Practical Innovation

A White Paper from Neterion, SBE and Force10 Networks

FULL LINE RATE
NEW BENCHMARKS
iSCSI OVER 10 GbE

Xframe® II

The Storage Area Network Goes Global

Combining the reliability of iSCSI with Full Error Recovery and the performance of 10 Gigabit Ethernet

October 2006

Neterion™

Summary

As the demand for bandwidth grows, so does the reach of networks. Users need to access data globally, from multiple locations around the world, quickly and transparently without breaking budgets. Data centers locked into servicing only local users unnecessarily constrain growth and expansion. These five factors — cost, bandwidth, reach, latency, and ease-of-access — are the driving metrics behind today's storage area networks (SAN).

Fibre Channel has served the early SAN market well. However, iSCSI and Ethernet technology have reached a maturity where replacing Fibre Channel with iSCSI is a compelling proposition. IETF RFC 3720 introduces complete error recovery mechanisms to iSCSI, providing robust and highly-available system operation across geographically diverse enterprise networks. With immediate availability of 10 Gigabit Ethernet adapters and switches, SANs can push to higher performance levels than ever before, as mass adoption drives costs down. This whitepaper will demonstrate, using off-the-shelf hardware and software components available today, how iSCSI can outperform comparable Fibre Channel implementations by up to 6x.

While virtually a single network, today's enterprise networks are often physically located in multiple locations. Compounded by the rising need of mobile users to access the network from anywhere in the world, storage area networks (SAN) are being called upon to provide access on a global basis. Providing global access requires a network that can minimize latency successfully while increasing available bandwidth and decreasing system cost without creating a network that is difficult to manage.

Many SAN deployments today are based on Fibre Channel technology. However, the distance limitations of Fibre Channel negatively impact network cost and make it impossible to cost effectively globally scale access to an organization's storage area network. Fibre Channel was designed as a campus-wide network, not to exceed 100 miles, and there are many difficult barriers to transporting Fibre Channel over the WAN because of Fibre Channel's sensitivity to delay. Additionally, while there are established protocols such as IFCP (Internet Fibre Channel Protocol) and FCoIP (Fibre Channel over IP) to allow Fibre Channel to extend further, Fibre Channel is a closed network and these alternatives are piecemeal approaches that require specialized hardware and introduce additional complexity and vulnerabilities into the network.

The Internet Small Computer Systems Interface (iSCSI) was created to extend the reach of storage area networks by enabling SAN functionality over the IP network. iSCSI uses SCSI over TCP/IP, enabling any requesting node on the IP network (initiator) to contact any remote dedicated server (target) and perform block I/O on it just as it would using a local hard drive. Because it is a native IP protocol, iSCSI has no distance limitations, can utilize existing network infrastructure, does not require specialized operator or administrator training, and can leverage the vast economies of scale of the enormous Ethernet market.

A slow start

Early deployments of iSCSI were, quite frankly, incomplete. None of the early vendors in the space offered a complete iSCSI stack with the high quality transport and fault tolerance that is required by Enterprise storage managers. Low cost was promoted to address the buying freeze triggered by the recession during which iSCSI evolved in order to enable it to compete against the market glut of incumbent technologies. Expectations for performance were set low and these first devices were sold primarily to unsophisticated users in the small-to-medium business (SMB) space. High-value qualities, such as standards compliance, interoperability, reliability, security, and integration were secondary. As point-in-fact, the IETF (Internet Engineering Task Force) only adopted the complete protocol in February 2003, yet many vendors were selling incomplete iSCSI enabled devices long before this. At the same time, Fibre Channel offered — at a premium price — predictable and reliable performance. As a consequence, early iSCSI product deployments worked only marginally well and failed to address true Enterprise storage needs, prompting Enterprise storage managers to adopt a wait-and-see stance. So what's changed with iSCSI?

10 Gigabit Ethernet and Error Recovery Level Two (ERL 2)

After a strong 2005, 10 Gigabit Ethernet is poised for steep growth in 2006, with the majority of IT managers in Fortune 1000 companies deploying 10 GbE at the network infrastructure level (switches and routers), and starting to connect the fast-speed line to their servers and storage systems. The Dell'Oro Group estimates that over 140 thousand 10 GbE switch ports were shipped in 2005, and this number will grow to nearly 500K in the year 2006.

Many applications can benefit instantly from upgrading the end-point connection to 10 GbE, like server consolidation and virtualization, backup acceleration, large size image & video transfers over the network, etc. Many vertical markets like Medical Imaging, Oil & Gas, Media-Entertainment, Military and Research, to name a few, are faced with exponentially increasing data volumes. And with the proliferation of Gigabit Ethernet to the desktop, the demand for bandwidth is only increasing.

With the completion of the entire iSCSI protocol as defined by IETF RFC 3720, iSCSI now has the full error recovery feature set required by enterprise storage managers (see Figure 1). While this standard has yet to be widely embraced by the industry, the features and reliability provided by completing the upper Error Recovery Levels of the iSCSI protocol will be essential to developing Enterprise market confidence in the reliability of iSCSI (see "5 key reasons to consider iSCSI with ERL2 and 10 Gigabit Ethernet").

