

# 10 Gigabit Ethernet: The Tipping Point for IP SANs

A Chelsio Communications White Paper



## Introduction

Internet SCSI (iSCSI) is a proven storage area network (SAN) technology - a number of leading storage vendors have been marketing and shipping iSCSI products for years. Until recently, iSCSI SAN performance has been limited by Gigabit Ethernet, compared to the 2 or 4 Gbps available with Fibre Channel. Now, 10 Gigabit Ethernet (10GbE) is available and practical, and with protocol offload enables iSCSI technology to surpass Fibre Channel in performance.

The ease and flexibility of IP make the case for 10 GbE that much more compelling. Since it leverages the same administration as Gigabit Ethernet and 10/100 Mbps Ethernet, the administration tools used today for managing the local area network (LAN) will work also for a SAN based on 10 Gigabit Ethernet. What's more, Ethernet provides compatibility across link speeds and vendors, which has always been a challenge with Fibre Channel. Because it is based on IP, iSCSI interconnects with every vendor's product. Users can freely mix and match speeds, such as connecting servers with Gigabit Ethernet links to storage systems which have 10 GbE connections.

## Enabling iSCSI in the Enterprise

Enabling iSCSI in the enterprise is greatly facilitated by the availability of iSCSI software and its reliance on the familiar TCP/IP/Ethernet infrastructure. Furthermore, it leverages much of the existing technologies used for Internet connectivity, such as authentication and security, at no additional cost. Taken together, these factors result in considerable cost savings.

All the major server operating systems support iSCSI. Microsoft's iSCSI initiator is freely-downloadable for Windows, and iSCSI is a component of Linux, HP-UX, Netware, Solaris, AIX, and MAC OS. Using built-in iSCSI initiator stacks reduces network attachment costs for servers, which can then connect to IP SANs without requiring specialized HBAs, saving \$600-to-\$1000 per server.

IP SANs can be fully redundant without expensive multi-pathing software on hosts. Microsoft incorporated its multi-path I/O (MPIO) technology in its iSCSI initiator, and iSCSI products are available that support the Windows MPIO implementation. This eliminates licensing fees for HBA-failover and HBA-load-balancing software that typically costs \$2-4K per SAN-attached server for a Fibre Channel SAN.

Security and network management are key concerns for IT administrators. In terms of security, IP offers richer features than Fibre Channel, including CHAP for secure user authentication, and IPSec for data encryption between the host and the iSCSI storage itself. In addition, administrators can continue to use their favorite TCP/IP network management tools for iSCSI SANs.

In terms of price-per-Gigabit, iSCSI costs less than half of 4G Fibre Channel: port prices for a 10GbE CX-4 switch or a 4G Fibre Channel are comparable, at about \$500 to \$600. However, 10GbE delivers more than twice the bandwidth for the same cost. Today, 10 GbE is already less expensive per Gigabit than Fibre Channel, and with the economies of scale Ethernet is famous for, the costs will continue dropping.

## iSCSI Performance

iSCSI has reached the point where it provides Fibre Channel-class performance, both in terms of speed and scale. Two years ago, traditional IP SAN implementations were typically 1 to 3 Terabyte deployments, mostly small and mid-tier environments. Today, it is common to see plans for IP SANs that are 200 Terabytes and beyond. The number of existing deployments confirm that iSCSI is real, proven, and now that it's powered by 10GbE, it is ready for mission-critical deployments in the enterprise.

The conclusion is clear: with the iSCSI price and performance advantages brought in by the arrival of 10 GbE with protocol offload, the discussion about iSCSI versus Fibre Channel is virtually over. With 10GbE, deploying an IP SAN no longer means compromising on performance in favor of cost savings. Quite the contrary, the savings are now accompanied by a doubling of the performance. Add to that the compatibility and simplicity of IP compared to Fibre Channel, the case for deployment of iSCSI SANs as a simple extension to enterprise LANs is very clear.

## 10GbE iSCSI Applications

**Storage Consolidation** 10GbE facilitates storage consolidation for mission-critical enterprise servers and Microsoft clusters into one IP SAN. In this configuration, up to ten servers can get iSCSI at GbE wire speed from a 10GbE storage array in tandem with your network attached storage. This enables getting the most out of email, Web, MRP and other business application servers. Figure 1 illustrates storage consolidation of enterprise applications such as Oracle, Exchange, and SQLServer.

