
COMPUTERWORLD
SNIA
SNW
April 12-15, 2010 | Rosen Shingle Creek Resort | Orlando, Florida

SNIA
SNW
COMPUTERWORLD
April 12-15, 2010
Rosen Shingle Creek Resort
Orlando, Florida

Data Center Challenges

- Cabling Complexity
- Policy and Regulatory Demands

- Reduce Carbon Footprint
- Data Center Limitations

- Scale Resources for Server Virtualization
- Expand Network Bandwidth

- Density Challenges
- Floor Space
- Power per Rack
- Simple Management

Deployment Complexity

Environmental Factors

Increased Computing Resources

Data Center Consolidation

SNIA
SNW
COMPUTERWORLD
April 12-15, 2010
Rosen Shingle Creek Resort
Orlando, Florida

10GbE Networking Trends

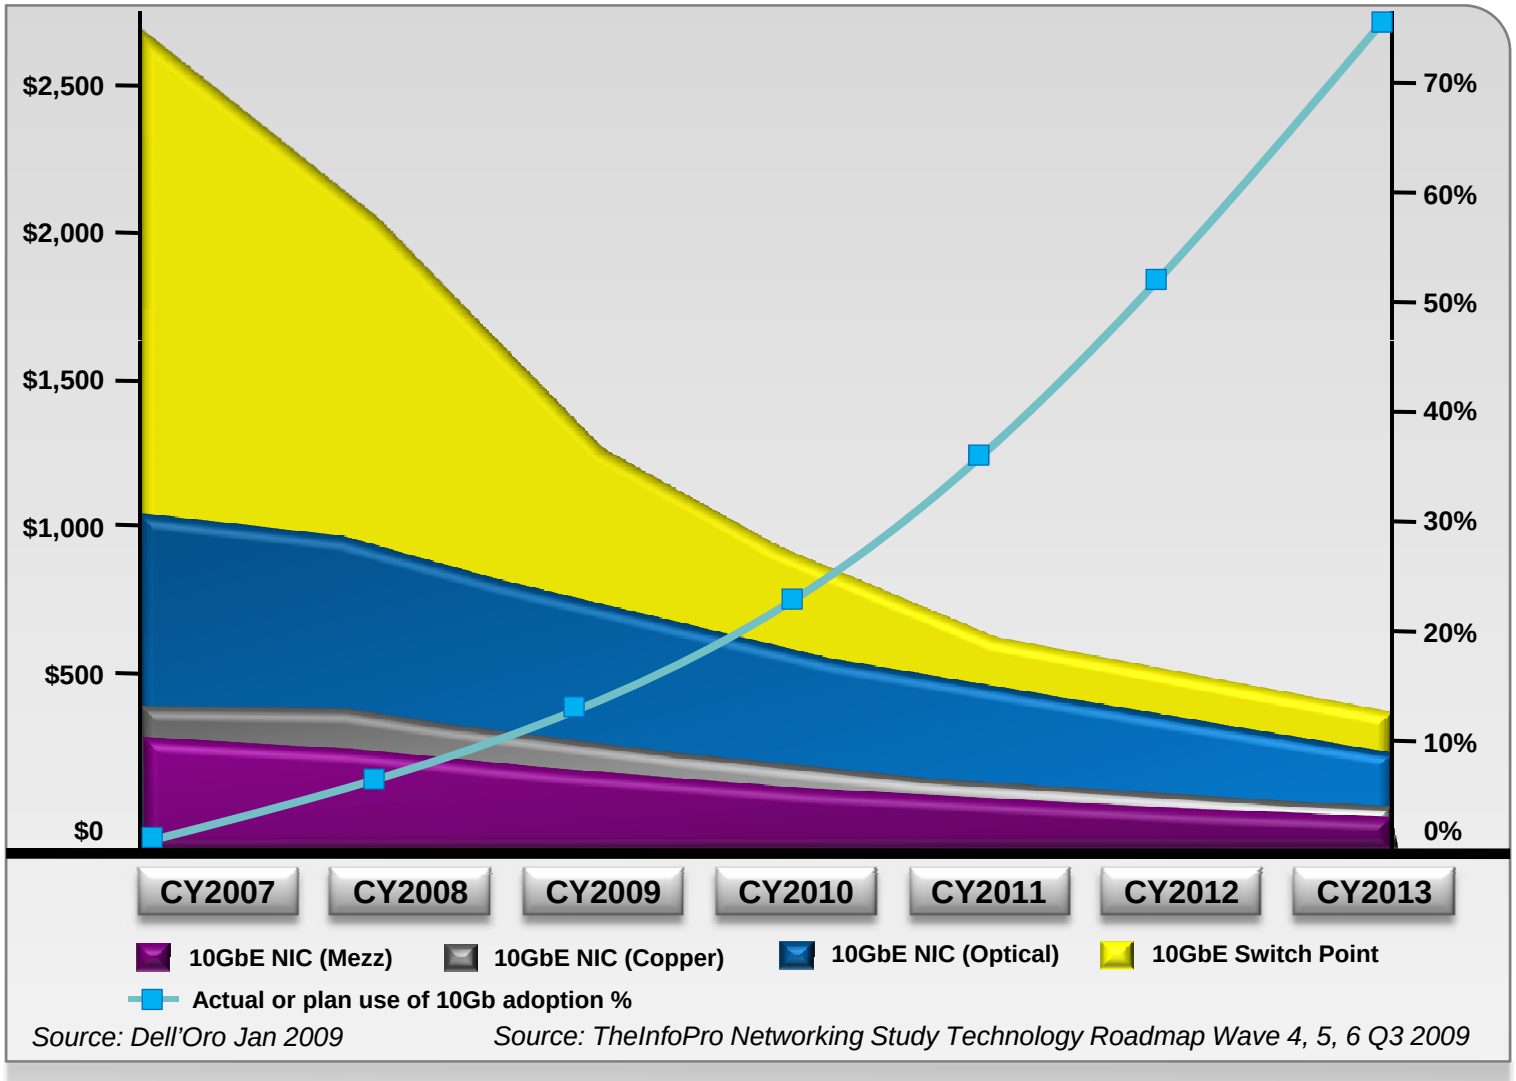

Source: Dell'Oro Jan 2009

Source: TheInfoPro Networking Study Technology Roadmap Wave 4, 5, 6 Q3 2009

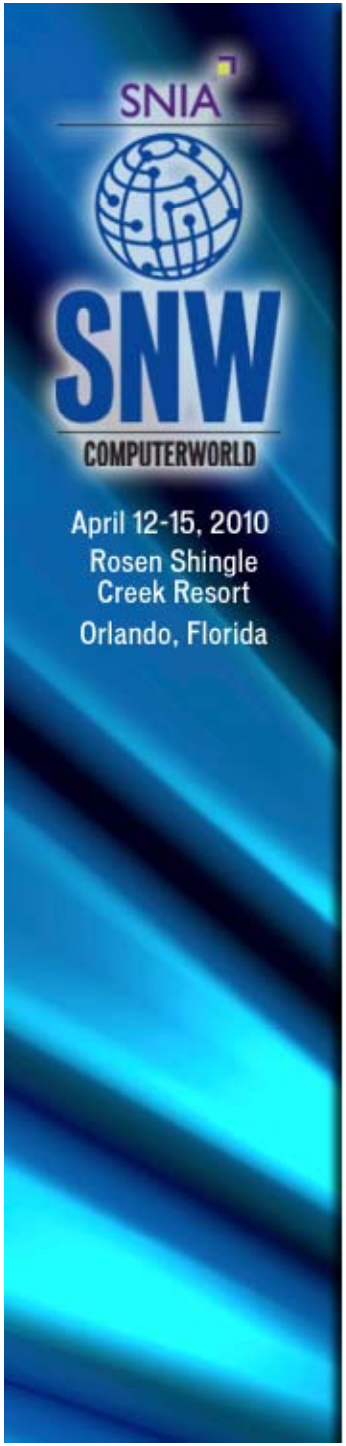

10 THINGS TO CONSIDER WITH 10GBE DEPLOYMENT

#1) Switch Architecture Top of Rack