When it was 1 GbE iSCSI with zero Error Recovery capabilities against 2 Gbps Fibre Channel, even with a cost advantage the poor reliability of early iSCSI implementations didn't put up a great fight. However, with the maturity and completion of iSCSI with full error recovery compounded by 10 GbE's throughput, availability, cost, distance, and economy of scale advantages, iSCSI with ERL 2 presents a much more serious challenge to Fibre Channel:

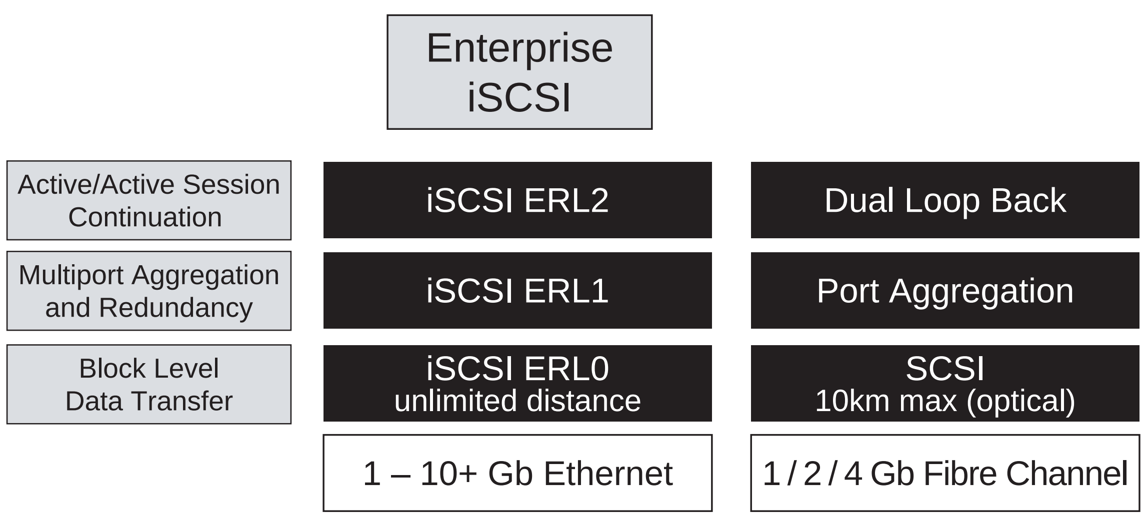

Figure 1: iSCSI ERL2

- Ethernet offers seamless integration with the LAN, MAN, and WAN, enabling native transport of data anywhere in the world

- The ubiquitous deployment of Ethernet across a wide range of applications provides economies of volume simply not possible with the limited deployment of Fibre Channel

- Ethernet is a robust technology with proven protocols, management tools, and trained experts that are the foundation of 85% of networks worldwide

- 10 GbE will continue to benefit from emerging Ethernet advancements, such as hardware-based iWARP support (RDMA over TCP/IP), which will add Remote Direct Memory Access to optimize performance and reduce CPU utilization

- Low-level compatibility with major operating systems eliminates the need to implement OS modifications when upgrading networks

- True quality of service (QoS) functionality enables traffic to be consistently classified, prioritized, and queued at line rate

- Single- and multi-mode fiber support gives network administrators greater throughput in short-reach environments and maximum flexibility through long-reach single-mode fiber in determining the location of server assets

- 10 Gigabit Ethernet coexists natively with already widely deployed Gigabit Ethernet infrastructures

- There is already a large base of available personnel trained to work with Ethernet. Their learning curve to learn the particulars of storage area networks is significantly less than that required to become an expert with Fibre Channel. Leveraging existing personnel expertise contributes to a lower total cost of ownership (TCO) for end-user customers

5 key reasons to consider iSCSI with ERL2 and 10 Gigabit Ethernet

1) Security, Reliability, Quality of Service of iSCSI ERL2 versus Fibre Channel's limited feature set
2) Performance of 10 Gigabit Ethernet versus 2 Gbps Fibre Channel
3) Cost of Ethernet vs. Fibre Channel (both from a hardware and a Total Cost of Ownership perspective)
4) Proven capabilities of Ethernet, such as unlimited distance, Server disk-less booting capability, and so on.
5) iSCSI with ERL2 enables robust disaster recovery, back-up, and replication over remote sites via existing IP WANs

History has shown that Fibre Channel equipment has been sold with a high margin, giving Fibre Channel a cost flexibility allowing it to continue to drop price to a level competitive with iSCSI. However, the majority of cost associated with networks is not the price of adapters but the cost of managing and upgrading the network over time. While the cost of adapters may be complementary between Fibre Channel and iSCSI, the salaries of trained personnel to manage them is not. In addition, the simplification that stems from a converged network on Ethernet results in great reductions in TCO (downtime, problems tracking, network management, etc.). These hidden "opportunity cost" are difficult to quantify but real for datacenter managers.

