

10 Gigabit Ethernet (10GbE) Storage Networking

Wednesday, November 9, 2005

Saqib Jang
Principal and Founder,
Margalla Communications
saqibj@margallacomm.com

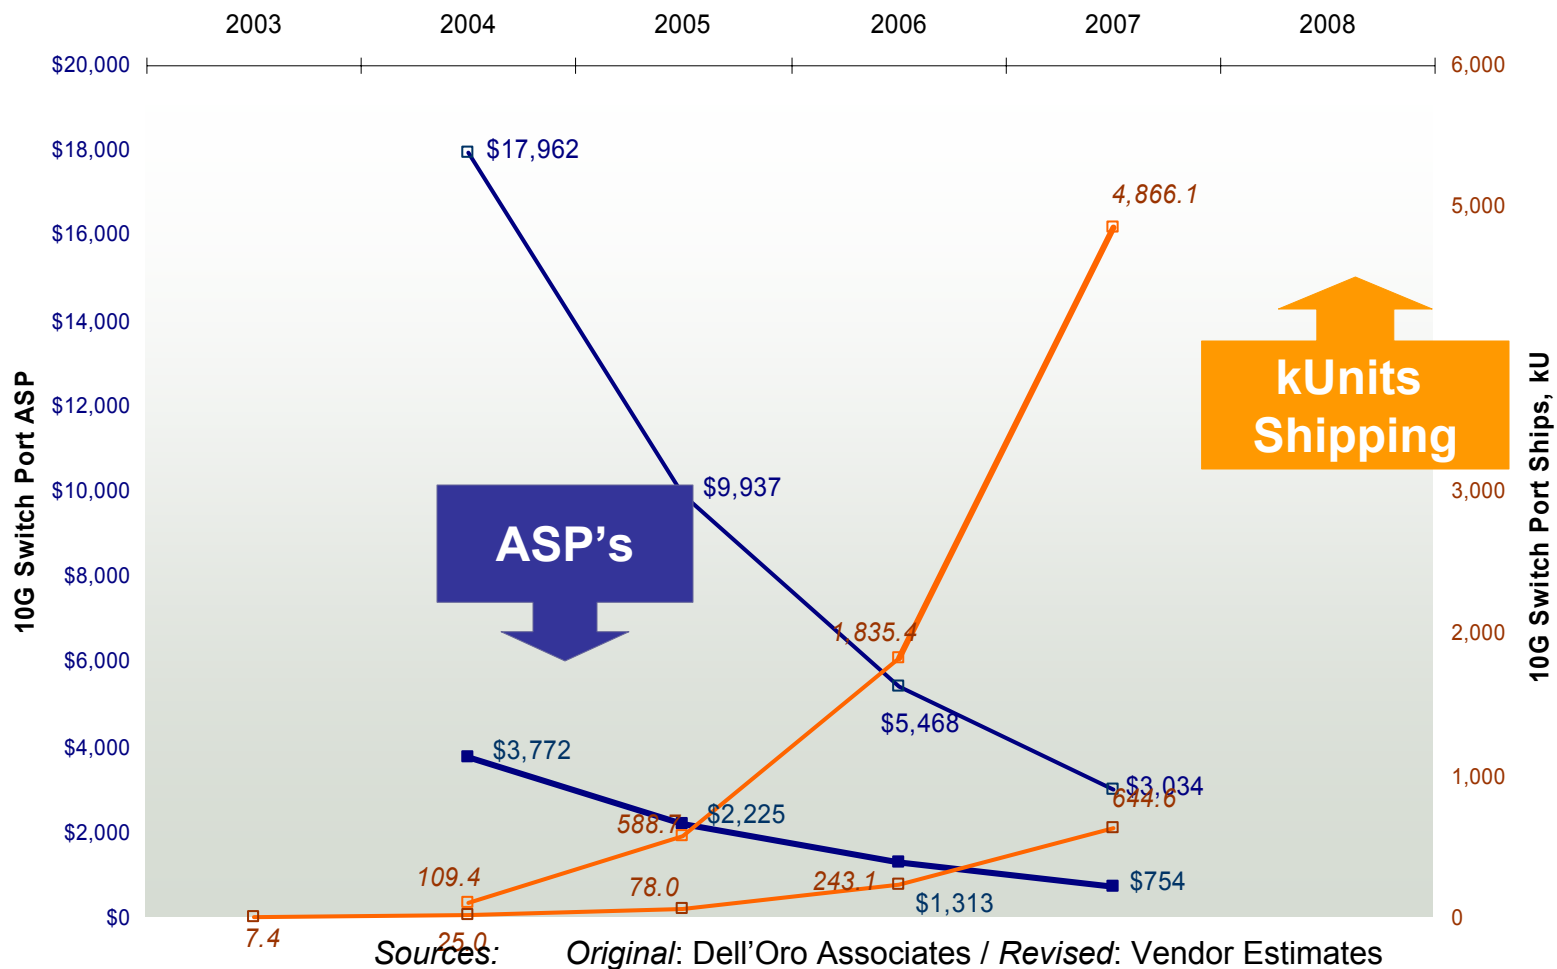
Agenda

- **10GbE Market Overview**
- **10GbE iSCSI Applications**
- **10GbE iSCSI Acceleration**
- **Summary**
- **Q&A**

10GbE Market Drivers

Criteria	Market Drivers / Enablers
Units growth	<ul style="list-style-type: none"> • 3x growth in the 10GbE NIC volumes in 2006 (synergy report) • 3x volume growth in 10GbE switch market in 2006 (Dell'Oro) • Quake has shipped its 200,000th 10G SerDes • iSCSI market growing at 50% per quarter • 58M GbE ports shipped in '05 (about 1M are multi-port NICs)
Infrastructure	<ul style="list-style-type: none"> • High-density 10GbE optical switches available now • High-density 10GbE CX4 switches available in 4Q05 (\$600/port)
Prices	<ul style="list-style-type: none"> • XFP over 12 months: \$500 to \$250; look to \$150 • CX4 switch port at \$700/port list price now • 10G CX4 RNIC expected to be at \$500/port by 1H06
Standards	<ul style="list-style-type: none"> • 10G CX4 (copper media) introduced and shipping • 10G-baseT silicon expected by early '06

