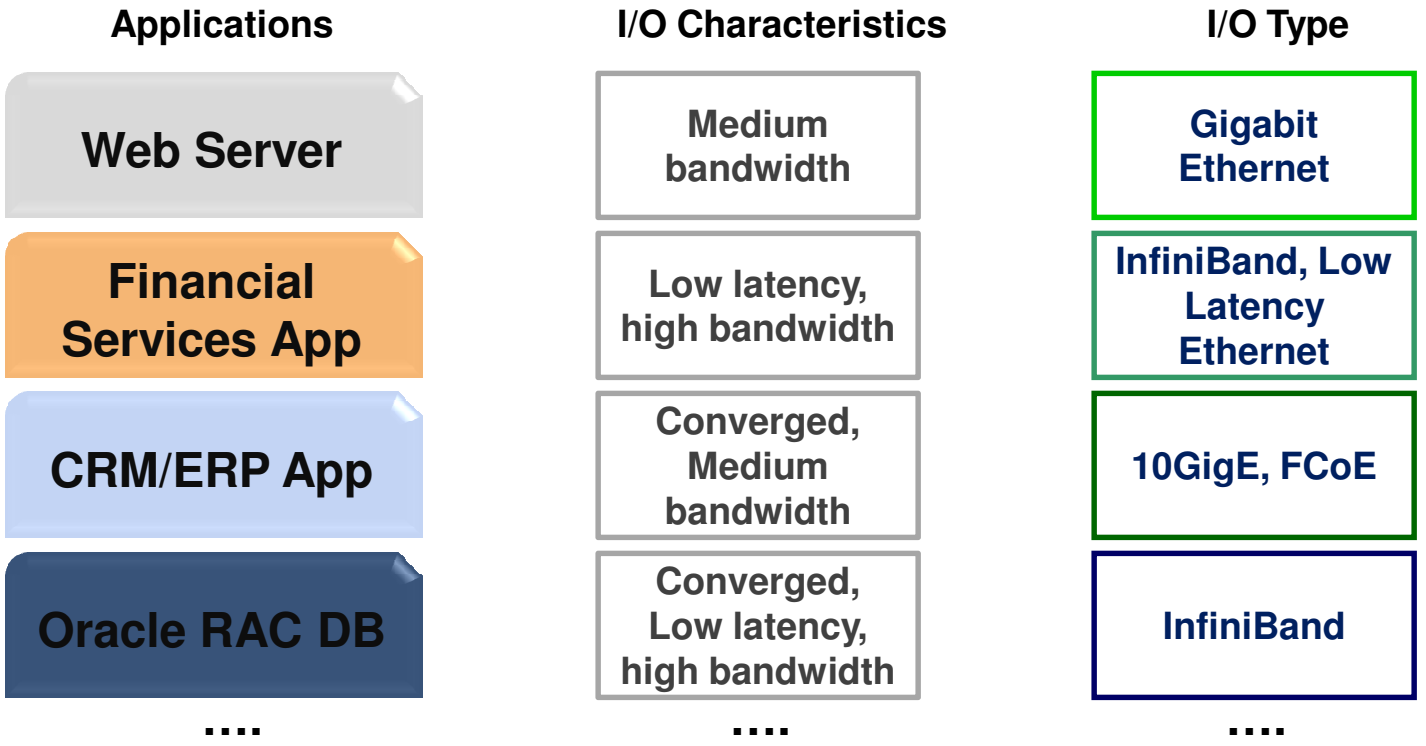

- Virtual I/O services must be adapted to VM apps
- Network agility = network services repurposing

Network agility must complement server agility

Network Agility Challenges

- Virtual I/O performance tied to network infrastructure
- Underlying network infrastructure is fixed