Computing the total cost of ownership of a network requires evaluation of more than just initial hardware costs. When counting capital costs, maintenance fees, system integration costs incurred with every change to the network infrastructure, ongoing support costs, physical space, port density, power, cooling, upgrades, system management, and soft costs like slow network performance and downtime due to system overloads, TCO becomes an important consideration. Annual recurring costs can range from 30-50% of initial hardware investments.

Optimized efficiency

Designed for virtualizing physical media, iSCSI is ideal for storage applications that serve up data that has already been created, such as is required for databases and streaming audio/video servers, as well as clustered computing. Without any distance limitations, iSCSI enables virtualized networks with nodes located physically around the world to appear as a single logical network. Because of its thin software layer, network administrators can turn almost any TCP-enabled device into an iSCSI initiator or target, giving near-universal access to a network while preserving existing Ethernet investments (see sidebar, "Stateless offload 10 GbE adapters make a great iSCSI solution").

Such flexibility changes the way people will design their networks. Since iSCSI can be implemented at the software layer, existing equipment can be upgraded with off-the-shelf software available today. For example, a server upgraded with a 10 Gbps Ethernet card with iSCSI ERL 2 opens access to the entire server farm, effectively converting the server farm into a SAN. In this way, iSCSI with full error recovery promotes a single IP network instead of segregating storage to its own network requiring independent equipment, applications, and staff.

Ethernet Switch/Routers are a perfect fit for iSCSI with ERL 2 at 1 or 10 Gbps

1) Unified Ethernet networks support multiple uses, saving money – both CAPEX and OPEX
2) High density switch/routers reduce the number of network devices to buy and manage
3) Non-blocking 1 GbE and 10 GbE ports optimize storage performance
4) High-availability switch/routers with redundant components improve uptime
5) Protocol-rich switch/routers connect seamlessly with multiple networks

More efficient use of available bandwidth is another optimization iSCSI offers over Fibre Channel. iSCSI, like Fibre Channel, writes and reads directly to and from the block level of a disk, thereby reducing the overhead for transporting multimedia or large engineering/ medical data files when compared to NFS systems or NAS devices. Additionally, unlike Fibre Channel where unused bandwidth in a dedicated channel goes to waste, complete iSCSI implementations with ERL 2 are able to make optimal use of bandwidth since connections automatically allocate bandwidth from connections that are currently not in use. More efficient use of available bandwidth translates directly to cost savings since more end nodes can be serviced without the need for new equipment, reducing capital and operational costs while simplifying network management.

ERL 2 iSCSI increases system reliability and adaptability as well (both iSCSI and FC transports take the same path through an operating system's file system layer.) Data centers require high-availability, and the ERL 2 feature set of iSCSI provides active/active task migration of connections from failed routes to available connections which prevents session and data loss — identical in behavior to the dual loopback featured in Fibre Channel implementations. Additionally, complicated network topologies introduce security vulnerabilities, as well as increase network management complexity. The flexibility of iSCSI with full error recovery enables networks to adapt as networks grow while maintaining complete end-to-end visibility and consolidated manageability of all data center assets.

Virtualization using iSCSI can also increase network robustness, especially as application traffic becomes less deterministic. Emerging web applications, service-oriented architectures, and clustered computing require interprocess communications between systems. Through its TCP/IP foundation, iSCSI with full error recovery enables users to allocate processes in different manners without compromising performance. When properly configured, IP SANs can be available to any authorized user at any location in the world with an Ethernet connection.

Consolidation factors

iSCSI with full error recovery also results in the consolidation of multiple networks. Many server banks go underutilized because they are the property of a single business group. In many cases, Fibre Channel server banks cannot be consolidated because of distance and interoperability issues. Because the reliability required by distributed database applications and tier-one storage managers was not available in early implementations of iSCSI, effective consolidation has been largely limited to the realm and reach of Fibre Channel.

With the zero distance limitations and flexibility of iSCSI with full error recovery, data silos can become shared resources (see Figure 2), and data centers can be combined physically — while retaining logical independence — leading to reductions in floor space, cabling, power, and other physical plant requirements as well as reduced management personnel and maintenance costs. Consolidation through virtualization also frees up scarce I/O slots to make room for future expansion.

