# 10GbE – Key Trends, Drivers and Predictions

# Ethernet Storage Forum Members











The SNIA Ethernet Storage Forum (ESF) focuses on educating end-users about Ethernet-connected storage networking technologies.

# Housekeeping

SNIA
Advancing storage & information technology

- Presentation with live Q&A at the end
- Questions submitted via web tool will be answered verbally
- Unanswered questions will be placed on www.sniaesfblog.org
- Request complete the 10 second feedback form

# Today's Panel

SNIA
Advancing storage & information technology

- Jason Blosil, SNIA–ESF Chair
- Gary Gumanow, SNIA–ESF Board of Directors
- David Fair, SNIA–ESF Board of Directors

# SNIA Webinar Agenda



- Trends driving new networking requirements
- Technologies enhancing 10GbE
  - New server platforms
  - I/O virtualization
  - Network flexibility (cabling and DCB)
- 10GbE market adoption
- Future of 1GbE / Fast Ethernet
- Where do we go from here

# Top Storage Trends

SNIA
Advancing storage & information technology

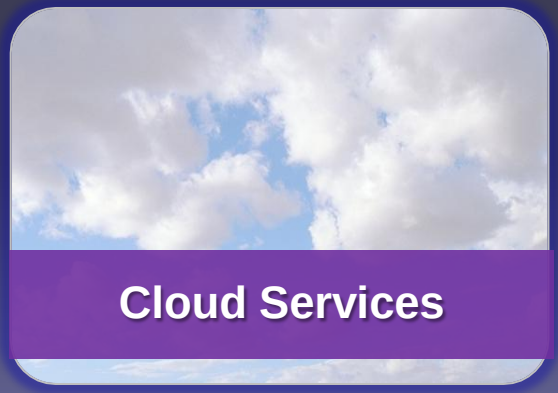





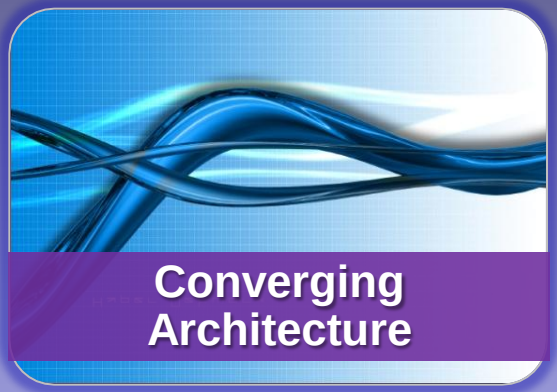

# Drivers of 10GbE

Economic Drivers

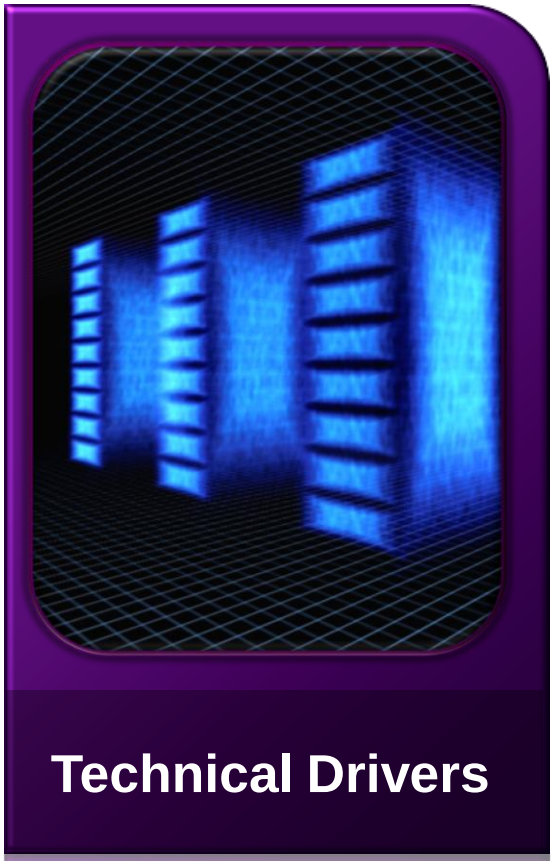

Technical Drivers


Competitive Advantage

**Increased Use of Virtual Resources for Real Work Loads!**

# Economic Drivers – 10GbE
# It's about $/Gb

SNIA
Advancing storage & information technology

- Hardware
  - 1GbE
    - Adapter – sub $100 / port
    - Switch – sub $75/ port
  - 10GbE
    - Adapter – sub $300 / port ($30 / Gb)
    - Switch – sub $500 / port ($50 / Gb)
- Green Initiatives
  - Data center avoidance
  - Data center efficiency