- Modular unit design – managed by rack
- Servers connect to 1-2 Ethernet switches inside rack
- Rack connected to data center aggregation layer with SR optic or twisted pair
- Blades effectively incorporate TOR design with integrated switch

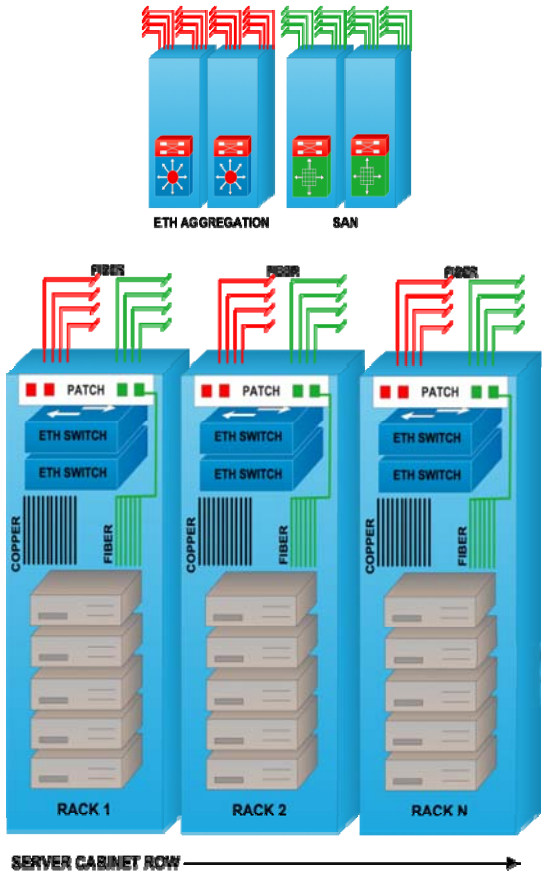

Switch Architecture End of Row

- Server cabinets lined up side-by-side
- Rack or cabinet at end (or middle) of row with network switches
- Fewer switches in the topology
- Managed by row – fewer switches to manage

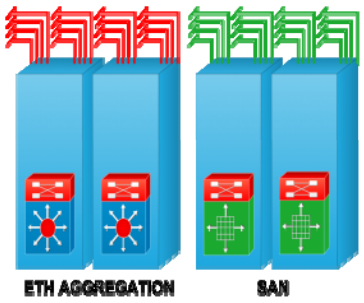

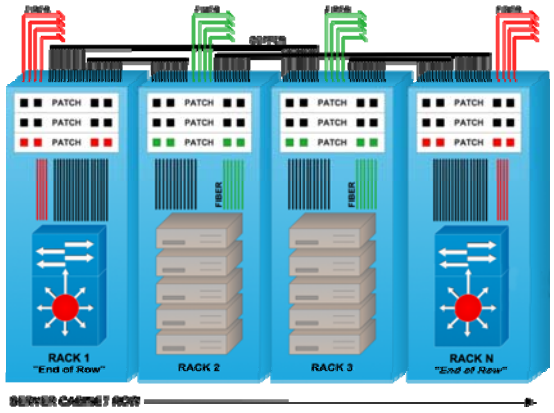

#2) Server Architecture

VS.

- Generally more I/O & memory expansion
- More expensive 10G physical interconnect
- Later 10G LOM due to 10GBASE-T

- Typically limited I/O expansion (1-2 slots)
- Much cheaper 10G physical interconnect – backplane Ethernet & integrated switch – leading 10G transition
- Earlier 10G LOM