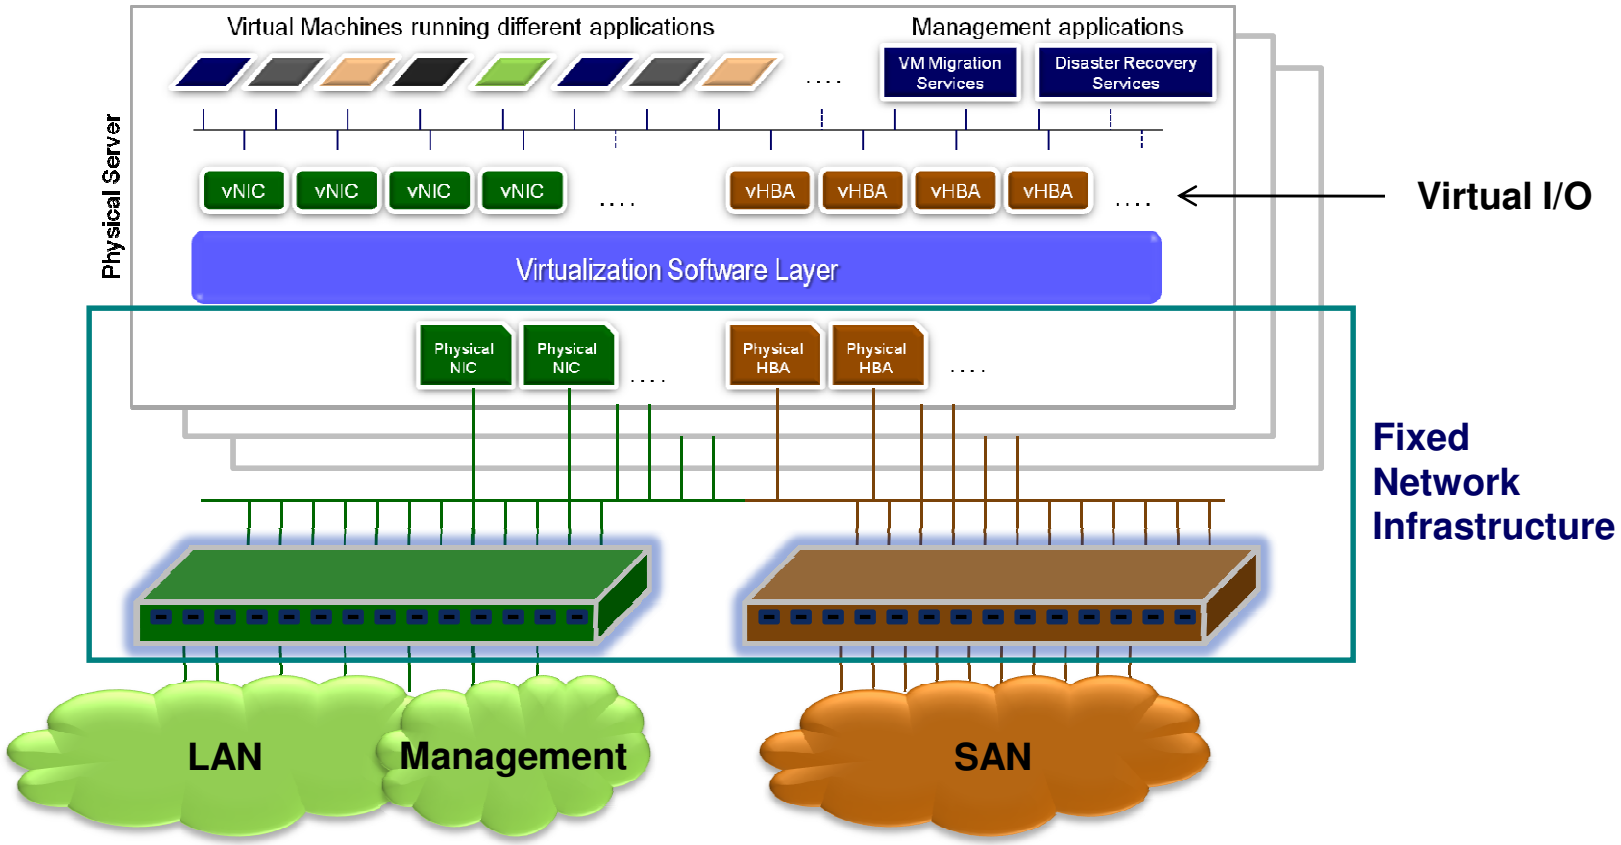

Need for on-the-fly network infra provisioning

Solution – IO Virtualization

- Physical I/O adapter ports repurposing
- Transparent virtual I/O for VMs and apps

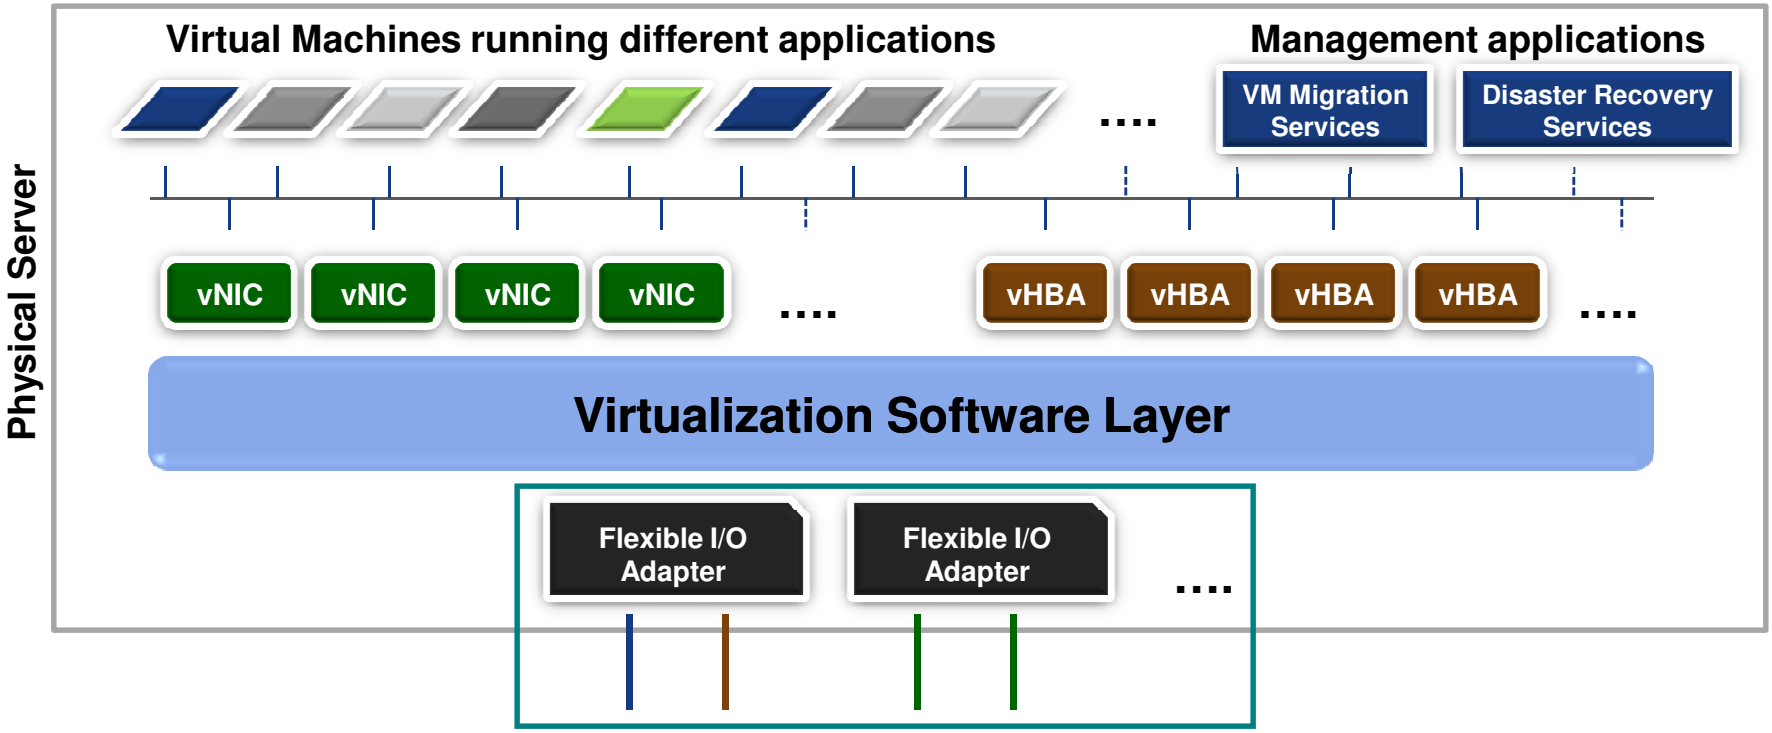

Physical ports can be repurposed to deliver different I/O characteristics and type

On-demand network adapter services for VM apps

Mellanox Multi-Protocol/VPI Connectivity Technology

Efficient, Flexible and Scalable for Maximum ROI

Financial | Clustered Database | Cloud Computing | Web 2.0 | Storage


Trading, Analytics

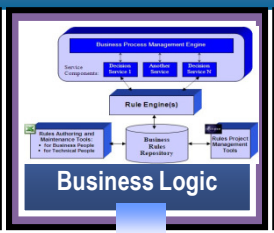

Business Logic

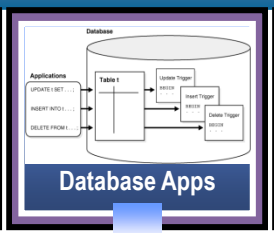

Database Apps


CRM, ERP Apps

.NET, JAVA


Web Services

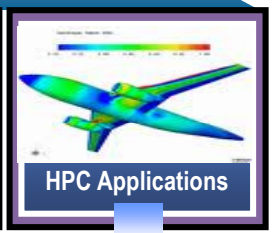
HPC Applications

Mellanox VPI Connectivity Solution

Applications Transparency enables Data Center Agility

Mellanox's Unique Value Proposition: 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