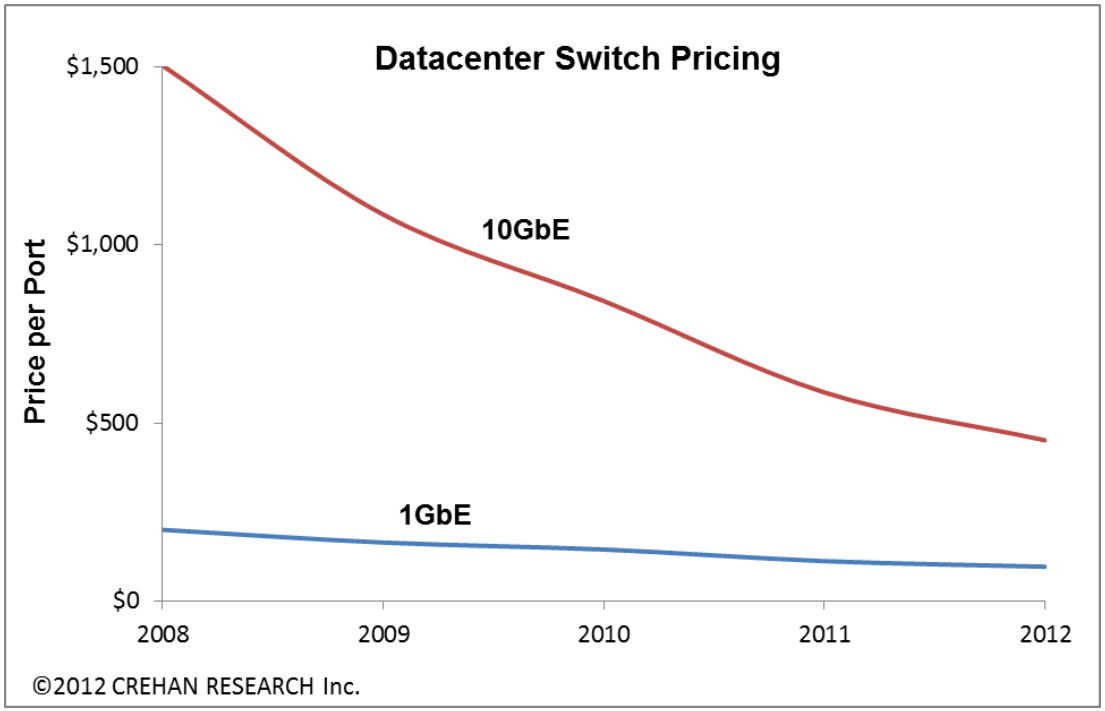


10GbE offers 10x the bandwidth at 5x the price

**OR**

a 50% reduction in $ / Gb of bandwidth

# Cloud Services

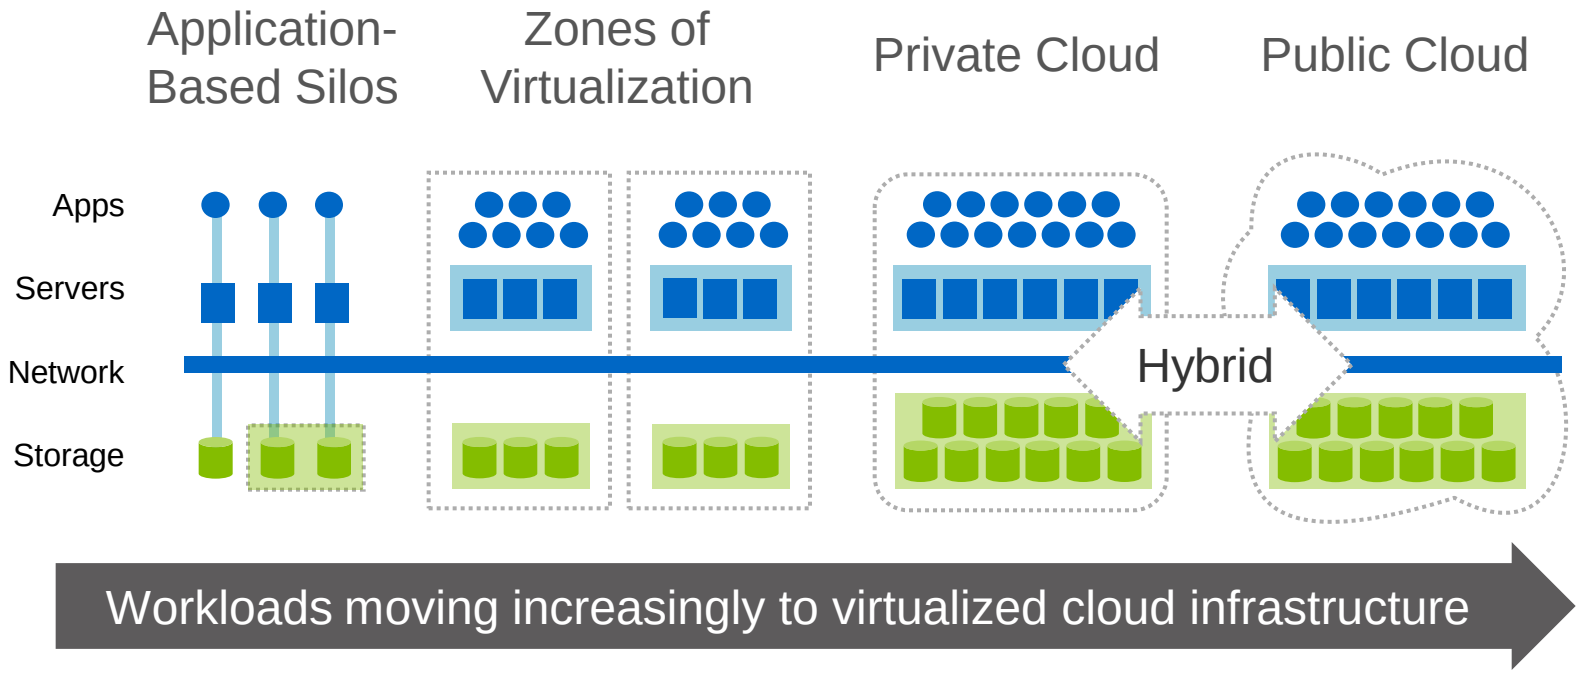


- Private Cloud: behind the firewall of an enterprise, closed to public
- Public Cloud: accessible via service providers for general purchase
- Hybrid Cloud: private clouds linked to public clouds

# Disruption of Flash Technology

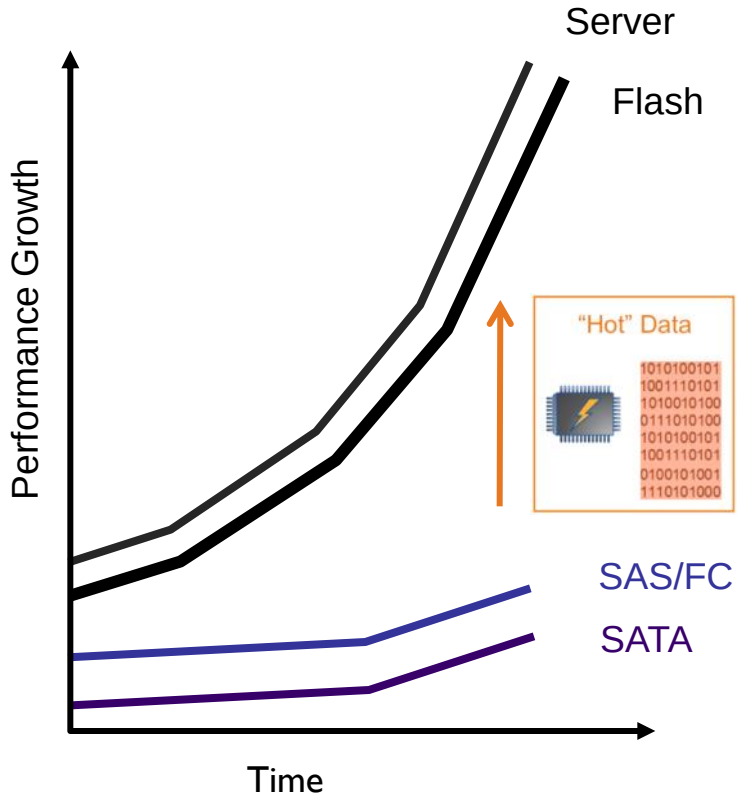


- Relatively small differences between HDD types
- Flash is a game changer
- Balancing cost and performance is key
- Flash performance relocates performance bottlenecks