The impact of consolidation should not be underestimated. The Meta Group, for example, estimates the cost of managing a five-chassis data center with an initial hardware cost between $750,000 and $1,200,000 to run between $252,000 and $372,000 annually, or about 1/3 the initial hardware cost. Consolidating underutilization of a five-chassis system to a single chassis drops initial hardware costs to between $150,000 and $250,000

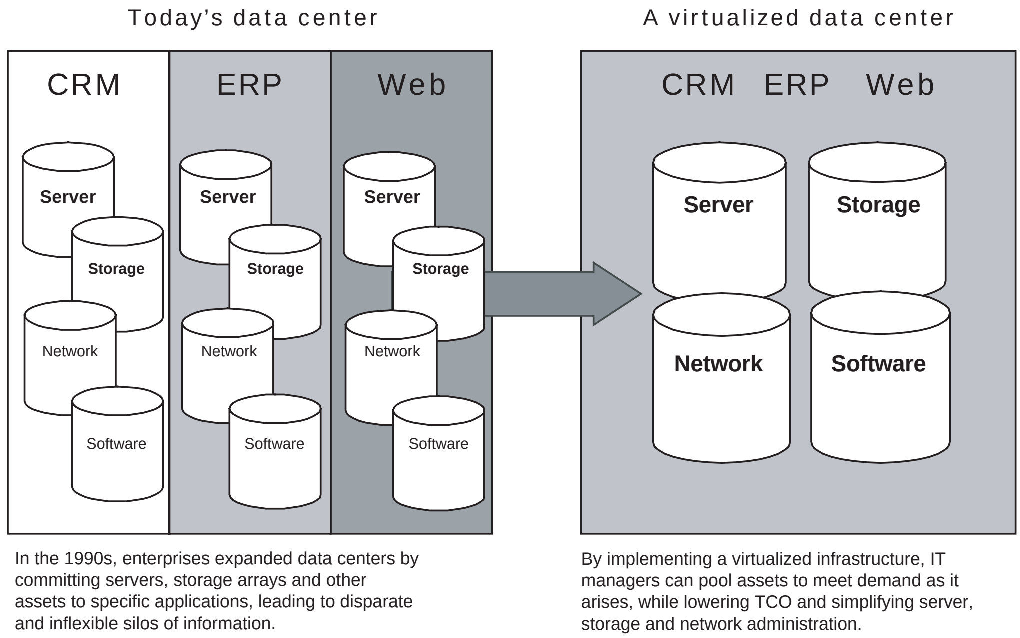

Figure 2: Maximizing Bandwidth in the Virtualized Data Center

with annual costs of $38,500 to $56,500. For global storage area networks, the cost savings of iSCSI using ERL 2 versus Fibre Channel can be tremendous.

Especially important to consider, iSCSI administrators will still be able to leverage their existing Fibre Channel investments. However, since iSCSI with full error recovery at 10 GbE offers so much more functionality and value, as well as bridging to Fibre Channel networks, new deployments will certainly favor iSCSI ERL 2. Fibre Channel is attempting to hold market share with a new 4 Gbps version of the spec, but this spec is still under development while 10 GbE is available for purchase and deployment today.

Real-world performance

iSCSI with ERL 2 brings important optimizations to storage networks. It's here, viable, and reliable, offering the kind of performance networks need today. Interoperability drops deployment costs to commodity levels and enables network administrators to select best-in-class components rather than having to purchase an entire system solution from one vendor. Today, all of the pieces are available for making deployment of ERL 2 iSCSI IP SANs a reality.

For example the S2410, a 24-port 10 GbE switch from Force10 Networks, is the world's first Ethernet device to deliver 300 nanoseconds switching latency - comparable to specialty interconnects. It interoperates seamlessly with Neterion's Xframe®, a family of 10 GbE adapters for PCI-X 1.0, 2.0 and PCI Express buses. iSCSI software from SBE provides Error Recovery Level 2, offering no single point of failure end-to-end with guaranteed quality of transport supported by multipath, session continuation, and dual loop back functionality. SBE iSCSI supports clustered file system applications, scales to support N-number of storage devices, and eliminates the need for upper-level switching at the hardware level. All of these technologies are available today and interoperate seamlessly with each other.

These market leaders came together to build a network system using these off-the-shelf devices to benchmark ERL 2 iSCSI performance at 10 Gbps (see Figure 3). Two configurations were set up for the benchmarks: first with a couple of Linux iSCSI initiators, the second with three Windows iSCSI initiators. In both configurations, the initiator systems accessed a single RAM disk storage target, running SBE's iSCSI target software for Linux (see Figure 4 for detailed hardware and software configuration). The results were measured using DiskTest and IOmeter, two standard network profiling tools available from the Linux Foundation and Microsoft.