10GbE Is Beating Forecasts



With CX4, HBA prices are at 2007 levels

10GbE Technology Roadmap

Bus Technologies

PCI-X 1.0

PCI-X 2.0

PCI-Express

Network Interfaces

Fiber Optics

10GBASE-CX4 Copper

10GBASE-T Copper

Protocol Offload

TCP/IP for Enterprise Data Center Networking

iSCSI for Enterprise Data Center Storage

RDMA for Clustering

Operating Systems

Linux TOE

Windows Chimney TOE

Proprietary TOE

Form Factors

Stand-Up Adapters

Blade Servers & Storage LOM

Server LOM

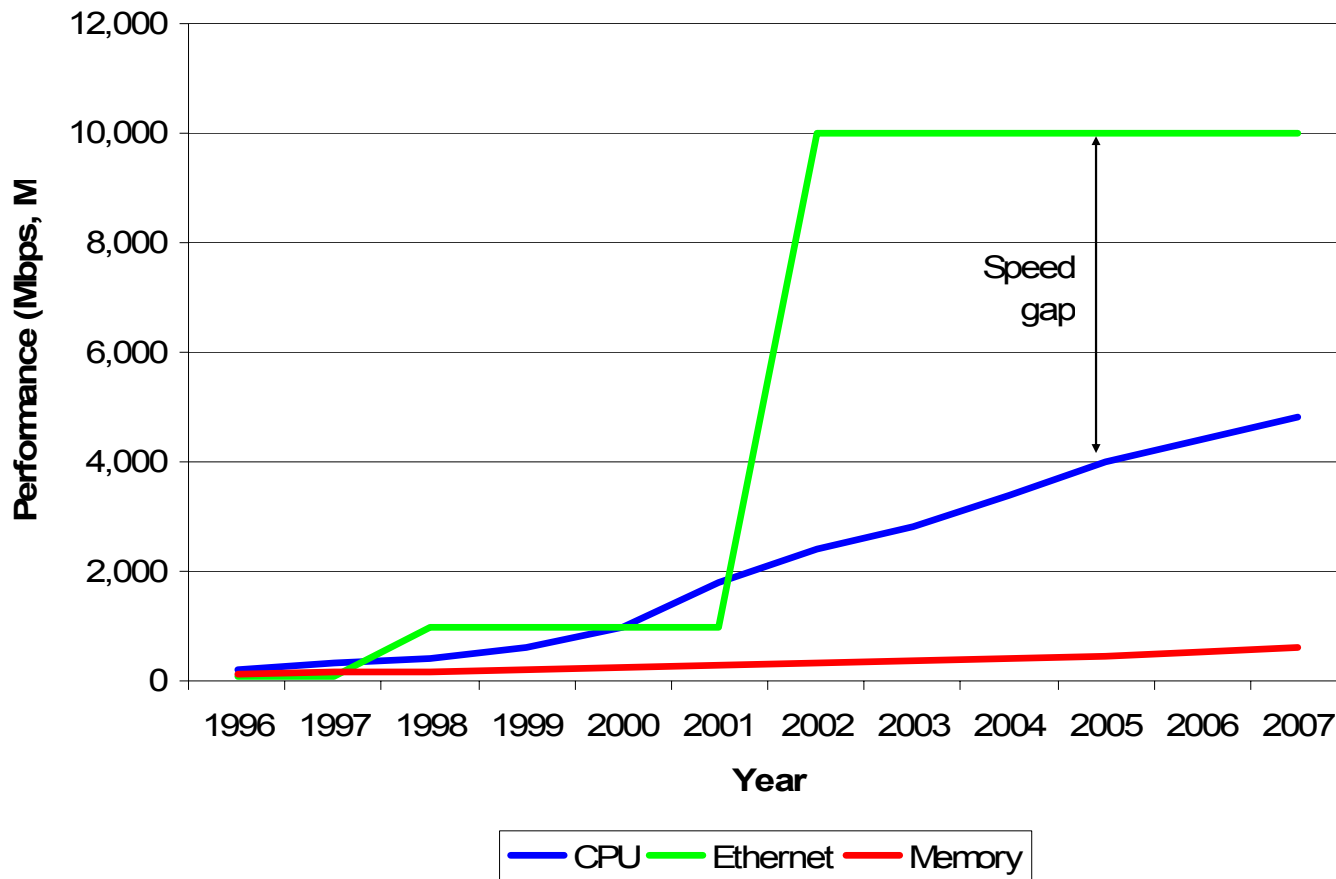
2005

2006

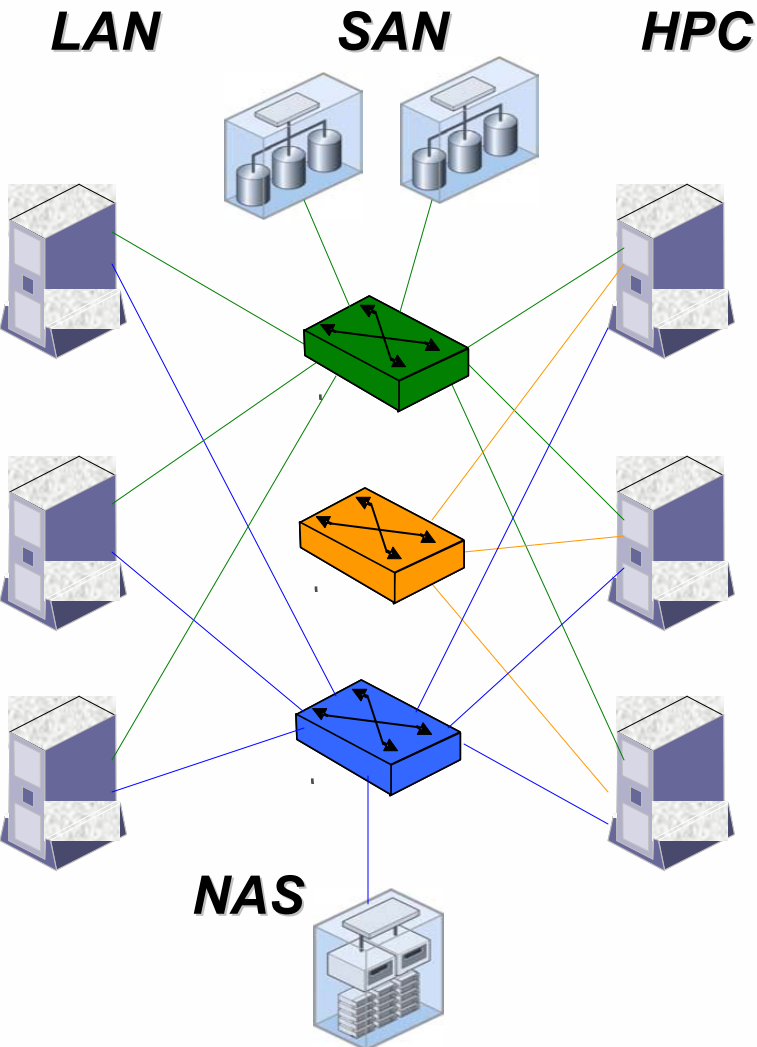
2007



Bridging the Network/System Speed Gap



Today's Data Center