10/40Gb/s, FCoE

InfiniBand: 10/20/40/56, FCoIB, EoIB


LOM

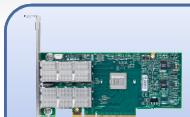
Adapter Card


Mezzanine Card

SwitchX™ VPI Switch

Unified Fabric Manager

Switch OS Layer

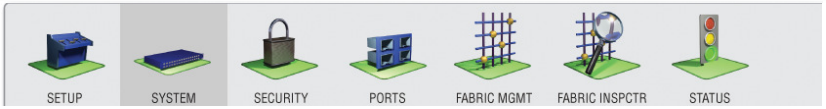

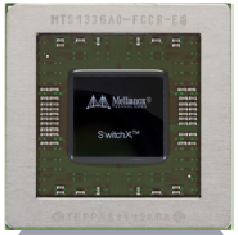

- 36 ports 40/56Gb/s
- 36 ports IB up to 56Gb/s
- FCoE, FCoIB
- 8 subnets

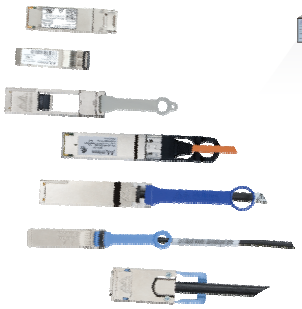

Industry's Only End-to-End 40GbE Provider

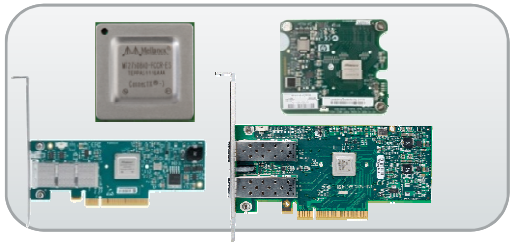

- Industry's highest performing Ethernet NIC
  - 10/40GbE w/FCoE with hardware offload
  - Ethernet industry's lowest end-to-end latency
  - Faster application completion, better server utilization

- Increasing ecosystem support momentum
  - Multiple Tier-1 OEM design wins (Dell, HP, IBM)
    - Servers, Blade Servers, and LAN on Motherboard (LOM)
  - Comprehensive OS Support
    - VMware, Citrix, Windows, Linux

- High density, low latency 10/40GbE switches
  - Up to 288 10GbE ports
  - Ultra-dense – 64 ports 10GbE or 36 ports 40GbE in a 1U system
  - Industry's lowest latency (230ns) and power per port
  - Sold through Tier-1 OEMs (HP, IBM)
  - Consolidation over shared fabrics

- Integrated, complete management offering
  - Service Oriented Infrastructure Management, with Open APIs

10/40Gb Ethernet Leverages Large Installed Technology Base

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

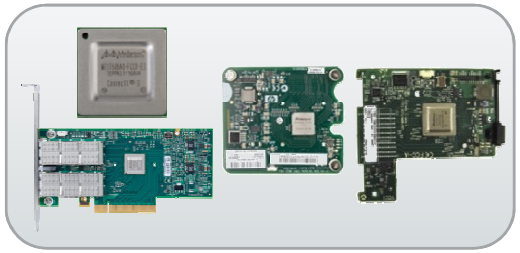

- Adapter market and performance leadership
  - First to market with FDR 56Gb/s adapters
  - Delivers next-gen application efficiency capabilities
  - Global Tier-1 server and storage availability
    - Bull, Dell, Fujitsu, HP, IBM, Oracle, SGI, Sugon, T-Platforms

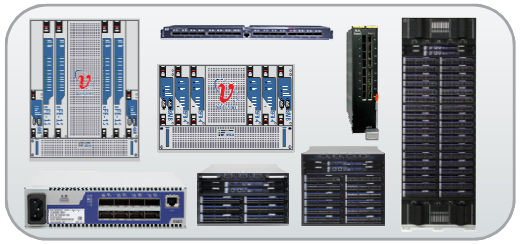

- Comprehensive, performance-leading switch family
  - First to market with FDR 56Gb/s switches
  - Industry's highest density and scalability
  - World's lowest port-to-port latency

- Comprehensive and feature-rich management/acceleration software
  - Enhancing application performance and network ease-of-use

- High-performance converged I/O gateways
  - Optimal scaling, consolidation, energy efficiency
  - Lowers space and power; increases application performance

- Copper and Fiber Cables
  - Exceeds IBTA mechanical & electrical standards
  - Ultimate reliability and signal integrity

InfiniBand is the Ideal Solution for Large, Bandwidth Intensive, Low Latency Applications