# SSD Performance

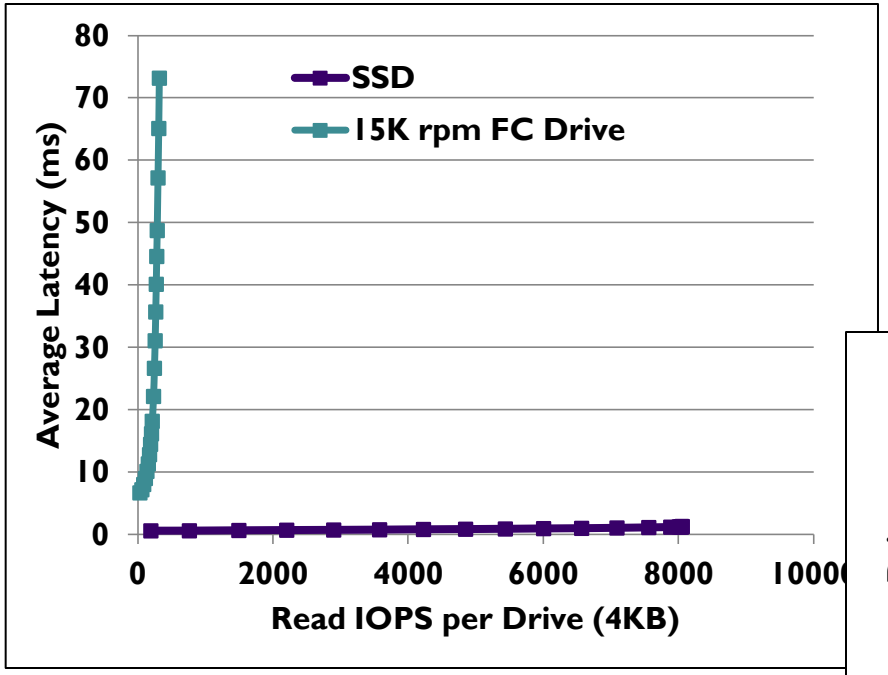


SSD versus 15k rpm FC, 4KB random reads

Sequential I/O throughput per drive

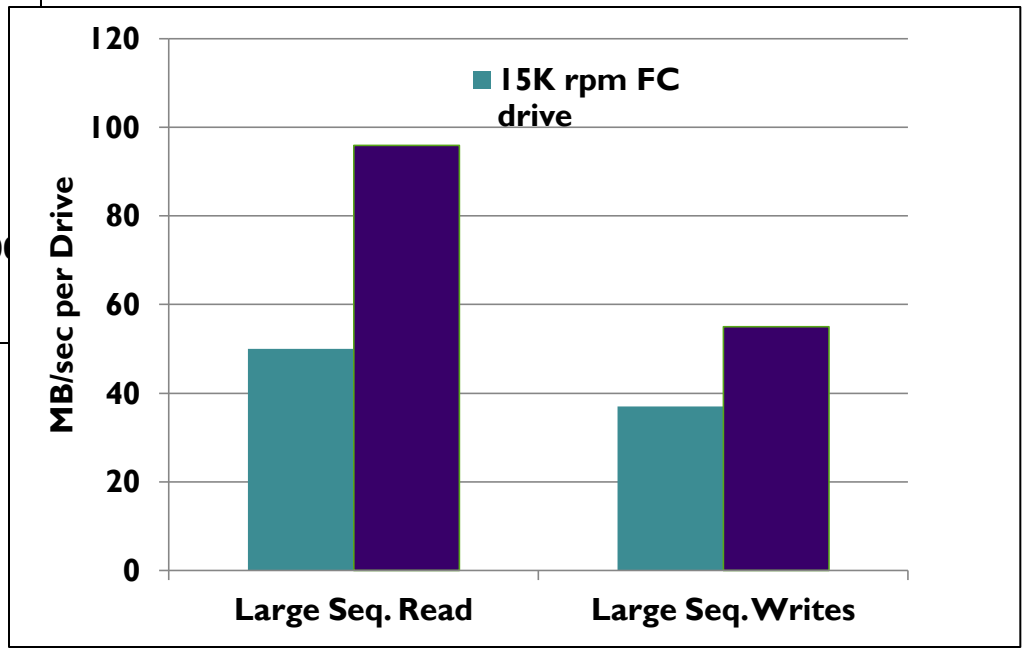

# End-to-End Flash Categories

SNIA
Advancing storage & information technology

## Host-side Flash Software

- Software only, may be tied to particular flash hardware

## Flash as DAS / Cache

- Flash hardware, stores persistent data
- May be combined with software to form cache