- **Multiple networks, different technologies, redundant equipment, increased capital costs**
- **Gateways & bridges required to interconnect disparate networks and incompatible speeds**
- **Staffing skill sets not leveraged, additional training required**
- **Complex network management environment, incompatible tools**
- **Vendor dependency, increased supplier base, increased operational support costs**

The New Converged Data Center

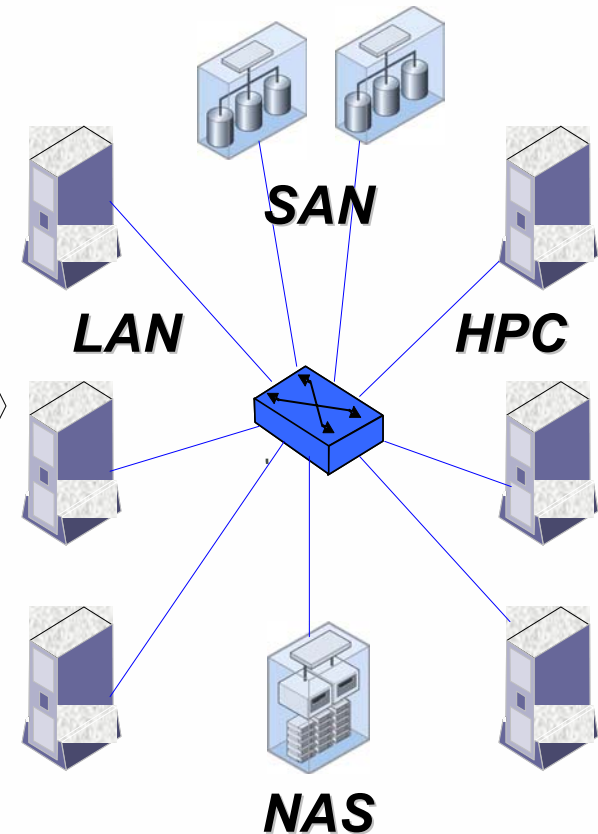
- **10GbE & Offload Lower TCO**
- **More performance from fewer servers**
- **Minimizes software licenses**
- **Data center consolidation**
- **Leverages staffing skills & tools**