VM Migration Acceleration

Faster VM Migration with VMware vMotion over ConnectX-2 EN 40GbE NIC

2X Boost

ConnectX-2

3.5X Boost

Migration of Passive VM

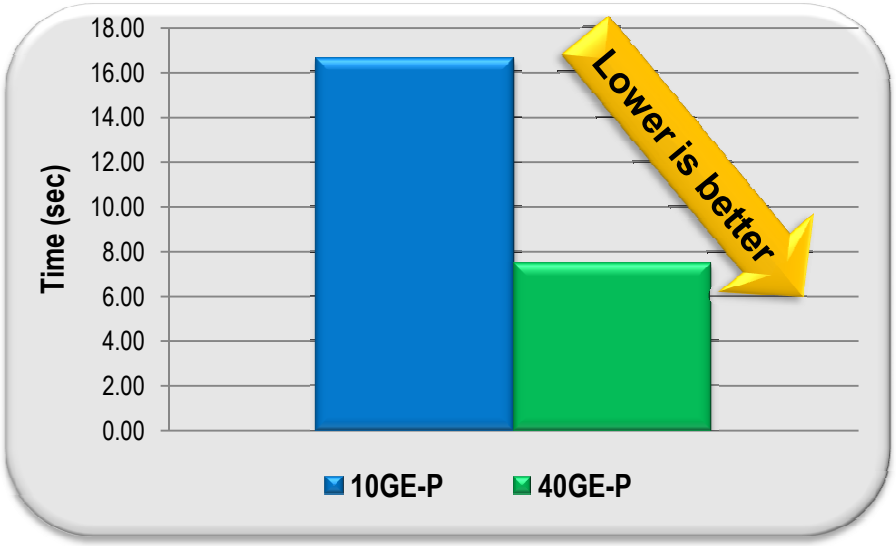

Migration of Active VM

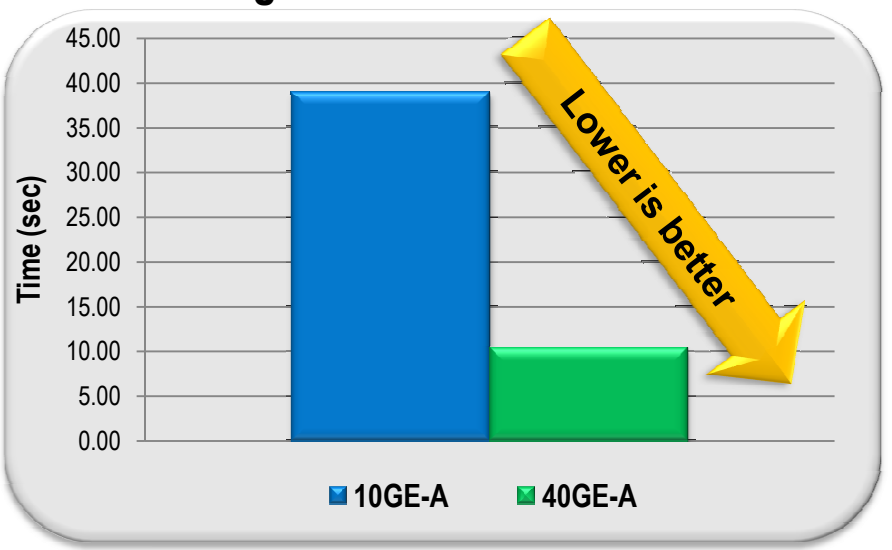

350% SLA boost in VM migration over 40GbE

SLA/ROI* Benefits with VM Migration Acceleration

Boosted SLA at a Lower TCO

32 VMs
4 x (8) Cores Servers
10GigE Connectivity

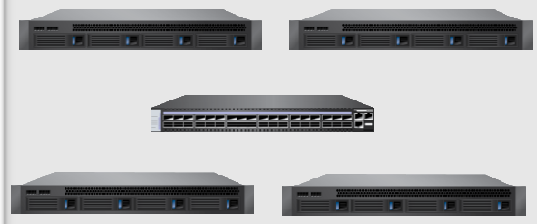

4 x 8 Cores Server
4 x 10GE Ports
4 x 10GE Switch Ports
4 x Cables
$15,800 Total Capital Cost
$987.50 Capital Cost per VM
$528.60 Power Cost per VM (3Y)

32 VMs
2 x (16 Cores) Servers
40GigE Connectivity

2 x 16 cores Server
2 x 40GE Ports
2 x 40GE Switch Ports
2 x Cables
$10,700 Total Capital Cost
$668.75 Capital Cost per VM
$295.65 Power Cost per VM (3Y)