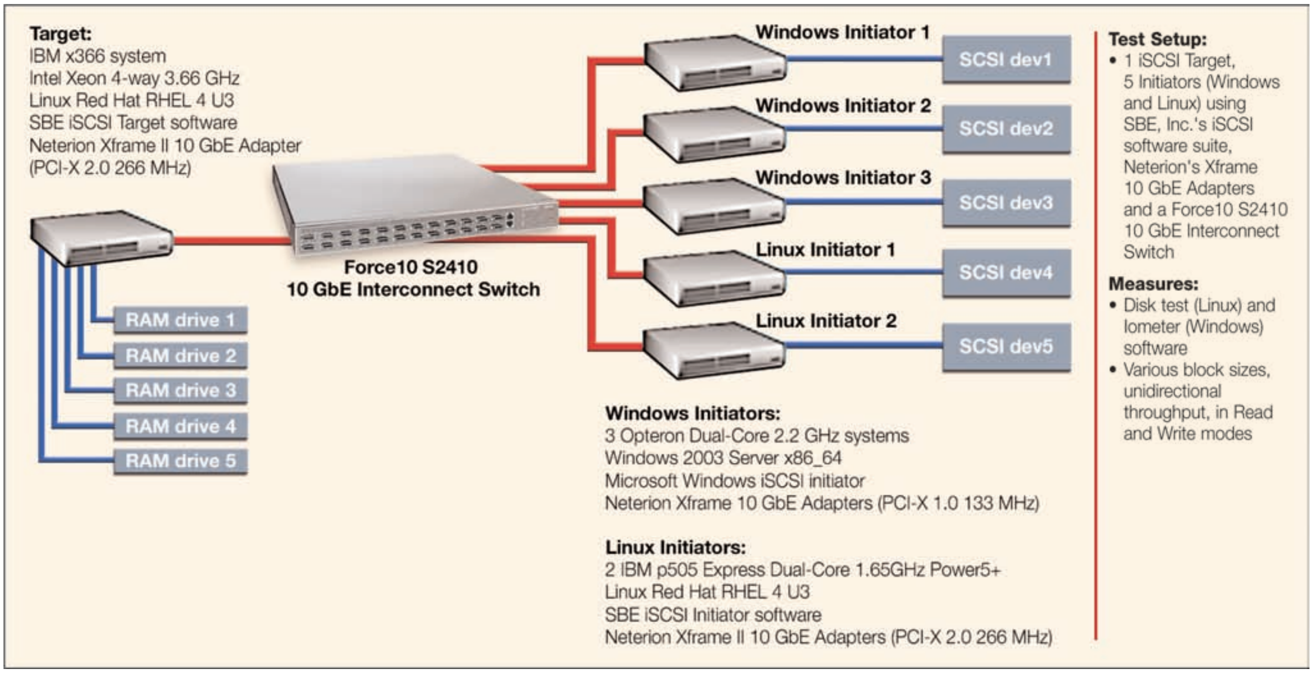

Figure 3: iSCSI Performance Benchmark Schematic

Details of the iSCSI systems configuration: Linux Target, 3 Windows Initiators and 2 Linux Initiators

Target

- IBM xSeries 366
- RHEL AS4 Update 3
- SBE iSCSI Target v2.5
- 4x Single-Core 3.66GHz Intel Xeon Processors with HT Disabled
- 8GB DDR400 Registered ECC Memory
- Xframe II PCI-X 2.0 266 10G (Driver rev. 5120)

Windows Initiator 1

- Windows 2003 Server x86_64
- Windows Initiator 2.02
- 2x Dual-Core 2.2GHz AMD Opteron 275
- 2GB DDR400 Registered ECC Memory
- 9x 300MB Ramdisk
- Iometer 2004.07.30
- Xframe I PCI-X 1.0 133 10G

Windows Initiator 2

- Windows 2003 Server x86_64
- Windows Initiator 2.02
- 2x Dual-Core 2.2GHz AMD Opteron 275
- 2GB DDR400 Registered ECC Memory
- 9x 300MB Ramdisk
- Iometer 2004.07.30
- Xframe I PCI-X 1.0 133 10G

Windows Initiator 3

- Windows 2003 Server x86_64
- Windows Initiator 2.02
- 2x Dual-Core 2.2GHz AMD Opteron 275
- 2GB DDR400 Registered ECC Memory
- 9x 300MB Ramdisk
- Iometer 2004.07.30
- Xframe I PCI-X 1.0 133 10G

Linux Initiator 1

- IBM p5 505 Express
- RHEL AS4 Update 3
- SBE iSCSI Initiator 1.6
- 2x Dual-Core 1.65GHz Power5
- 2GB DDR2 Registered ECC Memory
- Disktest v1.1.0
- 1x 300MB Ramdisk
- Xframe II PCI-X 2.0 266 10G (Driver rev. 5120)

Linux Initiator 2

- IBM p5 505 Express
- RHEL AS4 Update 3
- SBE iSCSI Initiator 1.6
- 2x Dual-Core 1.65GHz Power5
- 2GB DDR2 Registered ECC Memory
- Disktest v1.1.0
- 1x 300MB Ramdisk
- Xframe II PCI-X 2.0 266 10G (Driver rev. 5120)

Figure 4: Benchmark Test Configuration

Benchmark results, measured in Throughput, IOPS (IO per second) and CPU utilization on the iSCSI target and initiator sides, first with the Windows initiator systems, then with the Linux initiators