Server Architecture

Rack Servers Today

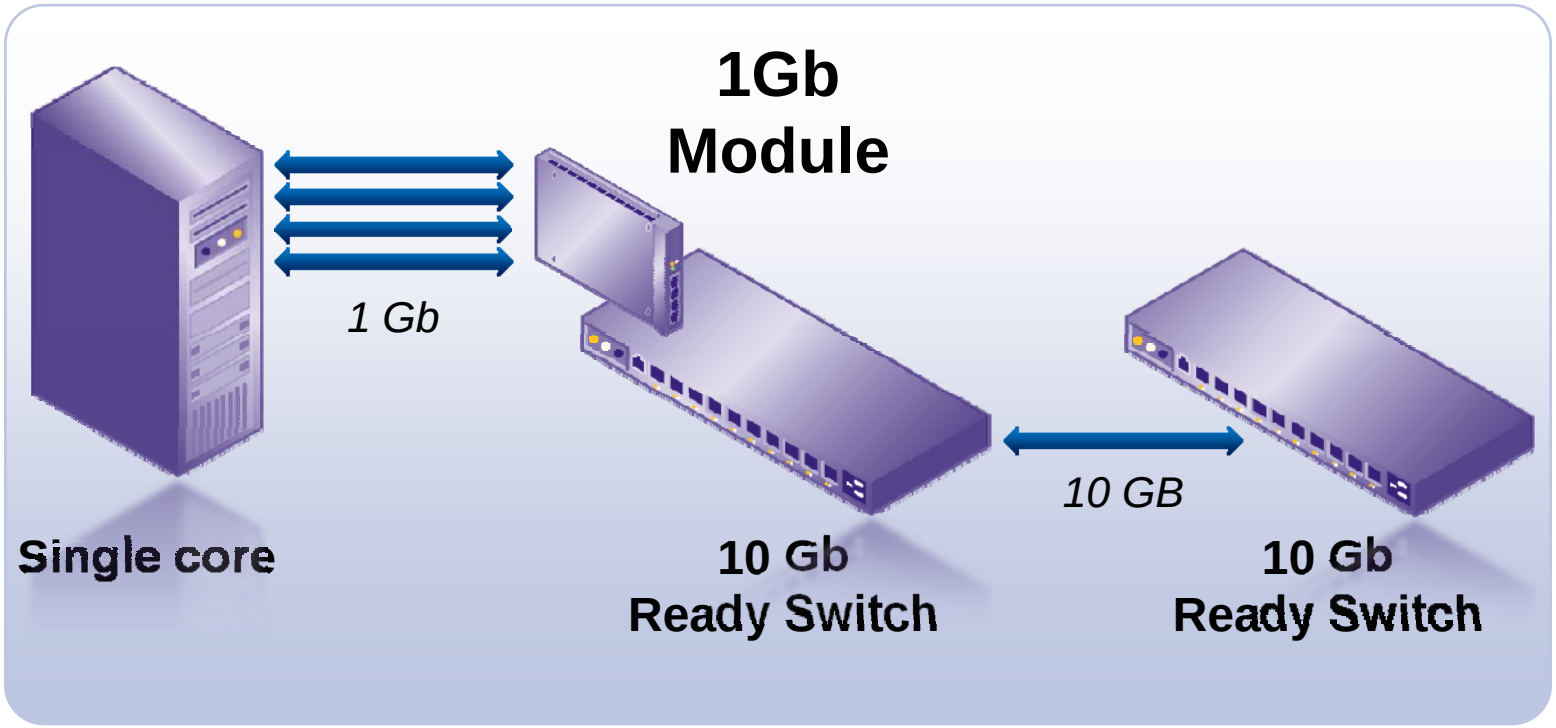

Server Refresh

Adding servers with 10GB connection to switch

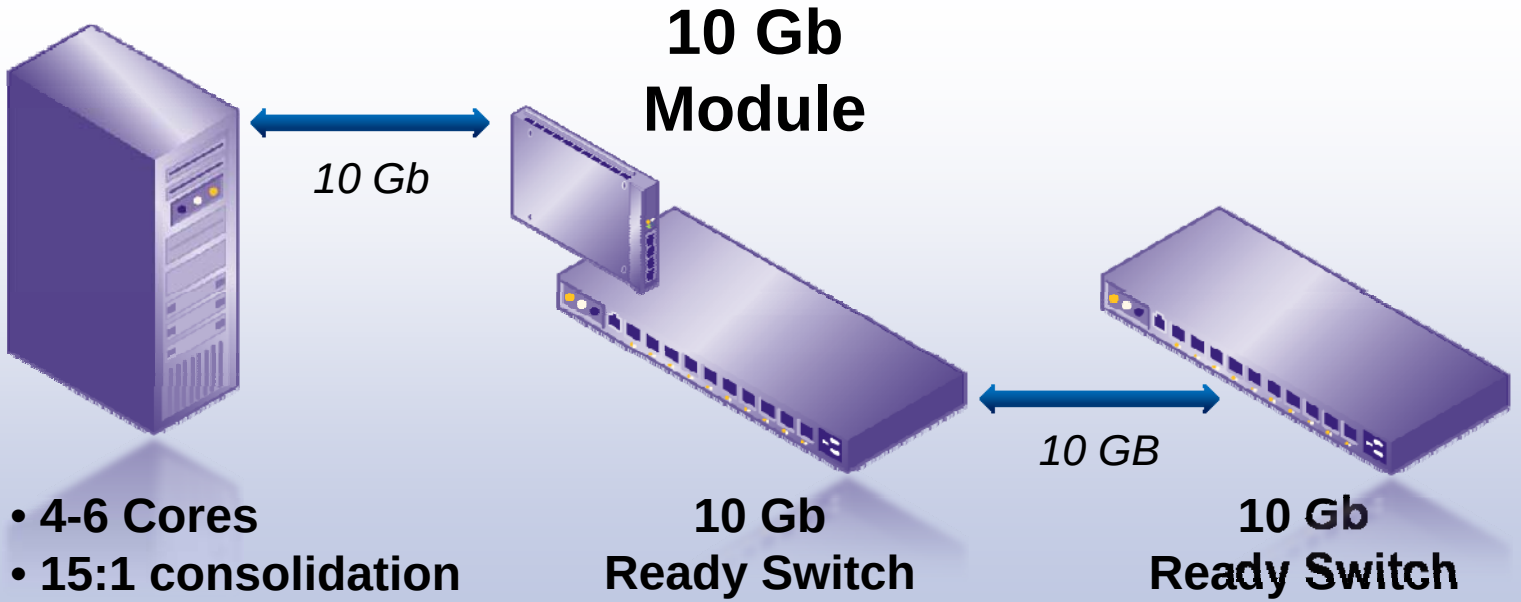

Blade Servers

PAIN POINTS:

- Limited real estate for Ethernet bandwidth

TODAY:

- 1 Gb LOM ports
- 1 Gb Mezz card

ONECONNECT:

- 10 Gb LOM ports
- 10 Gb Mezz ports

VIRTUAL NIC:

- 4 VNICs per port

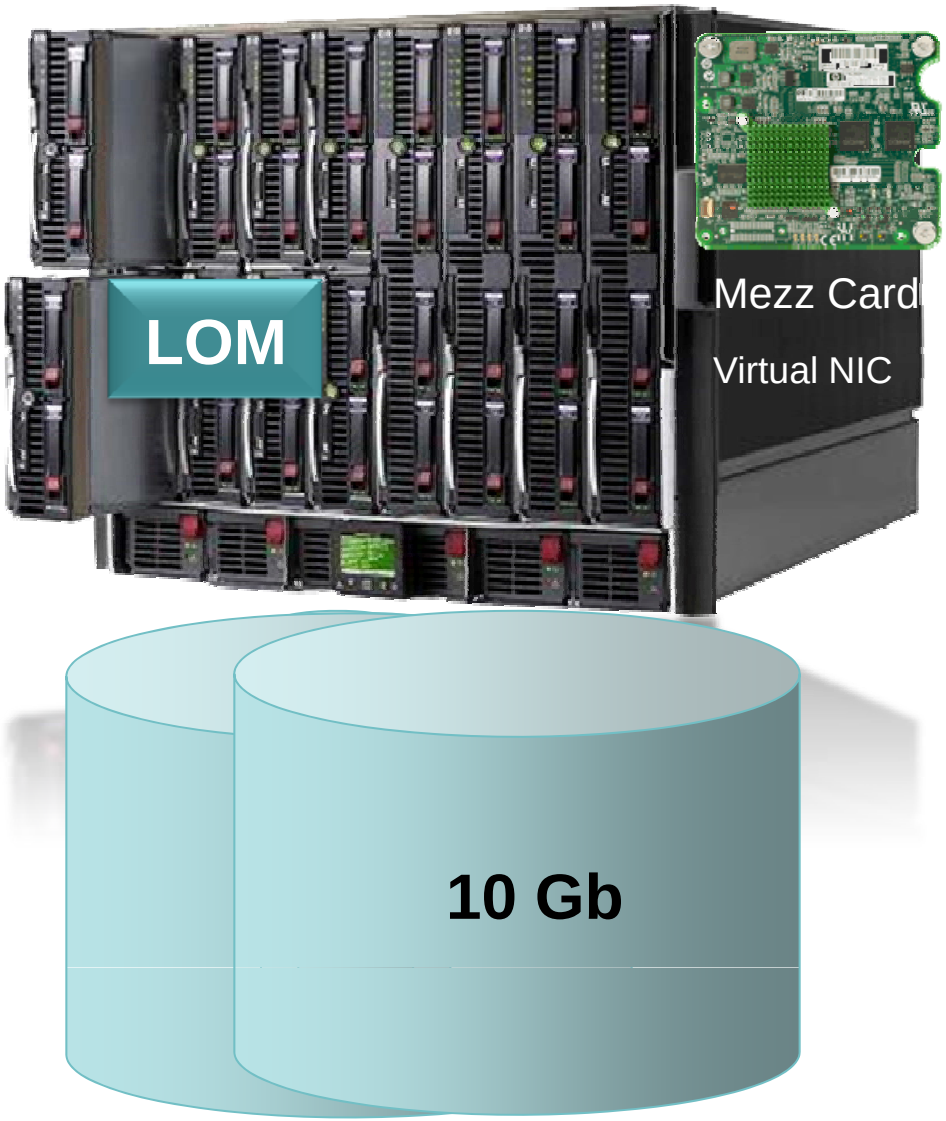

#3) Cables

- Today's Data Center
  - Typically Twisted Pair for Gigabit Ethernet
    - Some Cat 6, more likely Cat 5
  - Optical cable for Fibre Channel
    - Mix of older OM1 and OM2 (orange) and newer OM3 (green)
    - Fibre Channel may not be wired to every rack

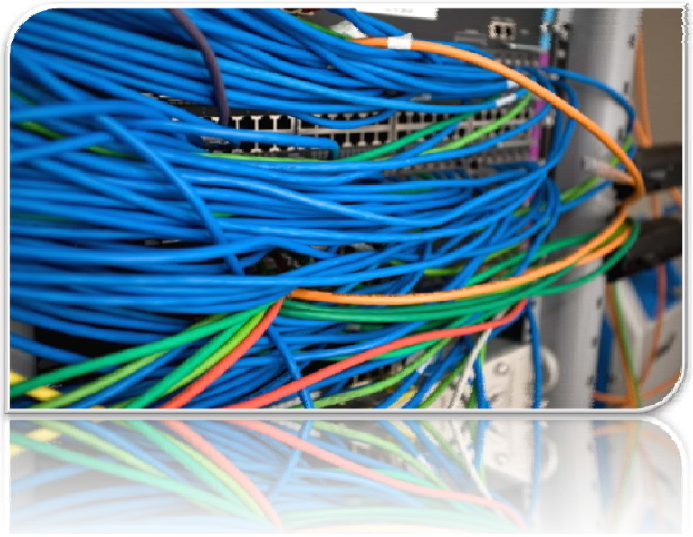

10GbE Cable Options

Mid 1980's	*Mid 1990's*	*Early 2000's*	Late 2000's
10Mb	**100Mb**	**1Gb**	**10Gb**
UTP Cat 3	*UTP Cat 5*	*UTP Cat 5* *SFP Fiber*	SFP+ Cu, Fiber X2, Cat 6/7

Technology	Cable	Distance	Power (each side)	Transceiver Latency (link)
SFP+ Direct Attach	Copper Twinax	10m	~0.1W	~0.1µs
10GBase SR (short range)	Optic Multi-Mode	62.5µm - 82m 50µm - 300m	1W	~0
10GBase LR (long range)	Optic Single-Mode	10 km	1W	Var
10GBASE-T	Copper Twisted Pair	Cat6 - 55m Cat 6a 100m Cat7 -30m	~8W ~8W 4W	2.5µs 2.5, 1.5 µs

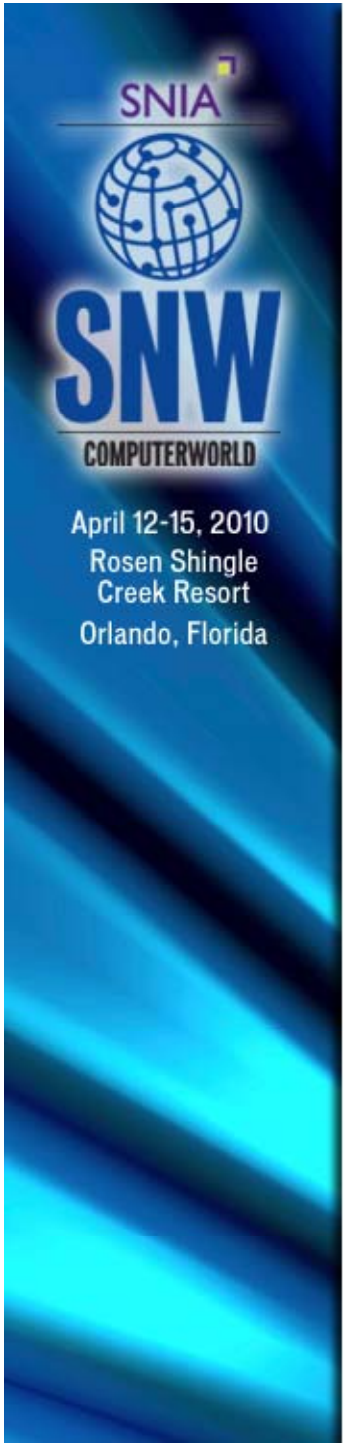

New 10GBASE-T Options

Top of Rack

- SFP+ is lowest cost option today
  - Doesn't support existing gigabit Ethernet LOM ports
- 10GBASE-T support emerging
  - Lower power at 10-30m reach (2.5w)
  - Energy Efficient Ethernet reduces to ~1w on idle

End of Row

- Optical is only option today
- 10GBASE-T support emerging
  - Lower power at 10-30m reach (2.5w this year)
  - Can do 100m reach at 3.5w
  - Energy Efficient Ethernet reduces to ~1w on idle

#4) Data Center Bridging

Ethernet is a "best-effort" network

- Packets may be dropped
- Packets may be delivered out of order
  - Transmission Control Protocol (TCP) is used to reassemble packets in correct order

Data Center Bridging (aka "lossless Ethernet")

- The "bridge to somewhere" – pause-based link between nodes
- Provides low latency required for FCoE support
- Expected to benefit iSCSI, enable iSCSI convergence