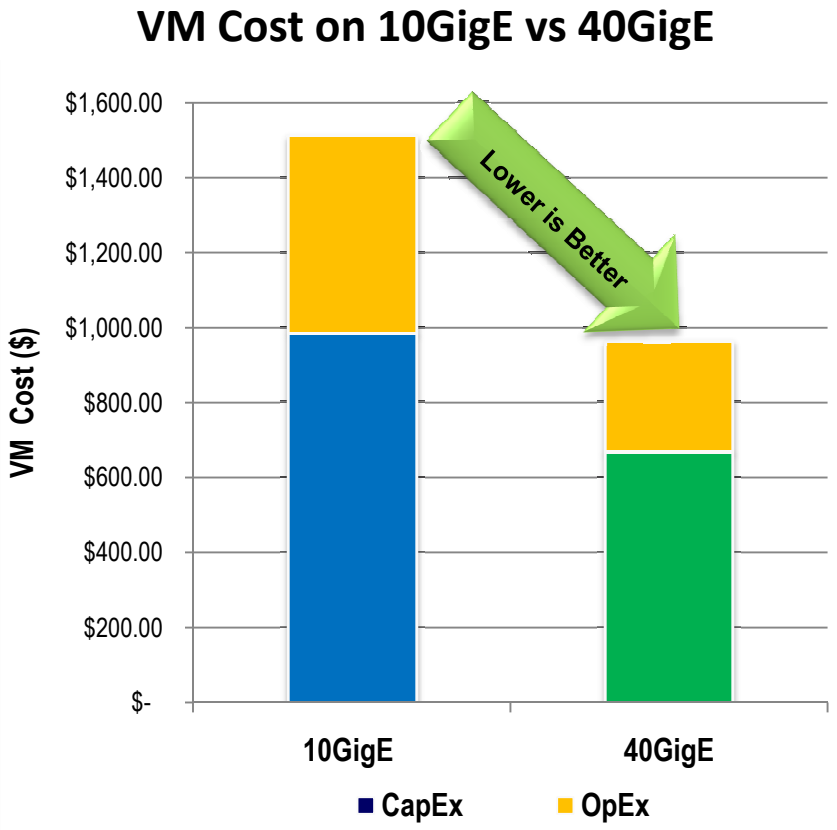

* Single VM Migration with 16GB memory

70%+ Operational cost saving over true 40GigE fabric

InfiniBand – The Most Scalable Server Fabric Technology

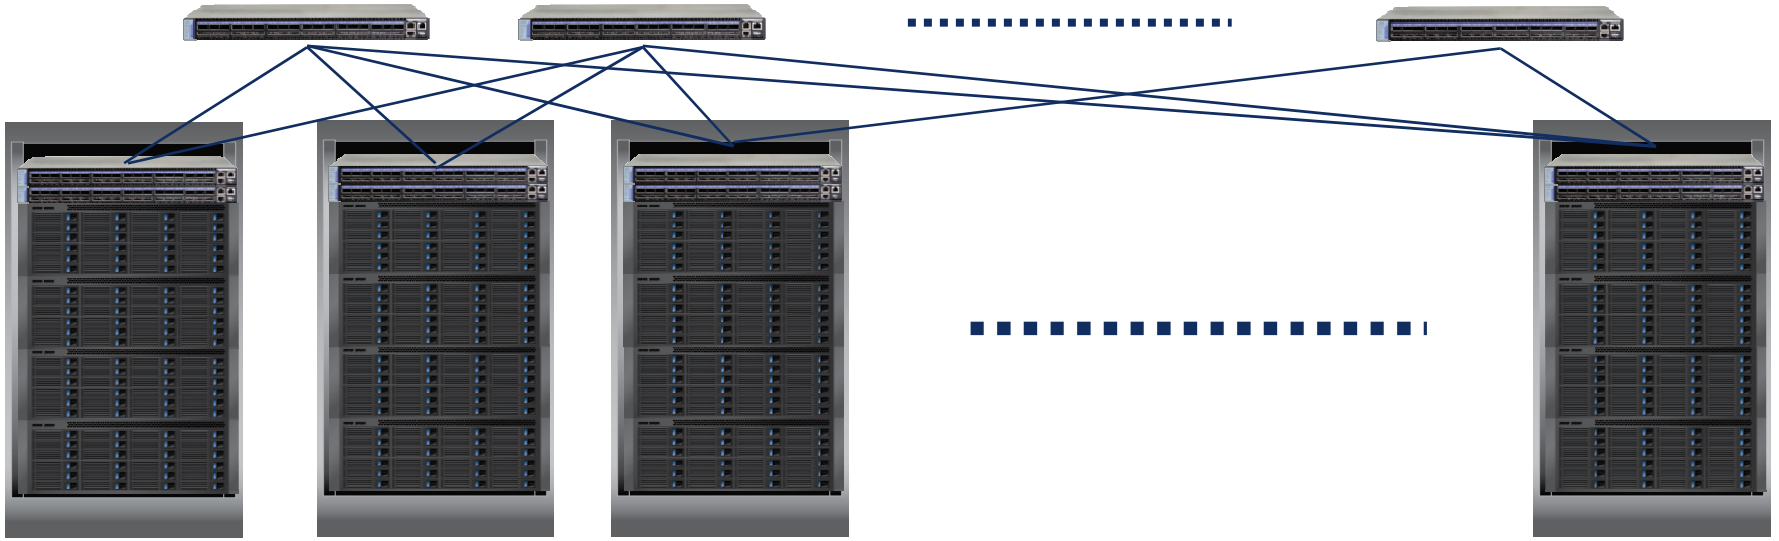

- Scalable ***large L2*** fabric – no broadcast domain
- Robust, standard ***multipathing***
- ***Highest throughput*** for traffic peaks and replication
- Robust, standard ***congestion management***
- Best ***price-performance*** – see next slide for details

The Mellanox Advantage: Superior Efficiency & Scalability

1000 Nodes 40Gb/s IB vs. 10Gb/s EN in the Cloud

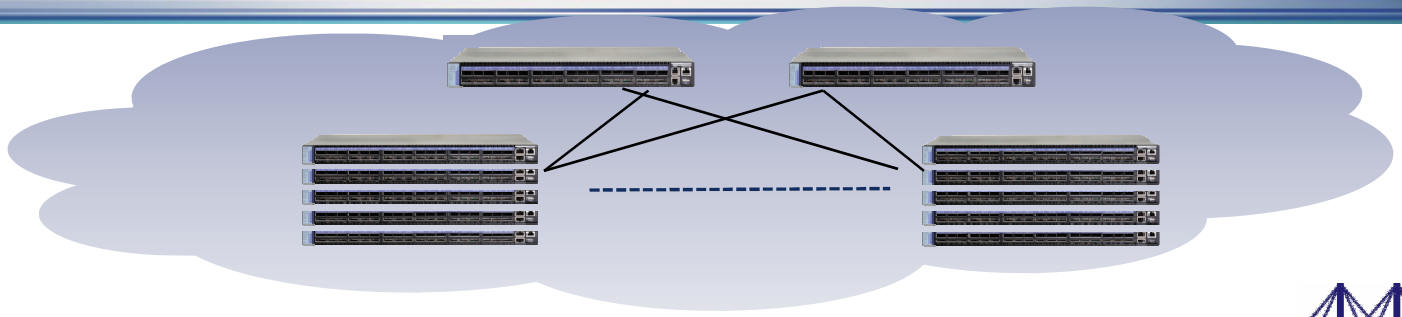

	Vendor A	Vendor B	IS5024
Protocol	10GbE	10GbE	IB QDR (40Gb/s)
Network Configuration	1 : 1 (24:24)	1 : 1 (32:32)	4 : 1 (28:8)
Tier1 + Tier2 ToR switches	42 + 21	32 + 16	36 + 8
Cost per ToR / NIC *	$22K / $460 (Low cost NIC)	$21K / $460 (Low cost NIC)	$6.1K / $560 (MHQH19)
Cost per Cable (2m copper to host /10m optic up)	$49 / $700 (2 * Cisco 10G SFP+ SR)	$49 / $700 (2*Cisco compatible SFP+ SR)	$58 / $300
$ Price per Node (NIC + Cbl + Switches)	$2.6K	$2.0K	$1.0K ✓
Throughput per Node (Avg/Peak)	10 Gbps	10 Gbps	10/40 Gbps
Power Per Tor / NIC	390W / 6.7W	250W ** / 6.7W	111W / 5W
Power Per Node	31W	19W	10W ✓
Scalability	No	No	Yes

The Mellanox Advantage: Up to 40-50% price-performance improvement

*** Street Price**
**** Published Number is 130W at 50% load**

Exalogic Elastic Cloud – What it is?

- Fastest Java performance
  - Application performance improved up to 10X
- Foundation for Mission Critical Cloud
  - Best for application consolidation
  - Best for elastic capacity on demand
- Standard hardware configuration
  - Servers and Storage connected by Mellanox QDR
  - InfiniBand to Ethernet gateway using BridgeX
- Standard Software Configuration
  - OS, Databases, Middleware, Applications
  - Oracle Virtual Machine (OVM) with Mellanox SRIOV technology over InfiniBand HCA
- Best Performance and Cost Performance
  - Lowest cost to deploy and operate
  - Highest performance

Oracle Exalogic Elastic Cloud

- 40Gb/s InfiniBand-enabled Exalogic delivers fast application runtime!
  - Accelerate the performance of Java applications by as much as 14x
  - Improve reliability and scalability beyond even the most mission-critical requirements
  - Reduce deployment effort by up to 95% and reduce costs by as much as 60%

- Ideal for mission-critical private cloud
  - Enables robust isolation for processes and near-instant failover of cluster nodes
  - Allows consolidation of multiple, unrelated applications
  - Extremely high reliability and elastic capacity