## Flash-based VSA

- Software

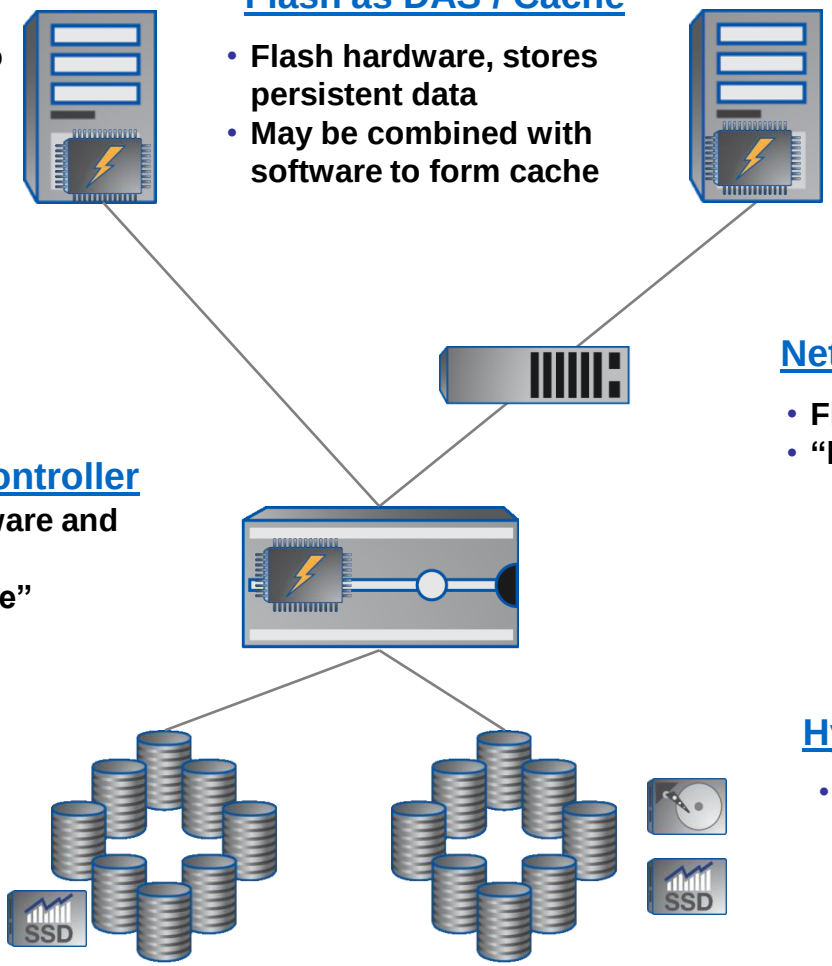


## Network-based Flash

- Flash Hardware and software
- "Bump in the wire"

## Flash in Controller

- Flash hardware and software
- "Behind wire"

## Pure Flash in Array

- All flash

## Hybrid Flash / HDD Array

- Mixed flash / HDD

# Vote #1

# 2012 Server Platforms Enable New Ethernet Capabilities

SNIA
Advancing storage & information technology

- New server platforms coming to market as of March 2012 let Ethernet accomplish more than ever
- Bandwidth per port doubles with PCI Express® 3.0
- Integrating PCIe3 onto the processor itself reduces latency
- What Intel calls "Direct Data I/O" changes the I/O paradigm for servers
  - The Ethernet controller talks directly with the processor's last-level cache
  - No longer is Ethernet traffic forced to detour through main memory on ingress and egress
- This new server platform has the I/O headroom to run up to 250 Gbps of Ethernet I/O per processor
  - More than a factor of three more headroom than any preceding server platform

# Flexible LOM Encourages Upgrades To 10GbE

SNIA
Advancing storage & information technology

- Similar to blade mezzanine cards, flexible LOM cards can be configured by OEMs at the time of order
- OEMs each have their own names for this form factor

| *"Classic" LOM* | *"Flexible" LOM* | *PCI Express® NIC* |
|---|---|---|
| **10GBASE-T LAN-on-Motherboard (LOM)** | **Custom PCI Express 10GBASE-T Converged Network Adapter** | **Standard PCI Express 10GBASE-T Converged Network Adapter** |