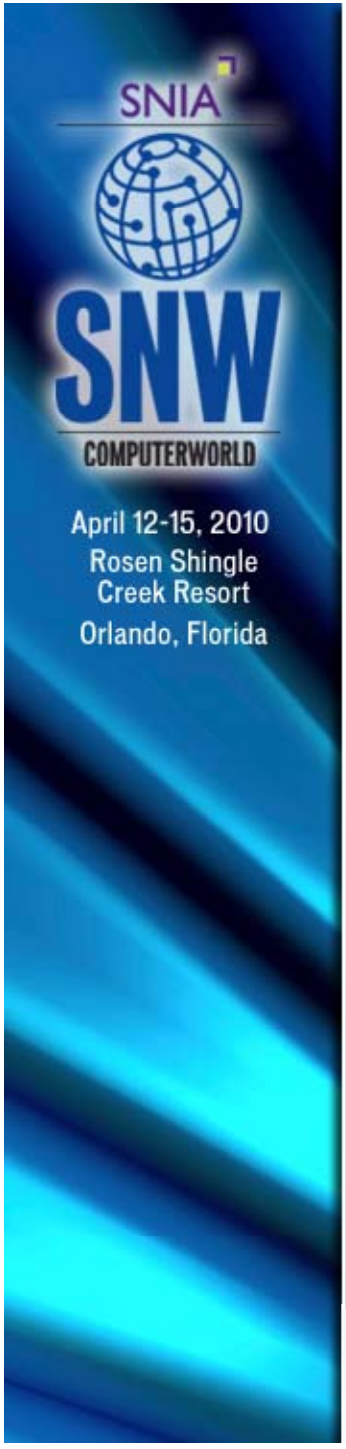

IEEE Data Center Bridging Standards

Feature	Benefit	Standards Activity
Priority-based Flow Control (PFC)	Enables lossless Ethernet – manages I/O between initiator and target on a multi-protocol Ethernet link	IEEE 802.1Qbb
Quality of Service (QoS)	Supports 8 priorities for network traffic	IEEE 802.1p
Enhanced Transmission Selection (ETS)	Allocate bandwidth to IP, iSCSI and FCoE traffic – managed with OneCommand 5.0	IEEE 802.1Qaz
Data Center Bridging Capability Exchange (DCBX)	Extends DCB network by exchanging Ethernet parameters between DCB switches	IEEE 802.1ab

#5) Network Convergence

CONVERGED FABRICS
Phase 1

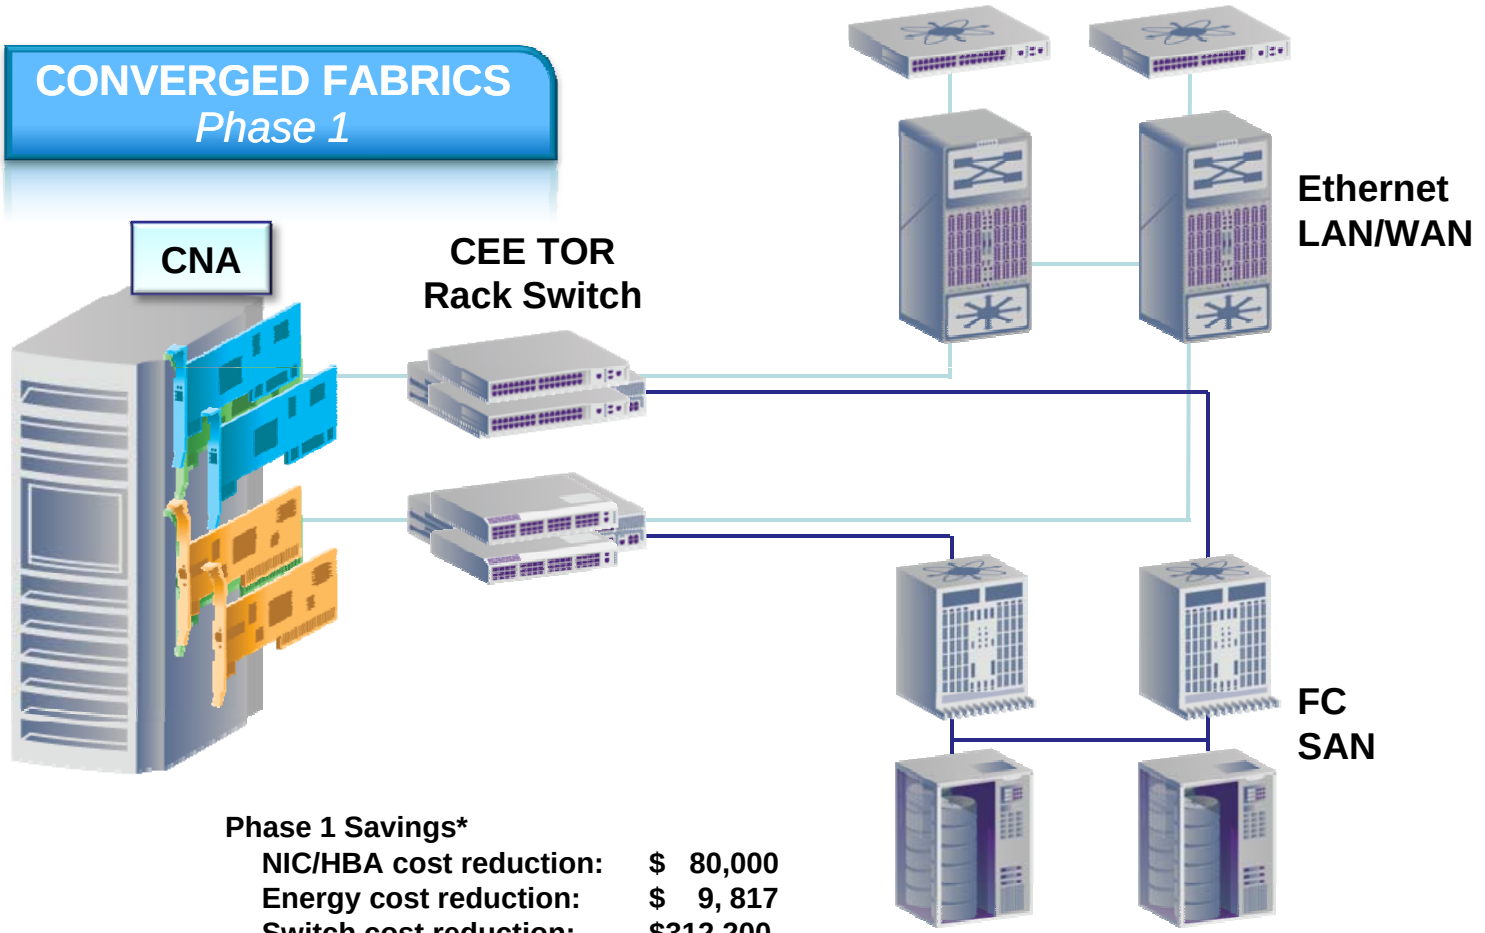

Phase 1 Savings*

NIC/HBA cost reduction:	$ 80,000
Energy cost reduction:	$ 9, 817
Switch cost reduction:	$312,200