Ethernet Everywhere – one technology, one network

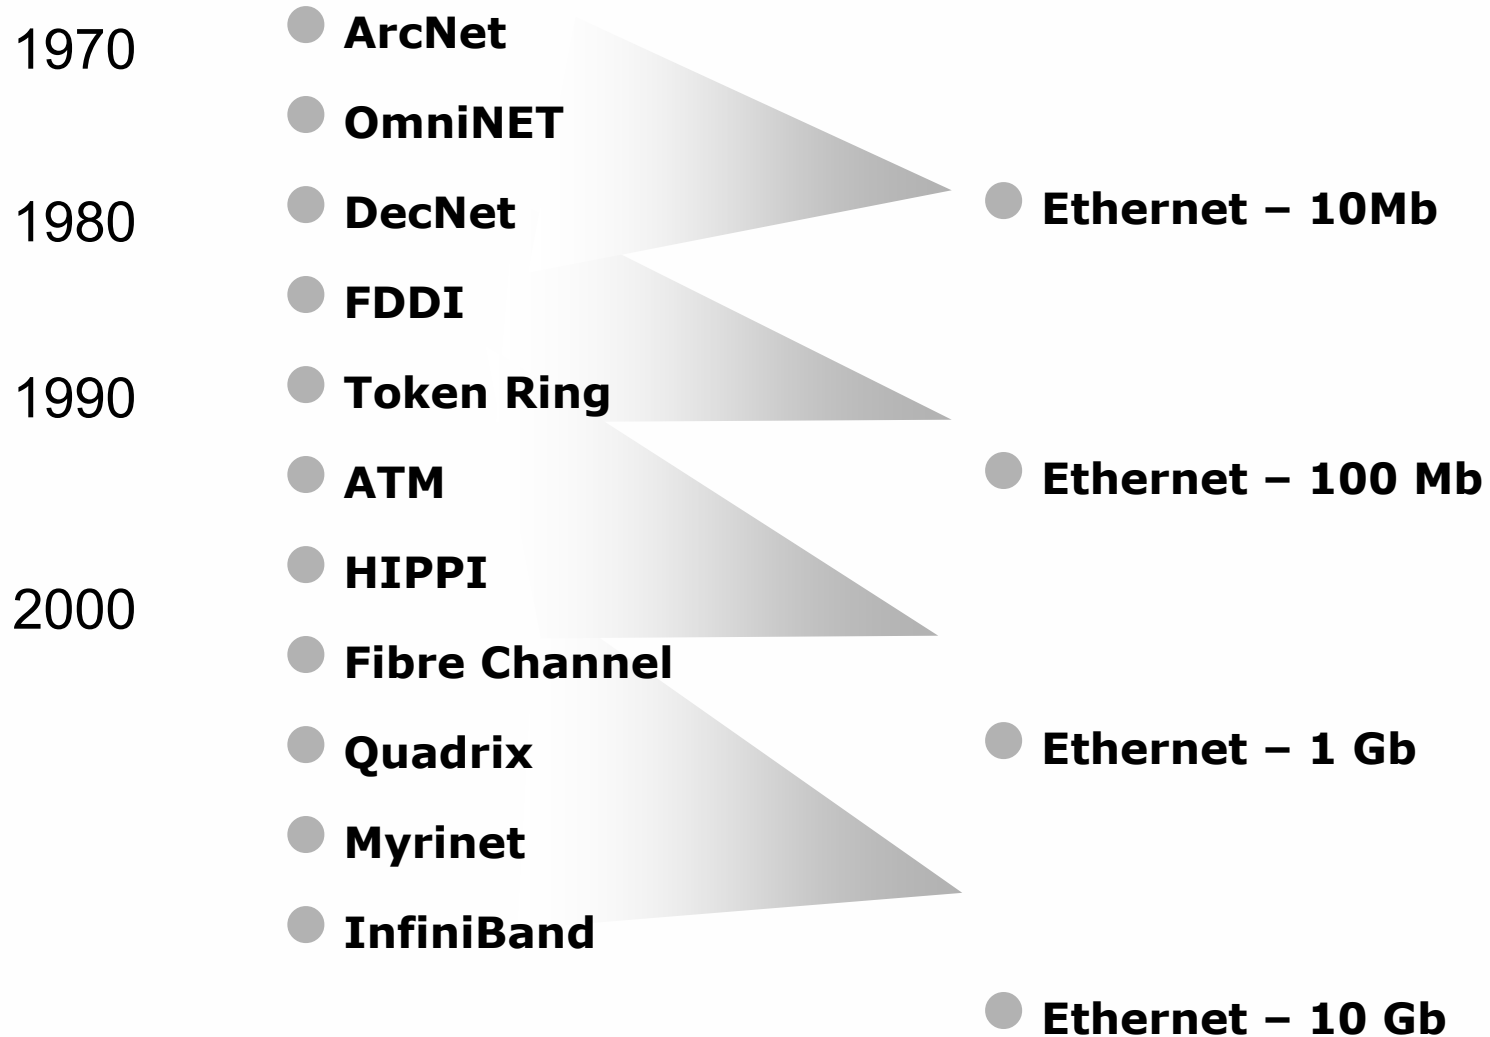
10Gb Ethernet

- **Beneficiaries of 10GbE & Offload**
- **High bandwidth apps**
- **Low latency apps**
- **Storage apps**
- **Faster backup and recovery**

Unified Network



Staying Power: Ethernet Just Gets Better



By a show of hands, which applications do you consider to be important drivers for 10GbE?