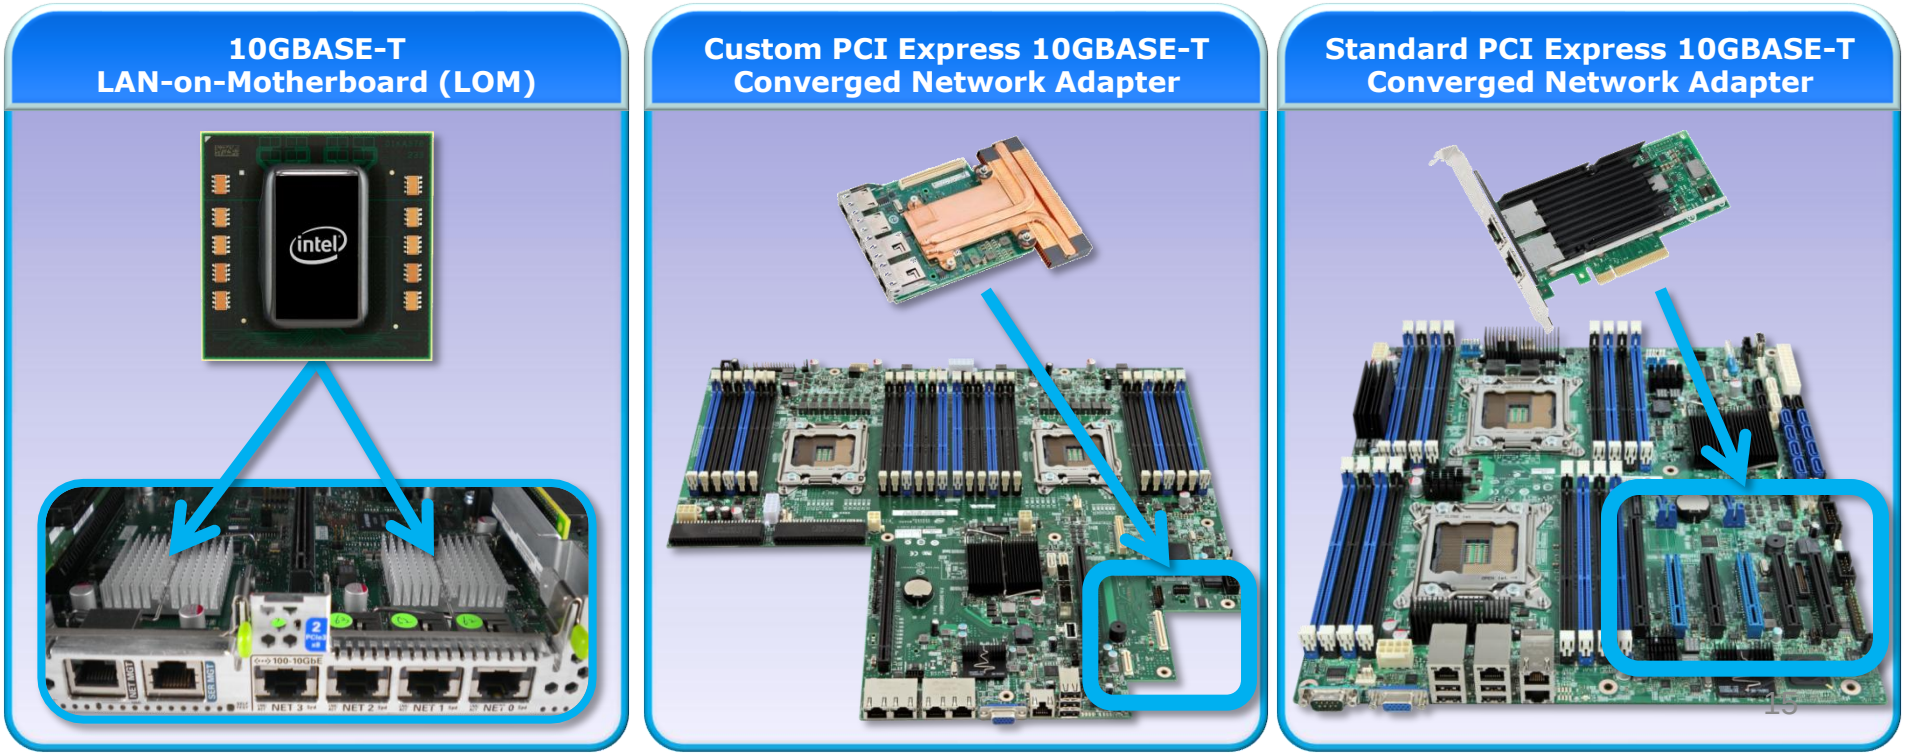

# 10GbE Brings Remote Data Memory Access (RDMA) Technology To The Data Center



- Remote Direct Memory Access (RDMA) is a technology that allows an application to transfer data directly to/from the data space of another application while bypassing the OS kernel
  - Permits low latency data transfers with low packet jitter, low CPU utilization, and traffic segregated into prioritized classes
- Two RDMA-over-Ethernet technologies being deployed today over 10GbE, iWARP (internet Wide-Area RDMA Protocol) and RoCE (RDMA over Converged Ethernet)
- High-performance computing (HPC) workloads written to the Open Fabrics Alliance stack for InfiniBand can run on a 10GbE network supporting either iWARP or RoCE
- Windows Server 2012 will take advantage of RDMA capabilities to support network Microsoft's file system called "SMB Direct 3.0."

OPENFABRICS ALLIANCE

# Ethernet Virtualization For 10GbE Overview

SNIA
Advancing storage & information technology

- Multiple approaches to "carving up" 10GbE ports and assigning the slices to virtual machines have emerged
  - Demand queuing
  - PCI-SIG® Single-Root I/O Virtualization (SR-IOV)
  - OEM-driven NIC partitioning: Dell's NPAR, HP's Flex-10, IBM's vNIC
  - OEM-driven network virtualization: Cisco's FEX
- Demand queuing is supported by Microsoft and VMware
- SR-IOV is supported in Linux Xen and KVM
  - Microsoft has committed to SR-IOV support in 2012
  - VMware is showing it in their ESX 5.1 beta
- Users have multiple and increasing choices for Ethernet virtualization of 10GbE ports