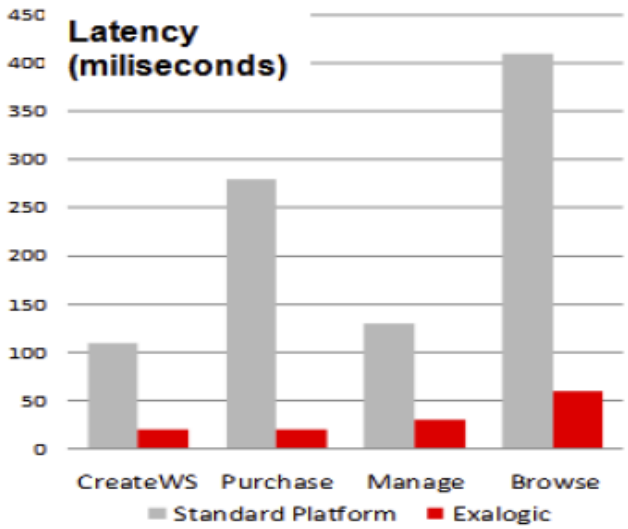

Mellanox IB-based Block Storage through Gateways

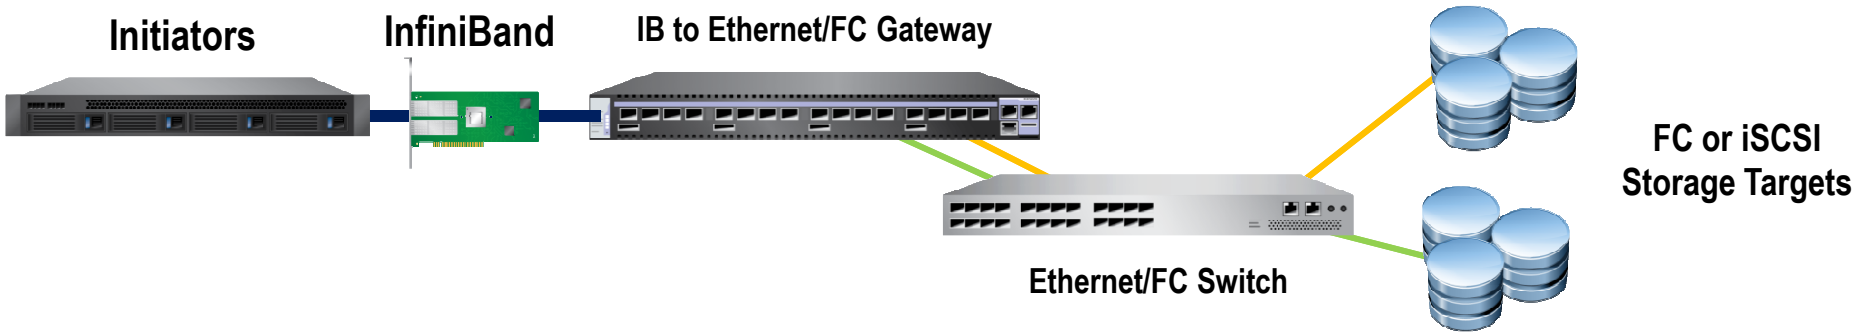

- Bridge InfiniBand clusters to iSCSI or FC storage
- Mellanox BridgeX solution
  - Uses FCoIB and EoIB on hosts
  - FC and Ethernet encapsulation/de-encapsulation over IB
  - FC and Ethernet management on hosts, stateless gateways (no termination)
- Mellanox Vantage-based IB to FC and Ethernet gateways
  - Uses iSER, initiator on hosts, termination in gateways, FC initiator in gateways
  - IPoIB for Ethernet, termination in gateways
- Xsigo IO Director solution
  - Uses vNIC and vHBA implementation over IB on hosts
  - IB Termination in gateway, FC initiator in gateway

VPI Case Study: Online Ticket Agency

- I/O Virtualization and Blade Servers with InfiniBand Yield Cost + Energy Savings and Enable Adaptive, On-demand Computing
  - Online Ticket agency can dynamically manage and balance loads and bandwidth between application and Web servers

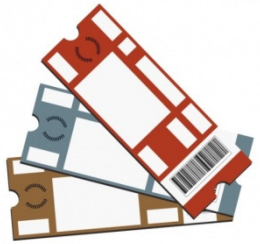

- Highlights of the new deployment include:
  - $1.5M capex savings on infrastructure alone
  - 50% to 75% reduction in deployment time
  - 75% less space required
  - Ticket agency accelerated their deployment per server up to 75 percent.

Enabling Cost-Effective Accelerated Services

Mellanox VPI Connectivity Solution at VMworld Lab

Powered by:

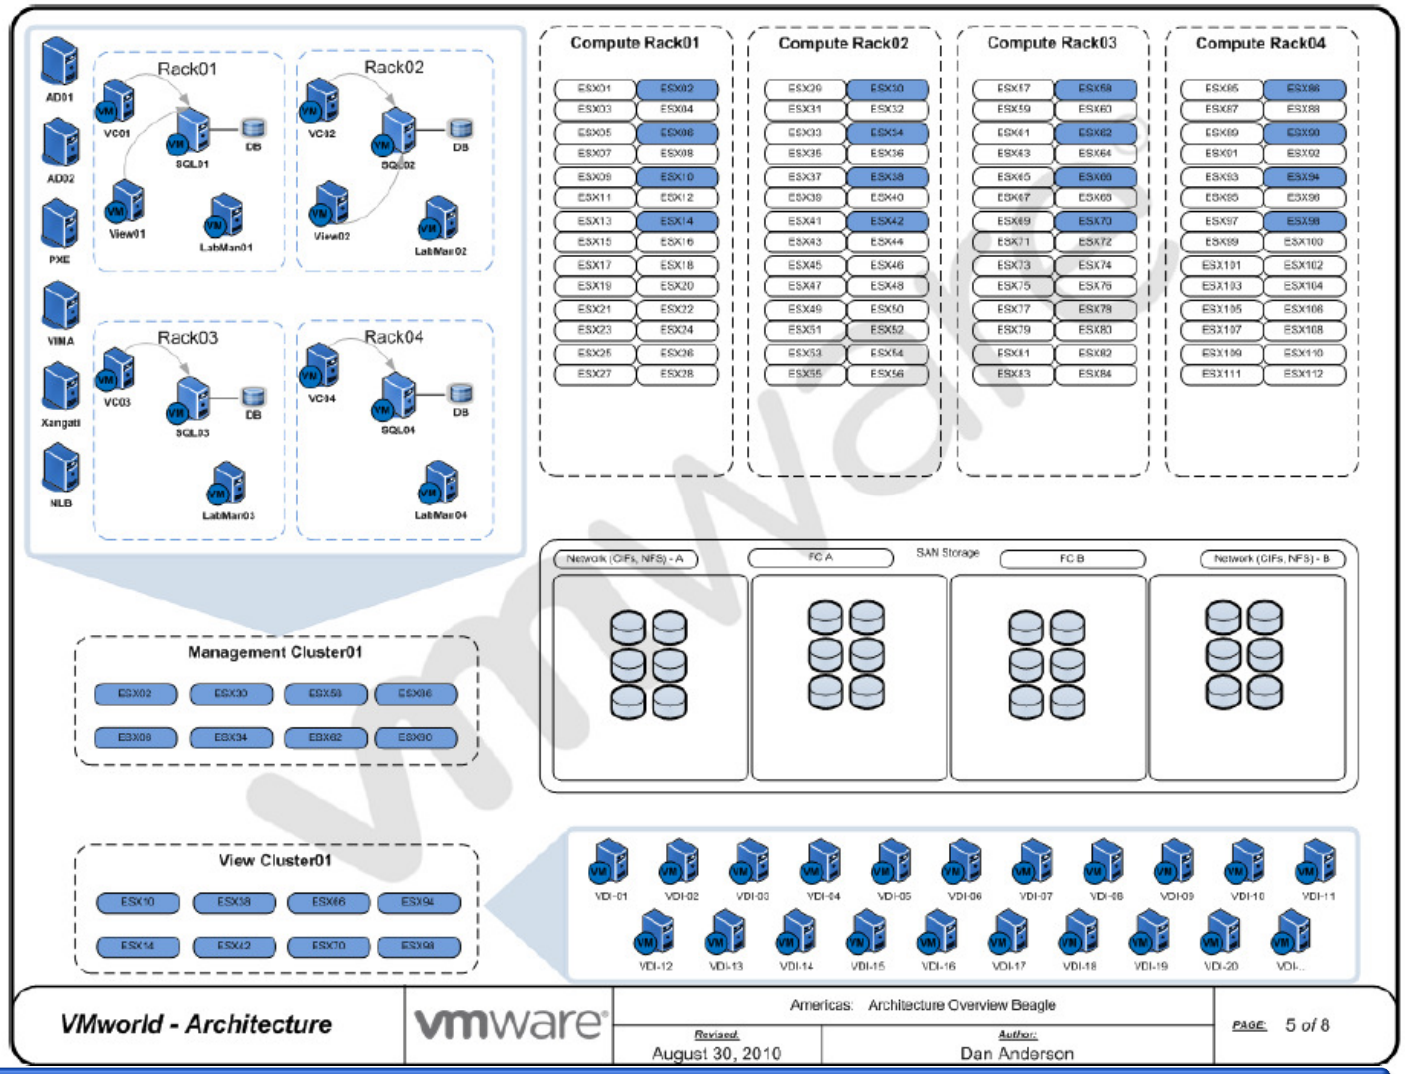

352 Servers, 736 CPUs, 3,072 Cores, 14.6TB Memory, 329TB Storage

Extreme Storage Performance using VSA (iSCSI RDMA)

- 1000 VMs per Rack, generate enormous random I/O load
- Traditional storage systems are the bottleneck in large VM deployments
- Mellanox VSA eliminate the bottleneck by accelerating I/O