- 1. Datacenter backup and disaster recovery**
- 2. Expanded storage consolidation fan-out**
- 3. NAS/SAN consolidation**
- 4. Support of high-performance applications (e.g. video, HPC)**

10Gb Ethernet & iSCSI

- **Highest-performance infrastructure for SANs**

Highest bandwidth

- **Twice as fast as 4Gb FC**
- **Not only eliminates traditional performance disadvantage of iSCSI...**
- **But enables iSCSI to surpass Fibre Channel**

Ease and flexibility of IP

- **Same administration as Gigabit Ethernet**
- **Full compatibility across vendors & Ethernet speeds**
- **10Gb to iSCSI storage, wirespeed Gigabit to multiple hosts**

More Advantages


- **Less software to buy and learn**
 - **Software initiator stacks for application servers**
 - **IP multipathing built-in to Windows and Linux**
 - **Use existing IP topologies, administration**
 - **More standards-based security options than FC such as CHAP, IPSec**

- **Lowest cost/gigabit for switch port**
 - **4 Gb Fibre Channel: \$125/Gb**
 - **10 Gb iSCSI: \$60/Gb**
 - **Comparable price/port on switches: \$500-600/port**
 - **QLogic Sanbox 5602 v. HP ProCurve 6400cl**

Enabling iSCSI in the Enterprise

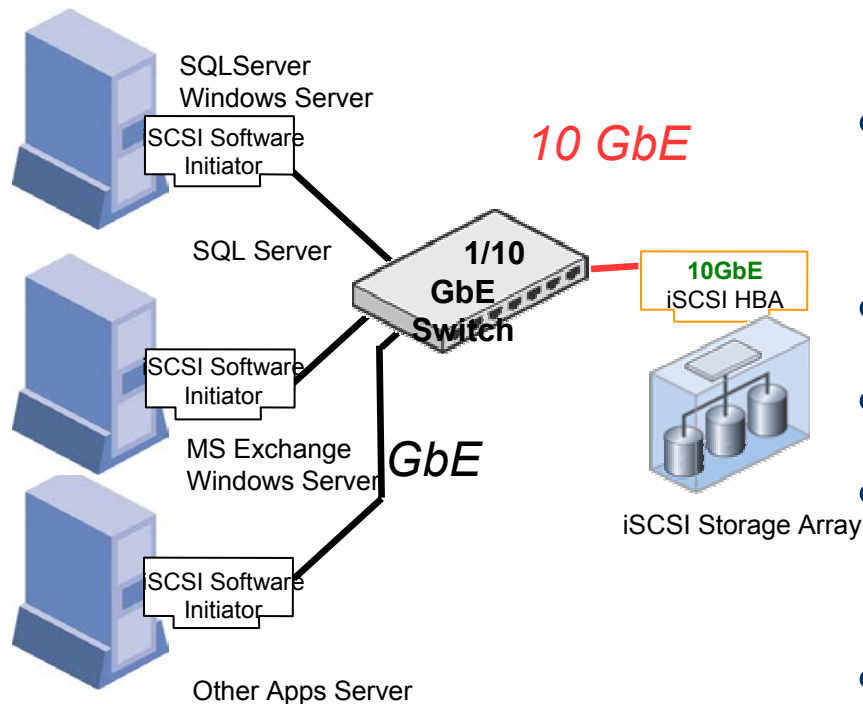
- **iSCSI is more proven than ever**
 - **Deployment ramp of IP SANs exceeding 200 TB**
 - **Mission-critical deployments abound**
 - **Broad iSCSI availability across OSES including Windows, Linux, HP-UX, Solaris, AIX, Netware**
- **iSCSI on 10 Gb beats Fibre Channel**
 - **Superior performance vs. 4G FC**
 - **Compatibility and simplicity of IP**
 - **10 Gb is less expensive and getting cheaper**