# Vote #2

# VM Density Drives More I/O Throughput

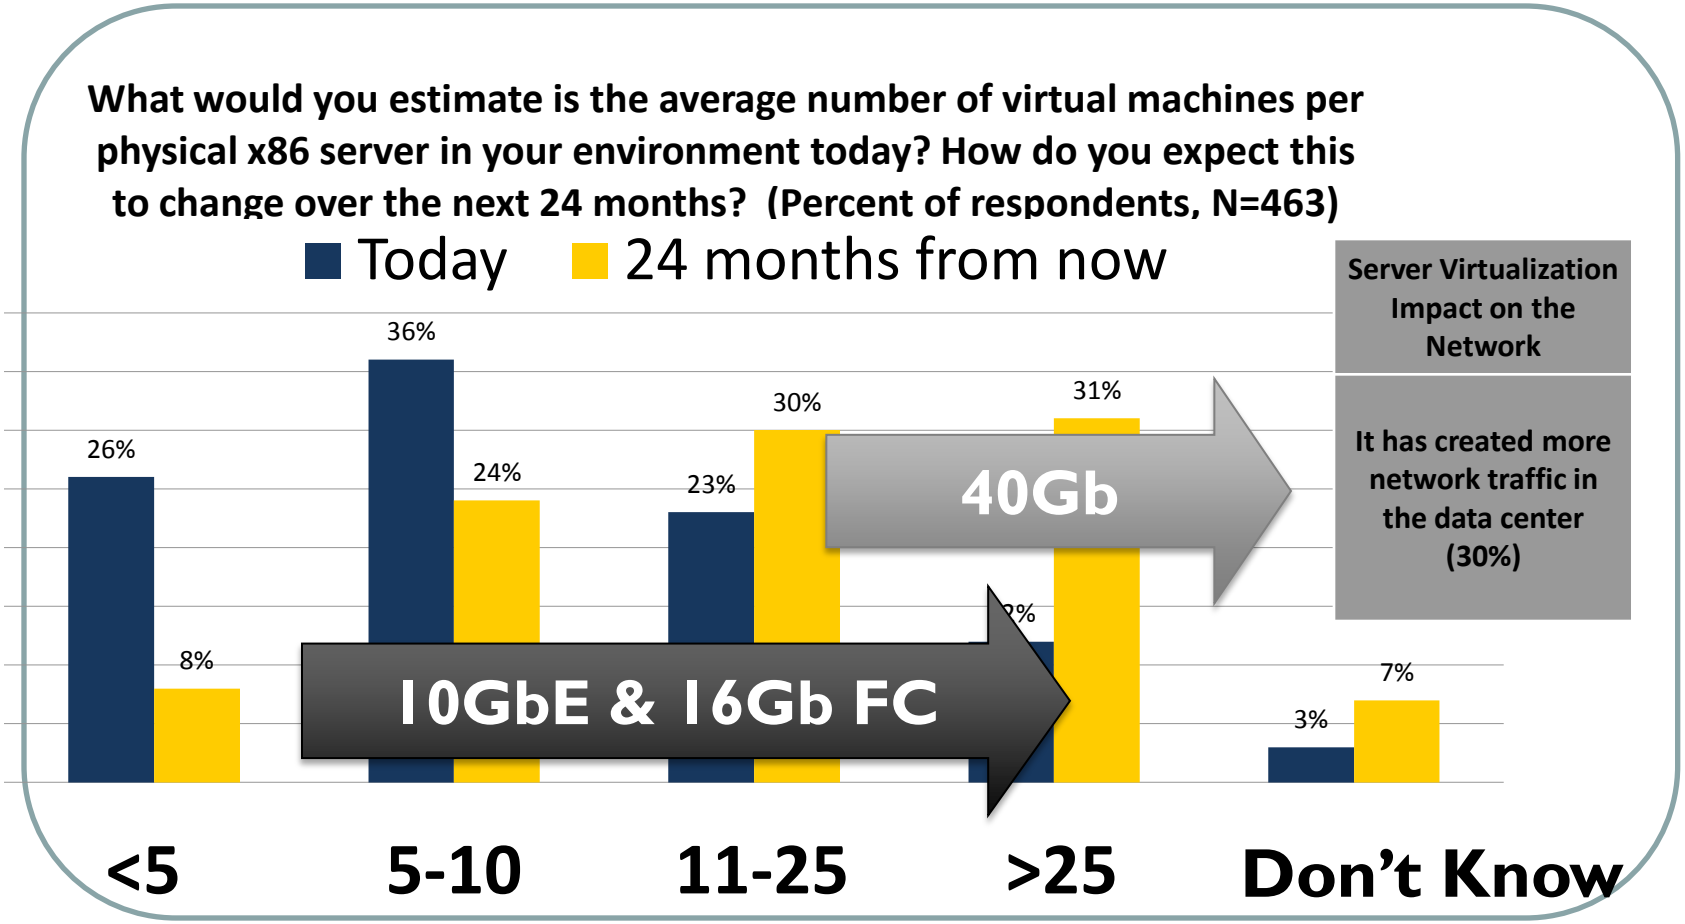

# Leverage existing infrastructure – Flexibility Network media options

SNIA
Advancing storage & information technology

- SFP+ Optical or Twin-Ax
  - iSCSI adapter with SFP cage
  - Twinaxial Cable with SFP+ modules on ends
  - Supports 5-7m distances
    - Top of Rack
- Optical versions available supporting 100m or more

- 10GBASE-T
  - Backwards compatible with Gigabit Ethernet
  - Uses Cat6/Cat6A/Cat7 unshielded twisted pair cabling with RJ-45 connectors
  - Supports 100m distances
    - Top of Rack, End of Row (EOR)

**Expensive** → **Less Expensive**

**Optical**
**Expensive optical transceivers**

**Twin-Ax**
**Expensive Cables**

**CAT6A**
**Less Expensive Cabling**
**Newer Switches**

# What is Data Center Bridging?



- **Enhancements to Ethernet**
  - Provided enhanced QoS support to Ethernet
  - **10GbE Technology today…**
- **What constitutes DCB Standards?**
  - **PFC** aka Priority based Flow Control (802.1Qbb)
  - **ETS** aka Enhanced Transmission Selection (802.1Qaz)
  - **QCN** aka Congestion Notification (802.1Qau)
  - **DCBX** aka Data Center Bridging capability eXchange
    - LLDP vs. DCBX
      - LLDP: Primarily a link level information exchange protocol
      - DCBX: Neighbors can configure parameters based on info exchange and state machine

# Data Center Bridging



- Data Center Bridging Capabilities eXchange Protocol (DCBX)
  - Supports centralized configuration of DCB and related protocols
  - Initiated by endpoints (hosts/storage) configure themselves
- Enhanced Transmission Selection (ETS)
  - Provides priority groups with bandwidth controls
- Priority Flow Control (PFC)
  - Enable lossless Ethernet operation to provide deterministic performance
- What's needed
  - HBA or Converged Network Adapter supporting DCB
  - 10G Ethernet Switch supporting DCB
  - 10G Storage Array supporting Ethernet storage over DCB