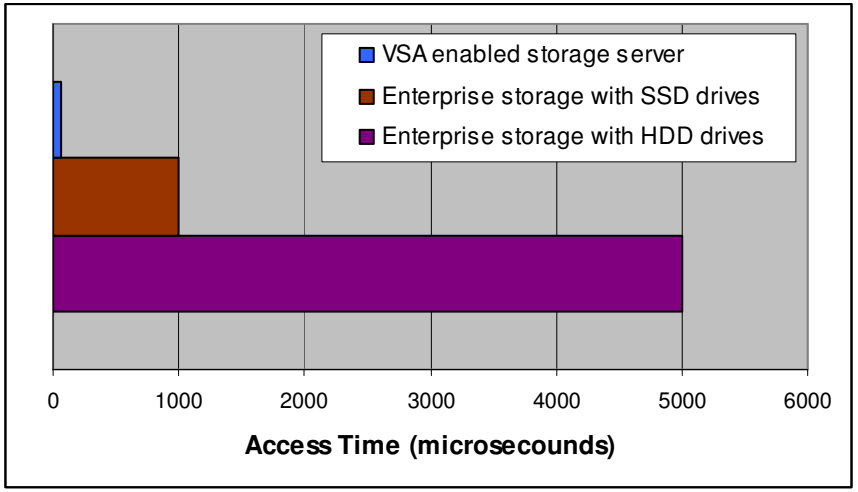

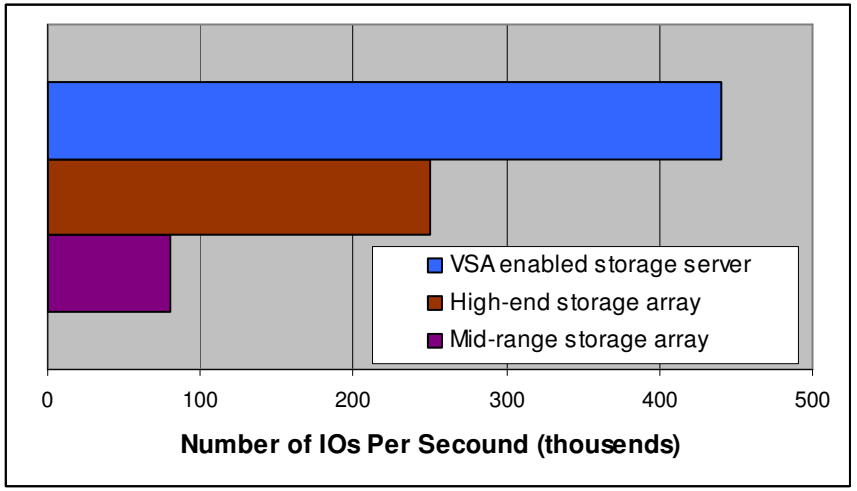

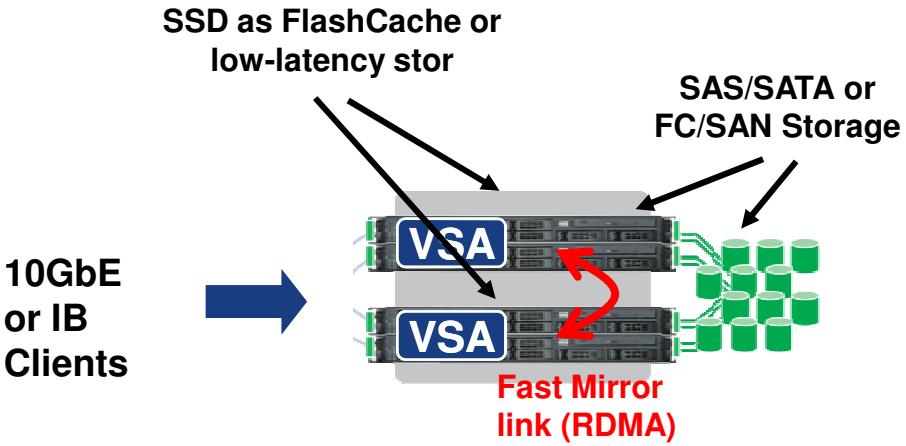

*** VSA performance on HP DL380-G6 + 2 Flash cards**

Exceptional Storage Performance with Mellanox RDMA

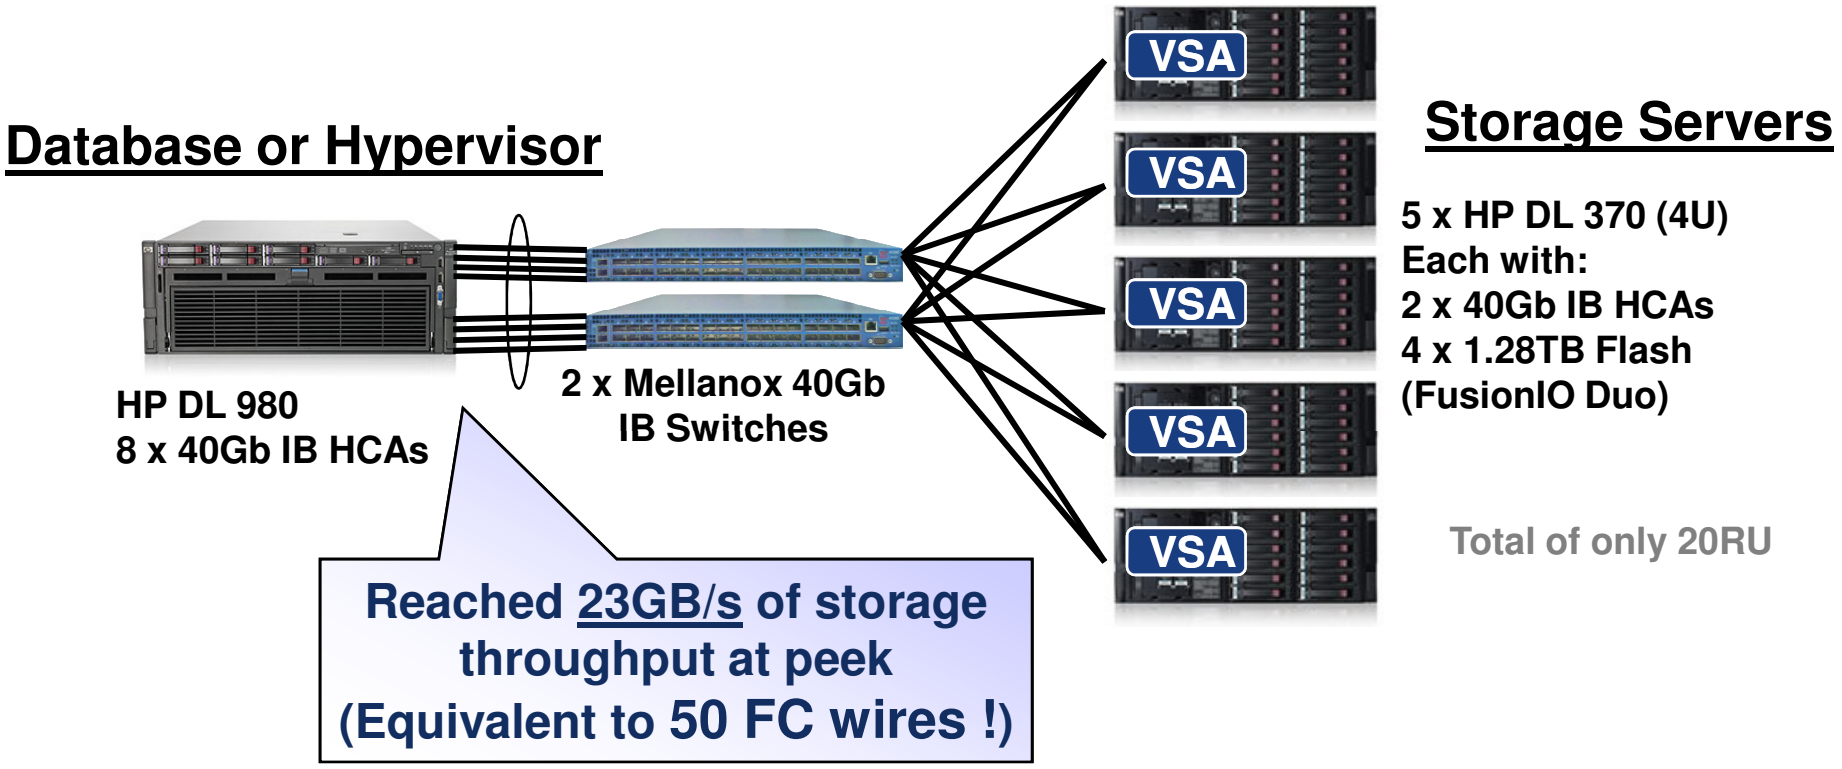

- Exceptional Application performance (single client) of 23GB/s
- Unmatched storage performance of 30GB/s, 2.5M Random IOPs
- 100x faster access time than traditional storage (only 50us)

UFM™ for Fabric IO Provisioning & Monitoring

- Eliminates Complexity
  - Central Dashboard with group activities
  - Virtual IO provisioning and Isolation
  - Automate connectivity migration
  - Rich API and SDK for fabric automation/integration
- Provides Deep Visibility
  - Real-time granular monitoring and mirroring
  - Online threshold based in-context alerts
  - Historical monitoring – for investigations
- Boosts application performance
  - Verify the network is not the bottleneck
  - Eliminate back pressure / congestion
  - Multicast/Unicast routing optimization

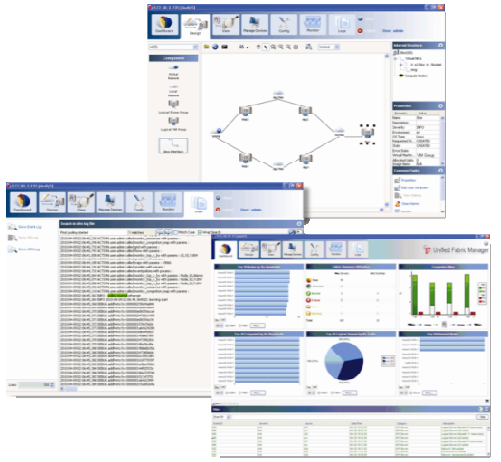

Our Solution

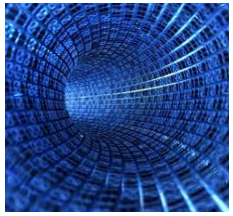

- Eliminate the Network and I/O Bottlenecks
  - I/O, Messaging, and Storage Acceleration
  - High-Bandwidth (40-56Gbps) and Low Latency fabric

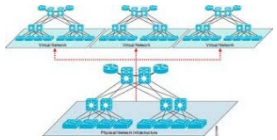

- Consolidate and Virtualize the Network and I/O
  - Less wires with more bandwidth
  - Hardware Based I/O and Network Partitioning (Virtualization)

- Deliver Best Network Efficiency and Scalability
  - Proven scalability without performance degradation
  - 3x lower power and half the cost compared to alternative solutions

- Adaptive Network Provisioning and Monitoring
  - End to end, policy based fabric provisioning
  - Visibility into physical, logical, and aggregated information
  - Open and extendable architecture with simple APIs

Mellanox Solution Competitive Advantage:
Convergence over high performance Virtual Protocol Interconnect (VPI)

Thank You