The 1st Phase of Network Consolidation will Occur in the Server Rack

*100 servers w/ 2 NICs and 2 HBAs,
5 RAID storage systems

Convergence in Fabric

CONVERGED FABRICS
Phase 2

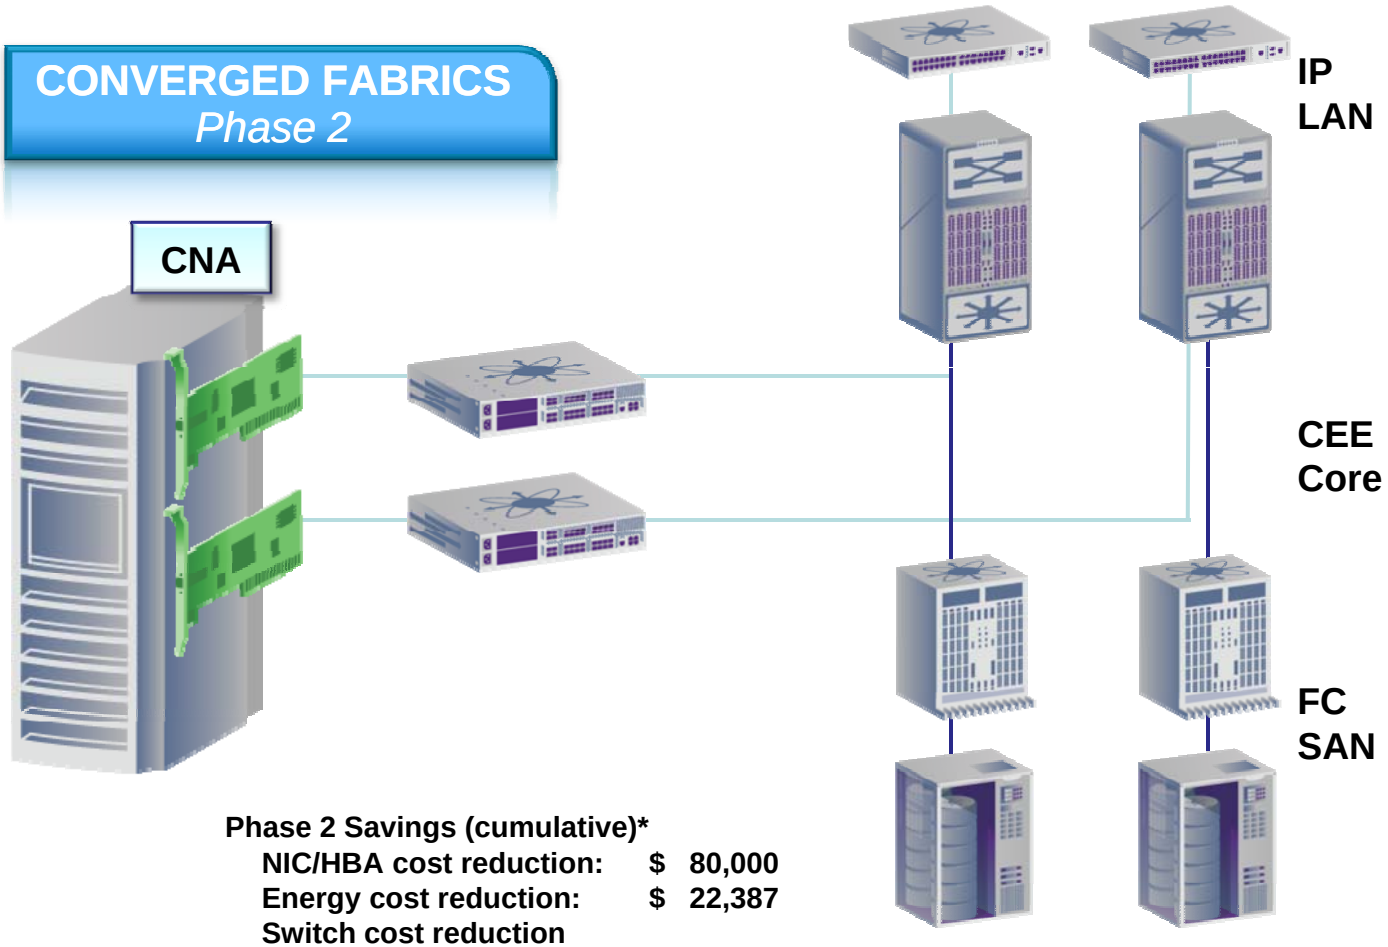

Phase 2 Savings (cumulative)*

NIC/HBA cost reduction:	$ 80,000
Energy cost reduction:	$ 22,387
Switch cost reduction	

The 2nd Phase will Consolidate TOR and IP Switching

*100 servers w/ 2 NICs and 2 HBAs,
5 RAID storage systems

SNIA
SNW
COMPUTERWORLD
April 12-15, 2010
Rosen Shingle Creek Resort
Orlando, Florida