10GbE uniquely suited for Storage




- **Free client scalability**
 - **Free software initiators**
 - **Free GbE ports with ALL servers**
 - **HBAs not required for GbE initiators**
 - **GbE speed adequate for servers**
- **Similar Target costs to IB, FC**
- **No change to existing Apps**
- **Little change to infrastructure**

10GbE is cheaper than IB or FC



- **FC/IB need special switch chassis**
 - **Must have Ethernet switch anyway**
- **FC/IB need special HBA/HCA**
 - **Grows with server count**
- **FC software is not free**
- **FC/IB need special training**
- **FC/IB do not allow convergence**
 - **Have to stock/purchase different HW**
- **FC/IB do not allow consolidation**
 - **More wiring**
 - **More power**
- **Ethernet has ultimate ease-of-use**


Fabric Comparative Analysis



• Pricing from CDW (Oct '05)


	10GbE	FC	IB-HCA	IB-LOM	Comments
Server	Free	\$450	\$300	\$80	Per client
Target	\$795	\$250	\$300	\$300	Per Target Array
Cable	\$80	\$20	\$80	\$80	Per Array
Switch Port	\$600	\$150	\$300	\$300	Per Array
Switch Chassis	N/A	\$6000	\$15,000	\$15,000	Estimates
Software	Free	?	?	?	
Training	Min.	New	New	New	
Fabric Total	\$2,950	\$40,800	\$49,000	\$38,440	48 clients + 2 Arrays

Enterprise Applications




- **Consolidate storage without compromise**
- **Use Gigabit switches with 10Gb ports**
- **10 Gb ports provide iSCSI storage connection**
- **Deliver wirespeed Gigabit SAN to 6-8 servers**
- **Maximize performance of file servers (NAS), email, web, ERP, and business applications**

Performance-Critical Applications



- **Rich media and HPC**
- **10 Gb enables greater storage consolidation**
- **Serve dozens of workstations instead of a couple from one central iSCSI storage array**
- **Combine with SAN file systems for a higher-performance, lower-cost solution than “high-end” NAS appliances**

High-Availability and Business Continuity for IP SANs




- **10G iSCSI allows superior price/performance SAN high-availability vs. proprietary (e.g. ESCON, FC) approaches**
- **10 GbE bandwidth enables high-performance asynchronous replication, remote snapshots, even synchronous mirroring**
- **10 GbE iSCSI connections increase data protection by increasing performance for real-time mirroring over distance**

● By a show of hands, what is the most important performance metric for 10GbE iSCSI?