# DCB Components

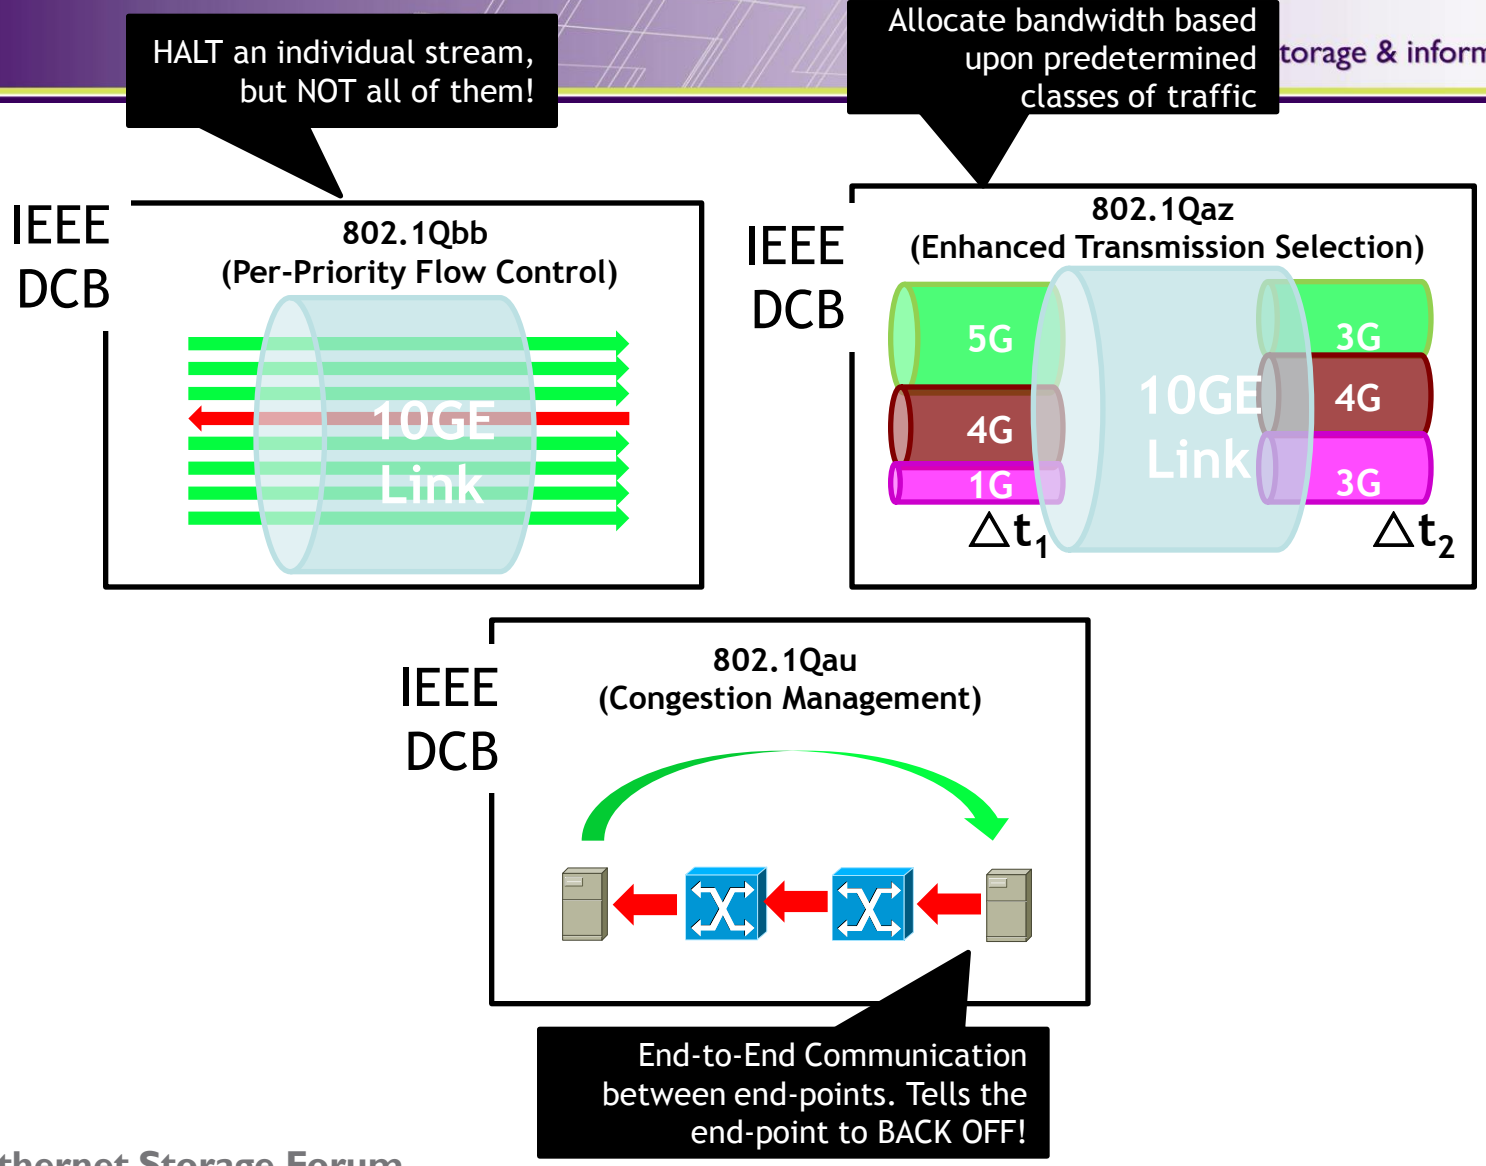

# DCB & 10GbE



- Enables convergence of LAN and SAN traffic
  - DCB enables mixed traffic
- Deterministic performance in converged, mixed traffic environment
  - Lossless Ethernet using Priority Flow Control enables support
- Ease of use
  - Centralized configuration via switch with DCBX support
    - Hosts, storage announce capabilities
    - Switch enables PFC, Jumbo-frames, other parameters via DCBX
- Easy server integration
  - Integrated LOM & Modular servers with DCB capabilities

# Ethernet Unifies the Data Center



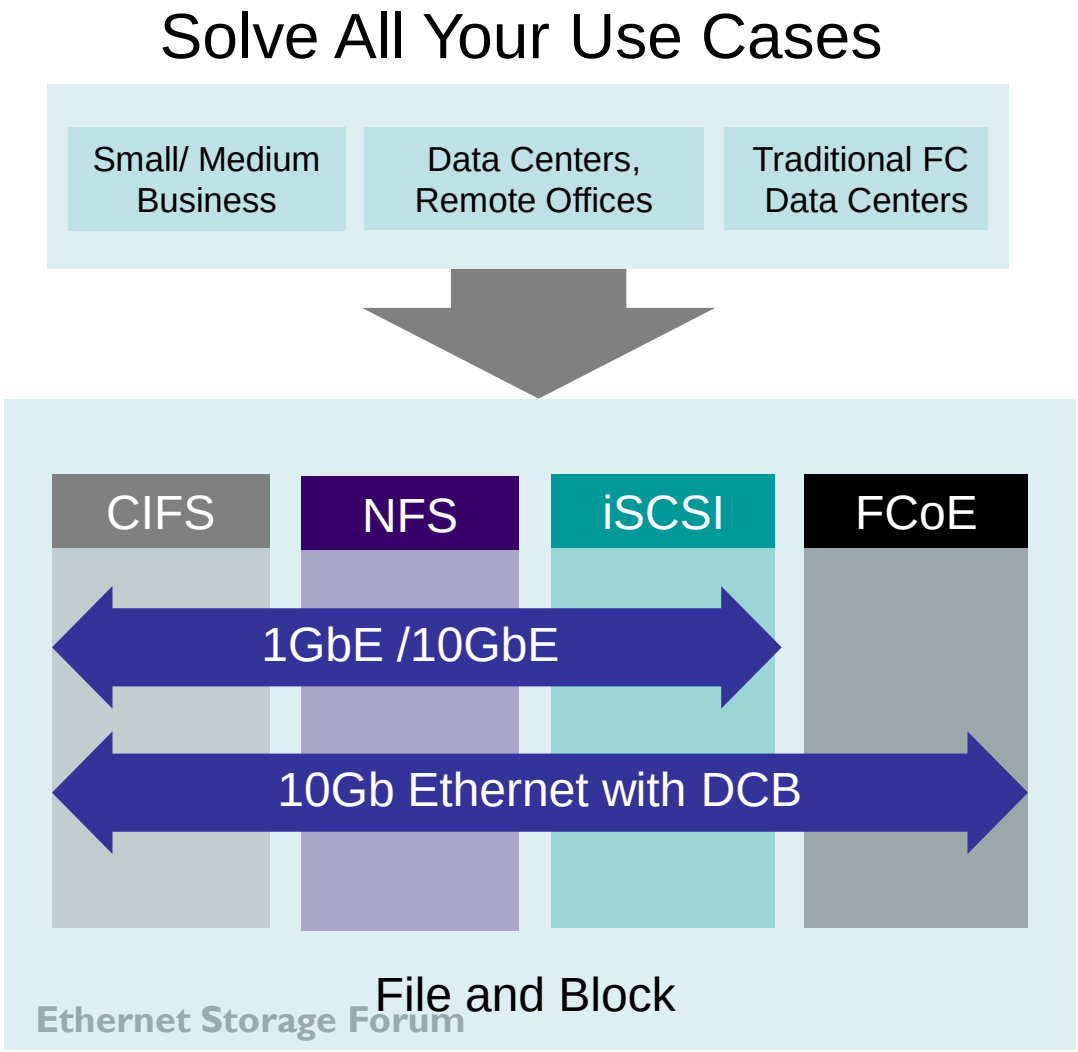


- Increased asset and storage utilization
- Simplified storage and data management
- Improved flexibility and business agility
- Reduced costs through consolidation
- Improved storage and network efficiencies