Windows Performance Results

	Target (x366)			Windows Initiator3			Windows Initiator2			Windows Initiator1		
	Throughput (MB/s)	IOPS	CPU	Throughput (MB/s)	IOPS	CPU	Throughput (MB/s)	IOPS	CPU	Throughput (MB/s)	IOPS	CPU
1KB Write	46.55	47668	78.9%	15.58	15953	12.6%	15.50	15877	11.3%	15.47	15839	11.2%
1KB Read	64.80	66355	65.5%	22.02	22548	17.6%	21.00	21514	16.7%	21.77	22293	16.7%
2KB Write	86.55	44313	74.6%	29.04	14869	11.4%	30.40	15565	11.2%	27.11	13880	9.7%
2KB Read	135.46	69357	65.6%	45.64	23366	18.8%	46.04	23573	18.9%	43.78	22418	17.8%
4KB Write	27.93	7149	10.6%	13.84	3542	2.5%	11.90	3046	2.5%	2.19	561	0.7%
4KB Read	147.75	37824	43.2%	48.62	12447	12.2%	52.25	13376	12.8%	46.88	12001	11.6%
16KB Write	434.89	27833	75.1%	146.18	9356	8.7%	146.24	9360	8.5%	142.47	9118	8.0%
16KB Read	429.01	27457	47.7%	140.55	8995	13.9%	150.97	9662	14.7%	137.50	8800	14.0%
32KB Write	555.70	17782	74.1%	182.98	5855	6.3%	188.33	6027	6.4%	184.39	5901	6.0%
32KB Read	869.72	27831	43.1%	291.03	9313	20.5%	291.72	9335	19.7%	286.97	9183	19.6%
64KB Write	576.86	9230	60.7%	193.34	3093	4.7%	191.30	3061	4.6%	192.22	3076	4.7%
64KB Read	1175.45	18807	34.9%	393.72	6300	23.2%	396.73	6348	24.2%	385.00	6160	22.3%
128KB Write	571.33	4571	59.1%	187.32	1499	4.0%	193.07	1545	3.8%	190.93	1528	4.1%
128KB Read	1186.30	9490	26.1%	397.10	3177	23.8%	401.37	3211	24.2%	387.83	3103	24.0%

Linux Performance Results

	Target (x366)			Linux Initiator2			Linux Initiator1		
	Throughput (MB/s)	IOPS	CPU	Throughput (MB/s)	IOPS	*CPU	Throughput (MB/s)	IOPS	*CPU
1KB Write	40.85	41980	74.6%	20.4	20928	32.1%	20.45	21052	28.1%
1KB Read	66.55	68246	79.5%	33.66	34514	14.9%	32.89	33732	15.0%
2KB Write	80.43	41211	72.9%	40.22	20625	25.7%	40.21	20586	26.5%
2KB Read	128.98	65994	80.8%	65.74	33439	15.7%	63.24	32555	14.1%
4KB Write	148.03	37930	73.2%	73.36	18772	27.3%	74.67	19158	26.8%
4KB Read	232.07	59260	80.7%	116.31	29771	14.8%	115.76	29489	13.6%
16KB Write	476.65	30486	90.1%	238.23	15232	23.4%	238.42	15254	24.2%
16KB Read	666.06	42612	87.8%	333.11	21302	16.8%	332.95	21310	16.7%
32KB Write	578.39	18545	61.7%	292.91	9373	21.3%	285.48	9172	21.4%
32KB Read	868.71	27823	77.6%	434.37	13905	12.3%	434.34	13918	12.4%
64KB Write	701.87	11230	56.9%	367.55	5881	18.2%	334.32	5349	18.3%
64KB Read	1143.1	18291	69.2%	573.95	9183	14.4%	569.15	9108	10.8%
128KB Write	739.64	5912	52.2%	369.89	2959	14.1%	369.75	2953	14.3%
128KB Read	1183.9	9471	47.8%	578.29	4626	11.9%	605.61	4845	18.1%

* CPU Utilization was a measure of the total iSCSI rx and tx threads only

Figure 5: Detailed Test Results

Detailed results of the benchmark can be found in the tables of Figure 5. Performance measures include Throughput (in MegaByte per second), IOPS (I/O per second) and percentage of CPU utilization. As an example, transferring sequential 128K sized blocks between the target and the Linux initiators yielded 1,184 MBps (9,472 Mbps) when reading, and 740 MBps (5,920 Mbps) when writing. Compared to the industry standard of 2 Gbps Fibre Channel, which can typically transport up to 200 MBps, 10 GbE iSCSI with ERL2 led to a performance increase by a factor of 6x.

10 Gigabit Ethernet is a key driver for iSCSI, delivering unparalleled performance for IP SANs, with the same quality of service and system robustness at a fraction of the total cost of acquisition and maintenance. Commodity off-the-shelf (COTS) hardware eliminates the need for expensive and proprietary equipment, and the prevalence of Ethernet through the network means that expert personnel, tools, and leading edge solutions will always be available. Finally, with iSCSI free of the distance limitations of Fibre Channel, the truly global SAN is now possible.