1. **IOPS**
2. **Throughput**
3. **Latency**
4. **CPU utilization**

10GbE Network Connectivity Options

10GbE Stateless Offload Adapters

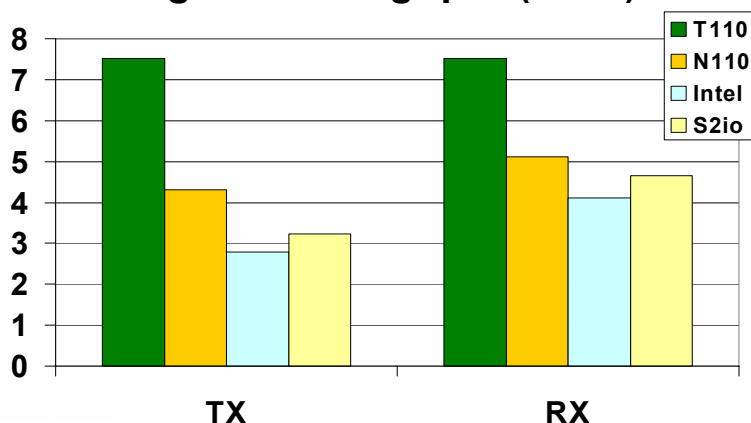


10GbE Protocol Offload Adapters
Offloaded Protocols: TCP, iSCSI, RDMA

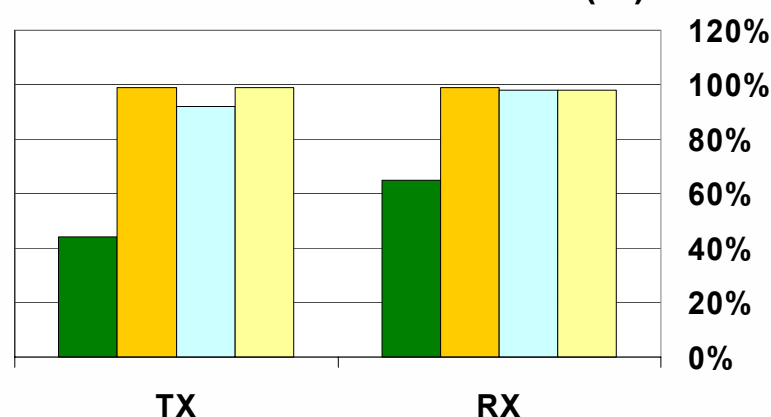


10GbE TOE vs. 10GbE NIC Network Performance

Higher Throughput (Gb/s)



Lower CPU Utilization (%)



Source: VeriTest, Inc.

Test tool: netperf

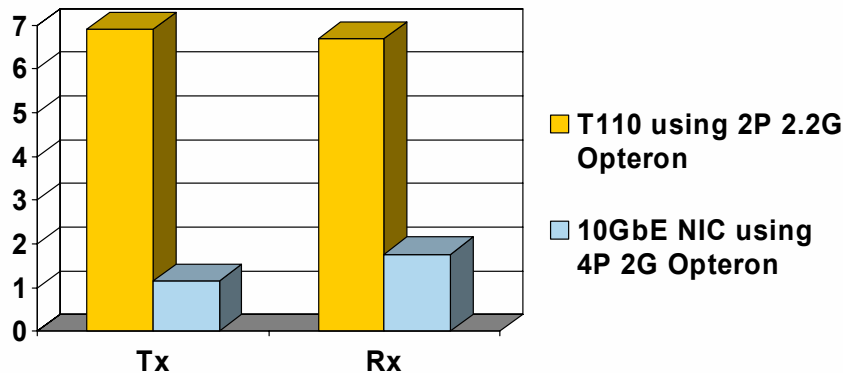
Test configuration: 2 systems connected through 10GbE switch running single TCP channel with 1500-byte Ethernet frames

System configuration: AMD Opteron 248 2.2GHz uniprocessor running Linux kernel 2.6.6


- **Performed May 2004 by VeriTest, Inc.**
- **10GbE TOE delivers ~2X network throughput vs. 10GbE NICs**
- **10GbE TOE shows ~1/2X CPU utilization vs. 10GbE NICs**
- **Net-net: 10GbE TOE show ~4X network efficiency vs. NICs**
- **No Jumbos, standard 1500B Ethernet frames only**

10GbE TOE vs. 10GbE NIC iSCSI Performance

Throughput (Gbps)



Average CPU Utilization (x2.2 Ghz Opteron)



Sources: T110 iSCSI Performance Analysis by Veritest and iSCSI Performance Analysis WP Published by Neterion
Note: Charts show performance at 4KB I/O size; iSCSI applications are transactional in nature using 2-4KB I/O sizes

- **10GbE TIE delivers ~4X iSCSI network throughput vs. 10GbE NIC**
- **10GbE TOE shows ~1/4X CPU utilization vs. 10GbE NIC**
- **10GbE shows ~16X iSCSI processing efficiency vs. NICs**

Key Takeaways

- **10GbE iSCSI is ready for prime time**
- **Protocol offload for 10GbE iSCSI**
- **10GbE iSCSI represents the lowest TCO and price/performance of any similar fabric**
- **10GbE is the ideal data center 'unified wire'**
 - **Ethernet will once again consume all Ethernet fabrics**

● **How soon and at what speeds do you anticipate deploying an iSCSI Target Array?**

1. 2x1GbE in 2006 or 2007

2. 4x1GbE in 2006

3. 4x1GbE in 2007

4. 10GbE in 2006

5. 10GbE in 2007

6. Not sure

Q & A

- *We are now taking your questions!*