Convergence in Storage

CONVERGED STORAGE
Phase 3

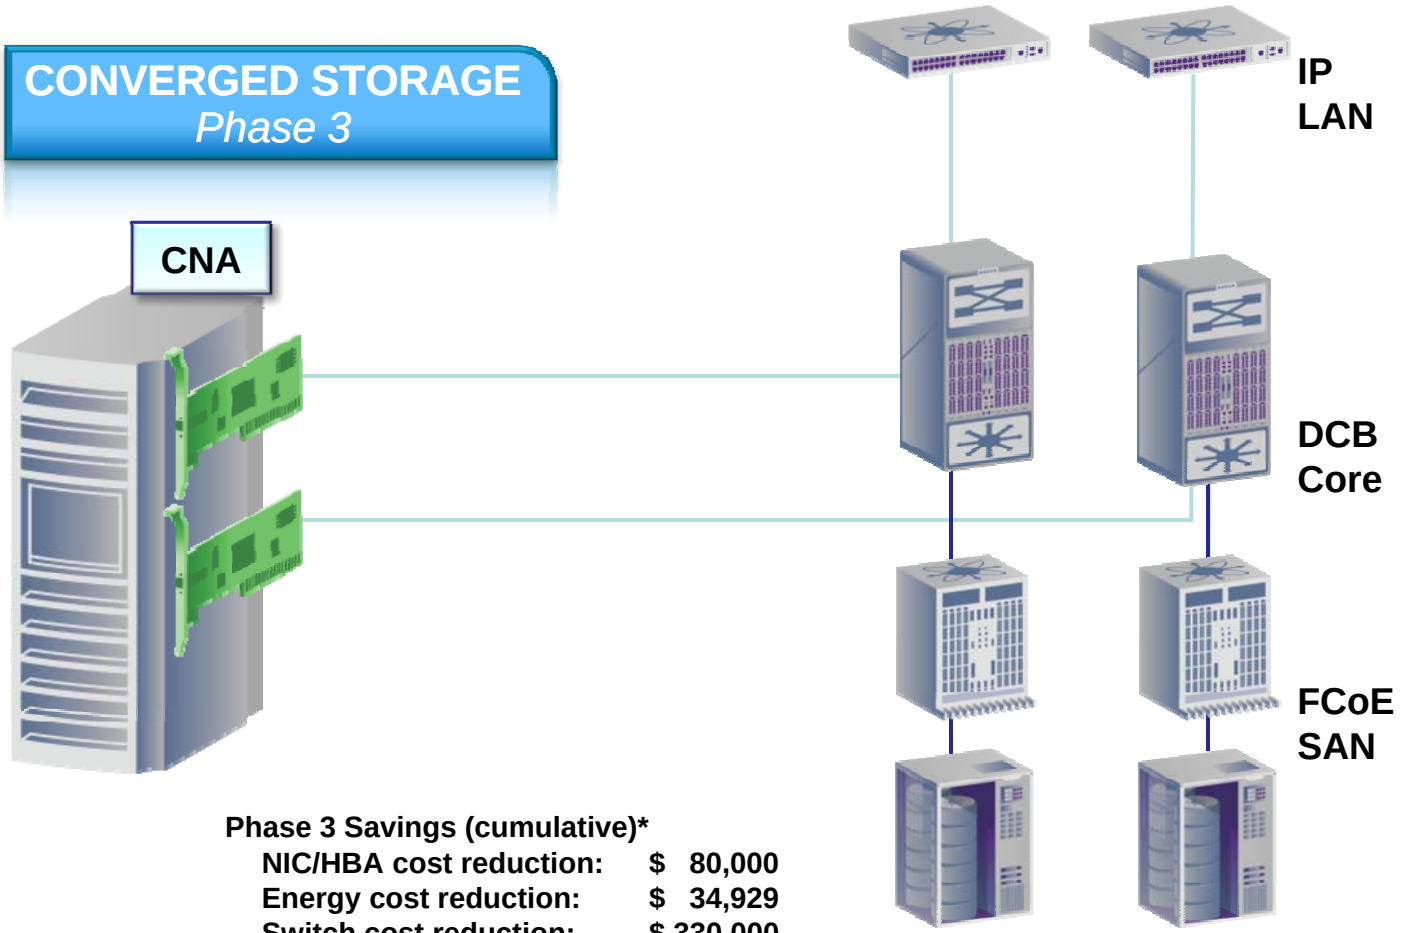

Phase 3 Savings (cumulative)*

NIC/HBA cost reduction:	$ 80,000
Energy cost reduction:	$ 34,929
Switch cost reduction:	$ 330,000

The 3rd Phase will Consolidate FC Switching and Storage

*100 servers w/ 2 NICs and 2 HBAs,
5 RAID storage systems

SNIA

SNW

COMPUTERWORLD

April 12-15, 2010

Rosen Shingle Creek Resort

Orlando, Florida

Savings with Convergence

Before Convergence 140 Cables

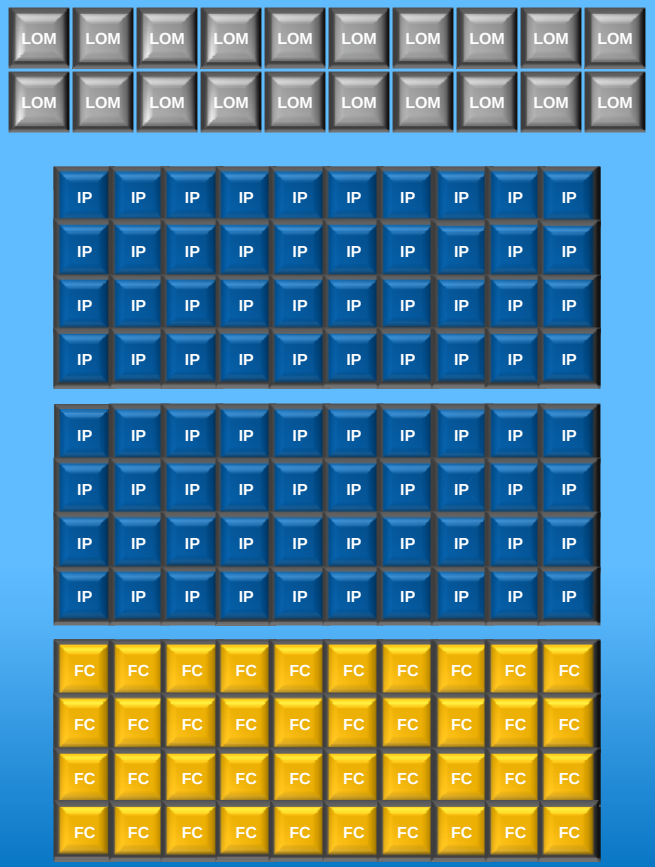

After Convergence 60 Cables

Savings Up To:

- 28% on switches, adapters
- 80% on cabling
- 42% on power and cooling

(Based on 2 LOM, 8 IP and 4 FC Ports on 10 Servers)

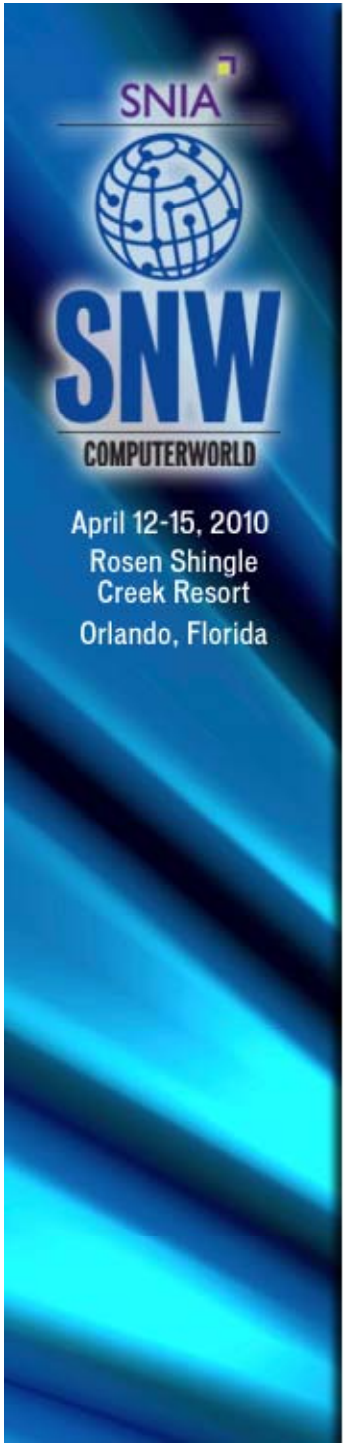

To Converge or Not To Converge?

Volume x86 systems

- Best business case for convergence
- Systems actually benefit by reducing adapters, switch ports and cabling

High End Unix Systems

- Limited convergence benefit
- Typically large number of SAN and LAN ports
- System may go from 24 Ethernet and 24 Fibre Channel ports to 48 Converged Ethernet ports
- Benefits in later stages of data center convergence when Fibre Channel SAN fully running over Ethernet physical layer (FCoE)

SNIA
SNW
COMPUTERWORLD
April 12-15, 2010
Rosen Shingle Creek Resort
Orlando, Florida