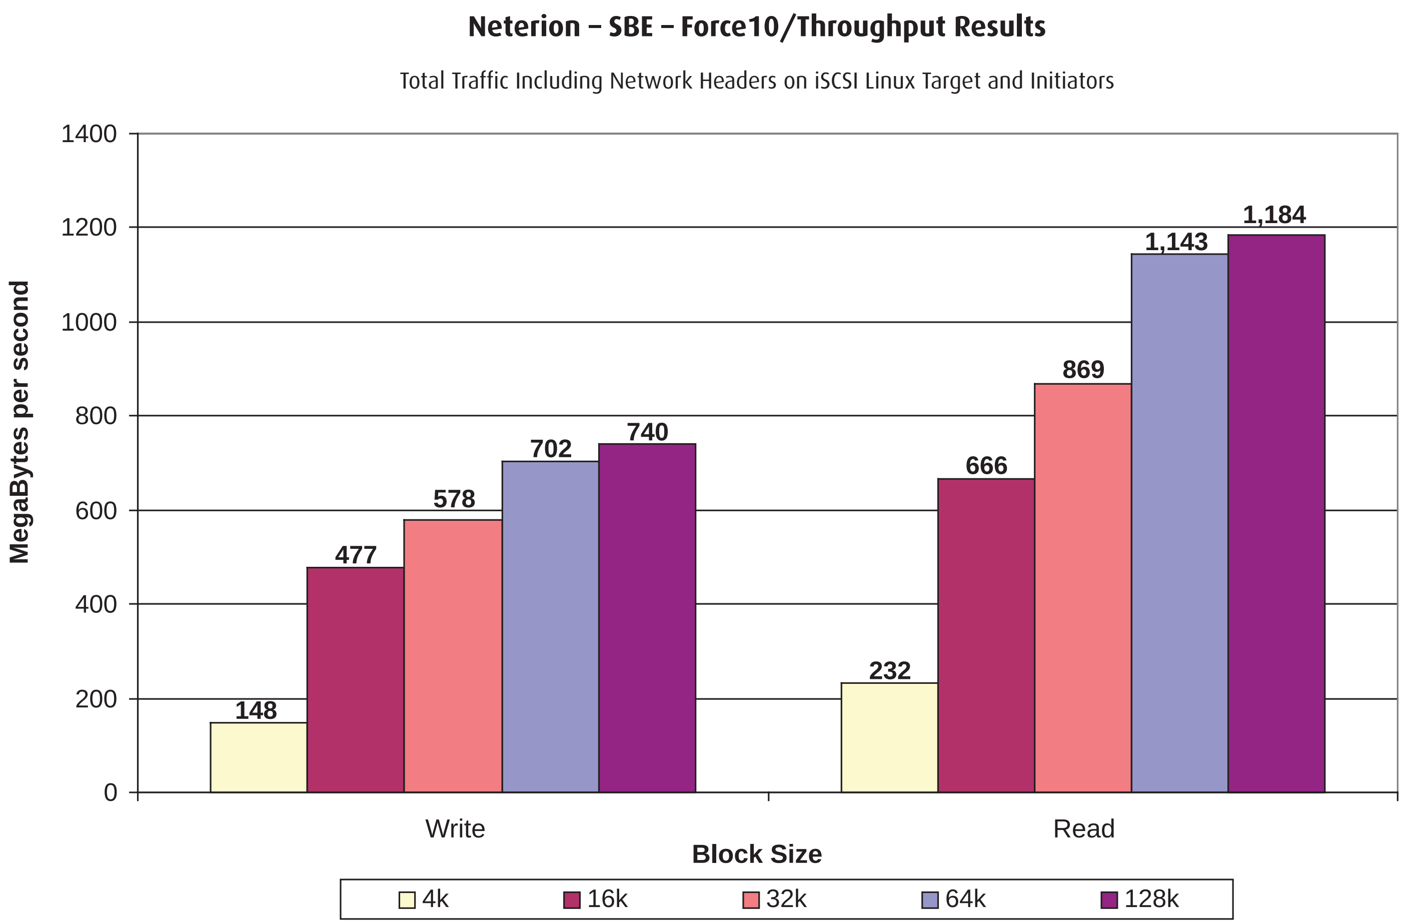

Figure 6: iSCSI Performance Benchmark Results

Stateless offload 10 GbE adapters make a great iSCSI solution

Best-in-class 10 Gigabit Ethernet adapters include comprehensive stateless (or "state-aware") offloads and assists, handling in hardware a significant amount of the TCP/IP processing road blocks, including Checksum Offloads, Large Send Offloads, Large Receive Offloads, and UDP "Checksum Over Fragment", to name a few.