# 10GbE Adoption

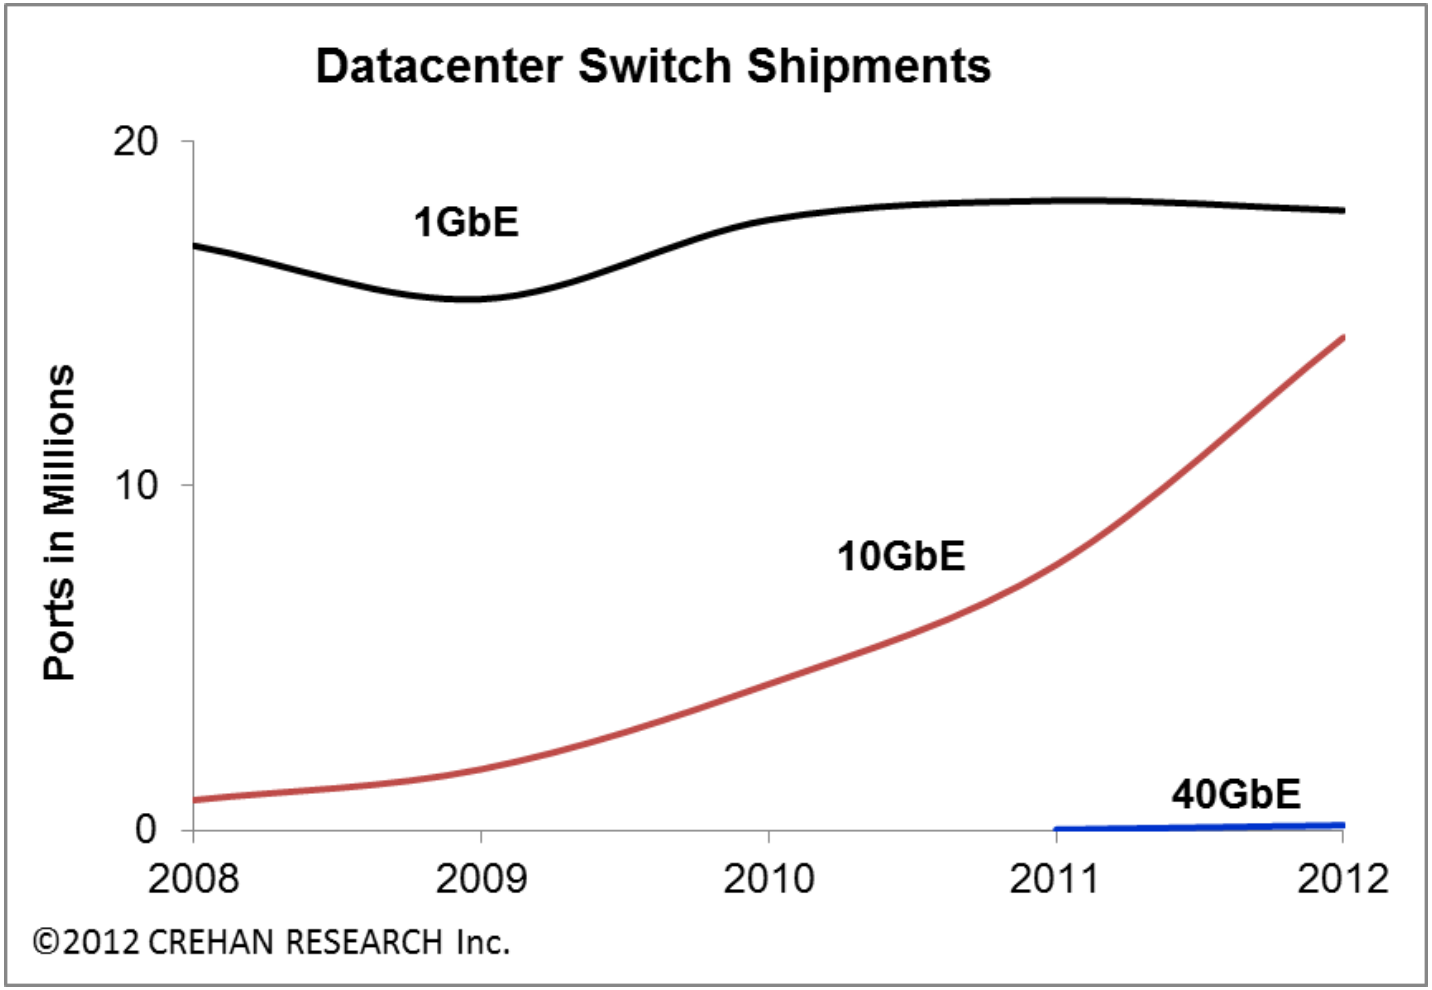

# Future of 1GbE / Fast Ethernet



- Data Center
  - Management networks for storage, servers, switches, etc.
  - Low demand storage applications such as print servers, basic file services, and active directories
  - Low demand or medium demand block storage, such as iSCSI storage for email applications
- Outside of the data center
  - Voice over IP (each desk requires wired connection)
  - Video surveillance
  - Virtual desktops (VDI)
  - And general client networking
  - Consumer

# Where Do We Go From Here?



- Data intensive applications are relentlessly driving the need for higher-speed connectivity
  - Cloud computing, mobile networking, high-speed consumer broadband
- 10GbE adoption is paving the way for 40G interfaces in the aggregation layer and 100G uplinks at the core layer
- 40GbE is the next logical step in the evolution of the data network
  - Forecasters expect 40GbE to begin significant edge server deployment in 2015-16

# Vote #3

# Questions?

SNIA
Advancing storage & information technology

- How to contact us:
  - Jason Blosil – jason.blosil@netapp.com
  - David Fair – david.l.fair@intel.com
  - Gary Gumanow – gary_gumanow@dell.com

- Full Q&A session from this Webcast will be posted on the SNIA-ESF Blog