#6) Server Virtualization

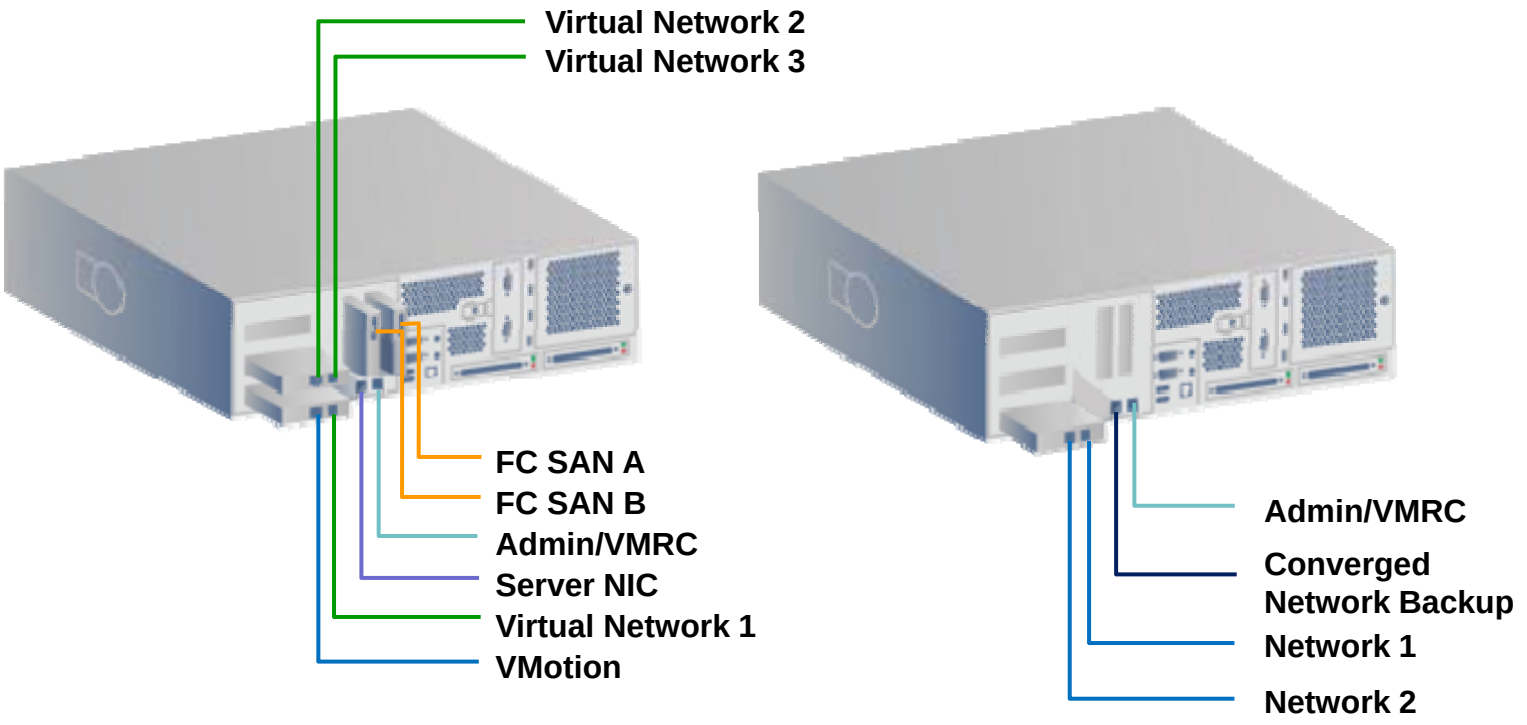

Separate LAN and SAN Connectivity With Many NICs and HBAs

Converged Network With CNA

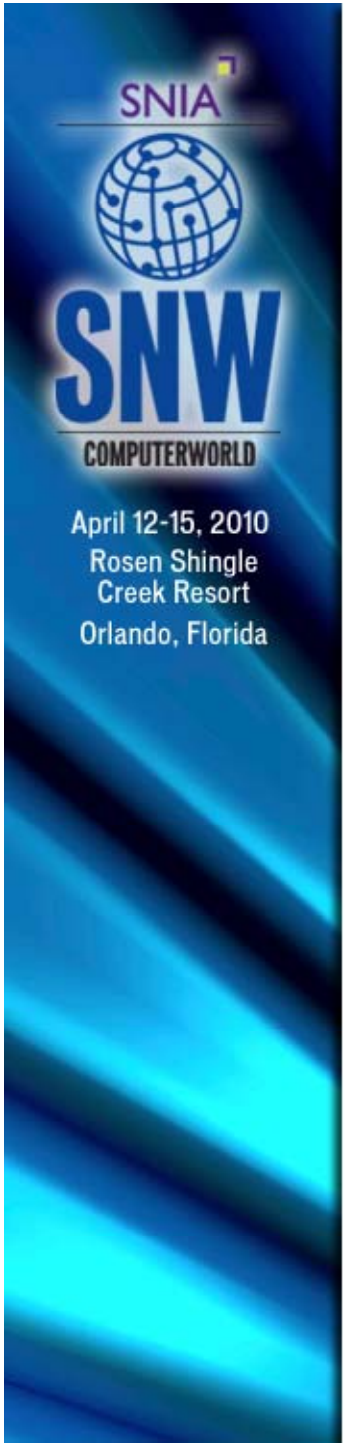

Virtual NICs

- Available today for blade servers
- Each physical port can be divided into multiple virtual ports
  - Look like physical ports to the OS or hypervisor
  - Looks like physical ports to the network
- Bandwidth allocated to virtual ports
- Virtual ports networked with virtual machines

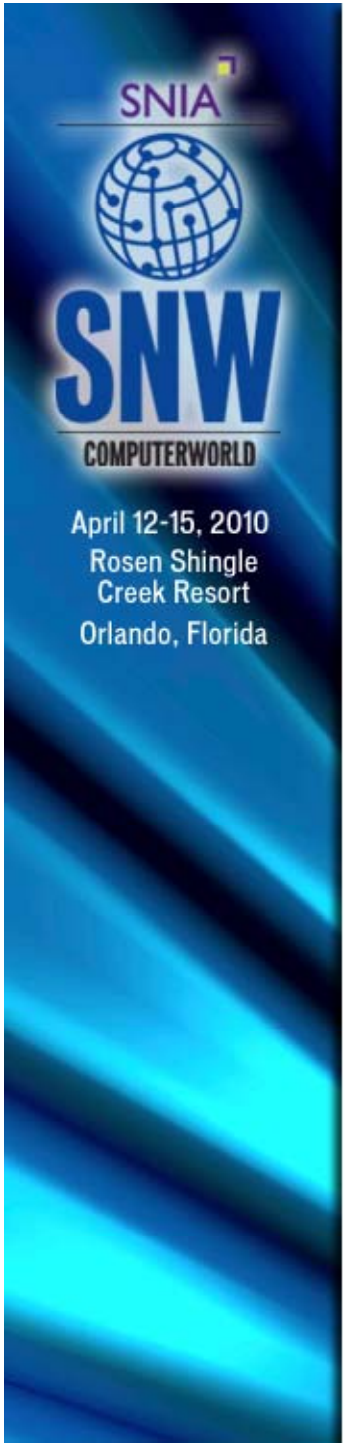

#7) Select 10Gb Adapter

- 10Gb LOM becoming standard on high-end servers
- Second adapter for high availability
- Should support FCoE and iSCSI offload for network convergence
- Compare performance for all protocols

Performance – TCP/IP