As the benchmark results show (see Figure 7), these adapters are fully capable of carrying 10 GbE iSCSI traffic today with limited CPU utilization rates. Additionally, they offer the following benefits:

1) Full compliance with all operating systems without an overhaul of TCP/IP stacks
2) Fast adoption and deployment of 10 Gigabit Ethernet technology by major server and storage vendors
3) Preserving the security and reliability of established TCP/IP stacks
4) Taking advantage of Moore's Law and the dozens of billions of dollars of R&D spent to improve processor speeds, from Intel to AMD, Sun/Fujitsu to IBM. As a consequence, processor speeds will always eventually match networking speeds, making stateless adapters faster than full-offload approaches
5) There is no need to wait for emerging TOE and RDMA incremental benefits since 10 Gigabit iSCSI is deployable with the current stateless, OS-compliant adapters available today, like Neterion's Xframe 10 GbE adapters.

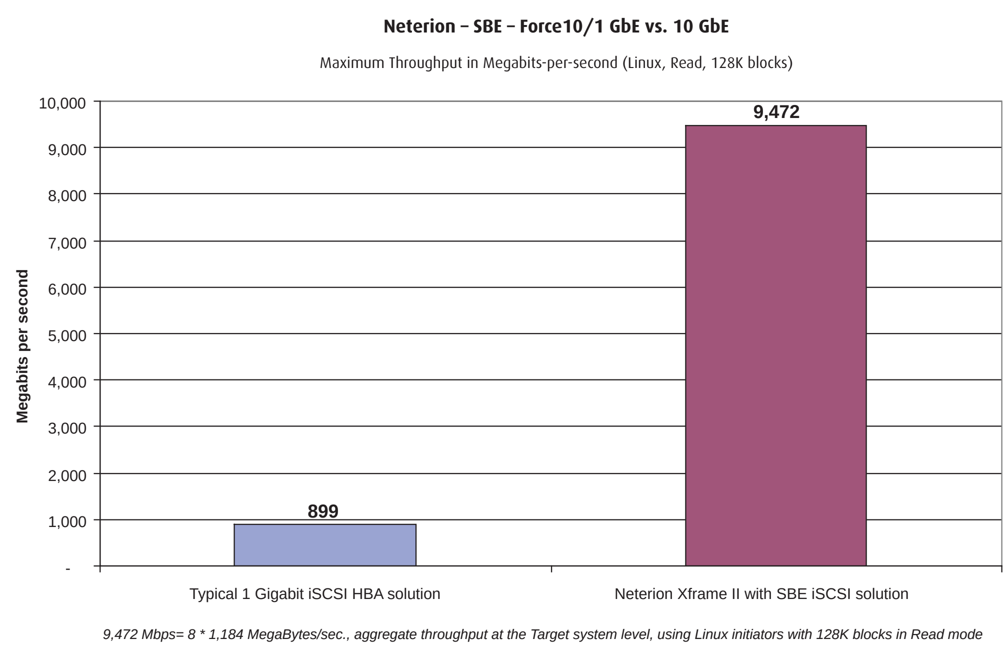

Figure 7: 10 GbE iSCSI Comparison

Source: Vendors benchmarks (posted on their websites) – more information available on demand

Contributing companies' information

Neterion

Founded in 2001, Neterion Inc. has locations in Cupertino, California and Ottawa, Canada. Neterion delivers 10 Gigabit Ethernet hardware & software solutions that solve customers' high-end networking problems. The Xframe® line of products is based on Neterion-developed technologies that deliver new levels of performance, availability and reliability in the datacenter. Xframe, Xframe II and Xframe E include full IPv4 and IPv6 support, and comprehensive stateless offloads that preserve the integrity of current TCP/IP implementations without "breaking the stack." Xframe drivers are available for all major Operating Systems, including Microsoft Windows, Linux, HP-UX, IBM's AIX, Sun's Solaris and SGI's Irix. Further information on the company can be found at http://www.neterion.com.

Force10 Networks

Force10 Networks is the pioneer in building and securing high performance switching and routing. Based on a revolutionary system architecture that delivers best-in-class resiliency and massive scalability, Force10's TeraScale E-Series switch/routers ensure predictable application performance, increase network availability, and reduce operating costs. Today, many of the world's largest Gigabit Ethernet and 10 Gigabit Ethernet networks depend on Force10 Networks. For additional information, please visit http://www.force10networks.com.

SBE

SBE designs and provides IP-based storage networking solutions for an extensive range of business critical applications, including Disk-to-Disk, Back-up, and Disaster Recovery. SBE delivers an affordable, expandable, and easy-to-use portfolio of software solutions designed to enable optimal performance and rapid deployment across a wide range of next generation storage systems. Based in San Ramon, California, SBE is a publicly traded company (NASDAQ: SBEI) with products sold worldwide through direct sales, OEMs and system integration partners. For more information, please visit the SBE website at http://www.sbei.com.

Neterion™ Practical Innovation

www.neterion.com

Neterion, Inc.
20230 Stevens Creek Blvd.
Suite C
Cupertino, CA 95014
Main Phone: 408.366.4600
Fax: 408.366.4650

Neterion, Corp.
340 Terry Fox Drive
Kanata, Ontario
Canada, K2K 3A2
Main Phone: 613.592.0457
Fax: 613.592.1802

Information: info@neterion.com
Sales contact: sales@neterion.com