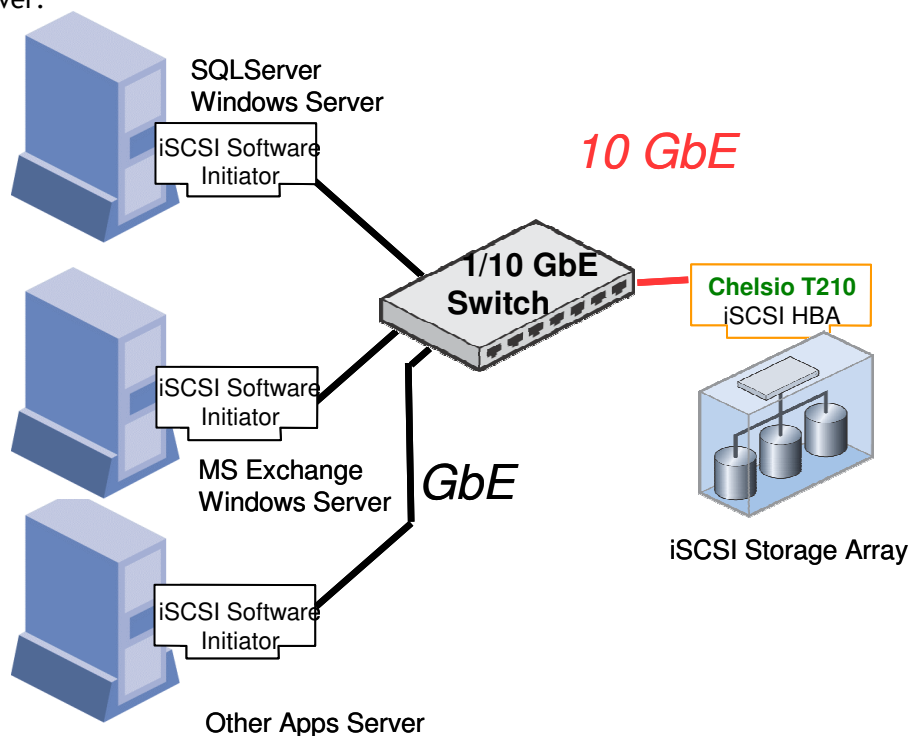


Figure 1: Storage Consolidation using 10GbE iSCSI

**Rich Media** In many networks built for processing high-definition media data, such tremendous storage performance is required that a storage system can only support a few workstations. This leads to the proliferation of “storage islands.” With the 10 GbE bandwidth advantage, storage systems can support many more users, and sites therefore require managing fewer boxes.

**HPCC** In high-performance cluster computing (HPCC) environments, SAN file systems, such as the Red Hat Global File System (GFS), have emerged to provide dramatically higher performance than NAS appliances. Using a SAN file system with 10 GbE iSCSI provides tremendous throughput for file sharing applications. Figure 2 shows the use of 10 GbE iSCSI for high-performance applications, such as rich media production and high-performance cluster computing (HPCC).




Figure 2: Storage Area Networking for Rich media Production and High-Performance Cluster Computing (HPCC)

**Disaster Recovery and Mirroring** 10 GbE facilitates data transport over long distances, providing a vehicle for disaster recovery, replication, and mirroring to reduce down time. If there is a disaster at one site, the data at the failover site is more current and more quickly available when 10 GbE is used as the mirroring interconnect.

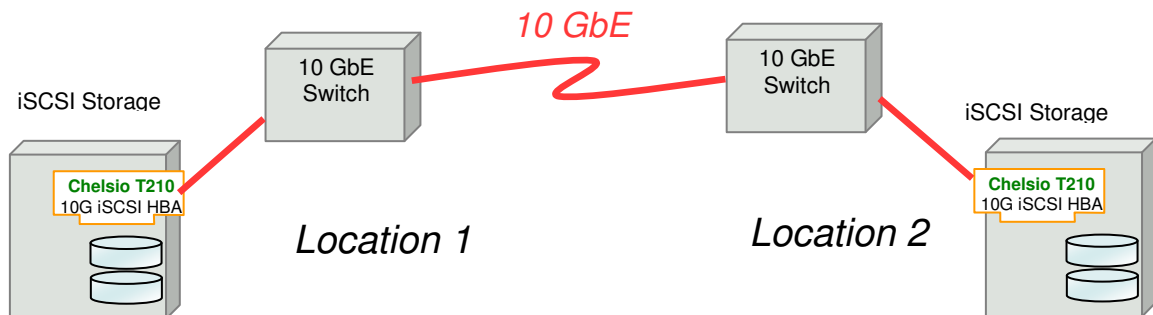


Figure 2: High-Availability and Business Continuity Using 10 GbE iSCSI

## SAN Fabric Cost Comparison

Ultimately, with multiple contending technologies offering satisfactory performance, the decision of which SAN fabric to deploy will invariably come down to cost, and here is where the true value of a unified 10 GbE-based fabric emerges.

Consider the following topology which may be typical in data center installations within the next few years:

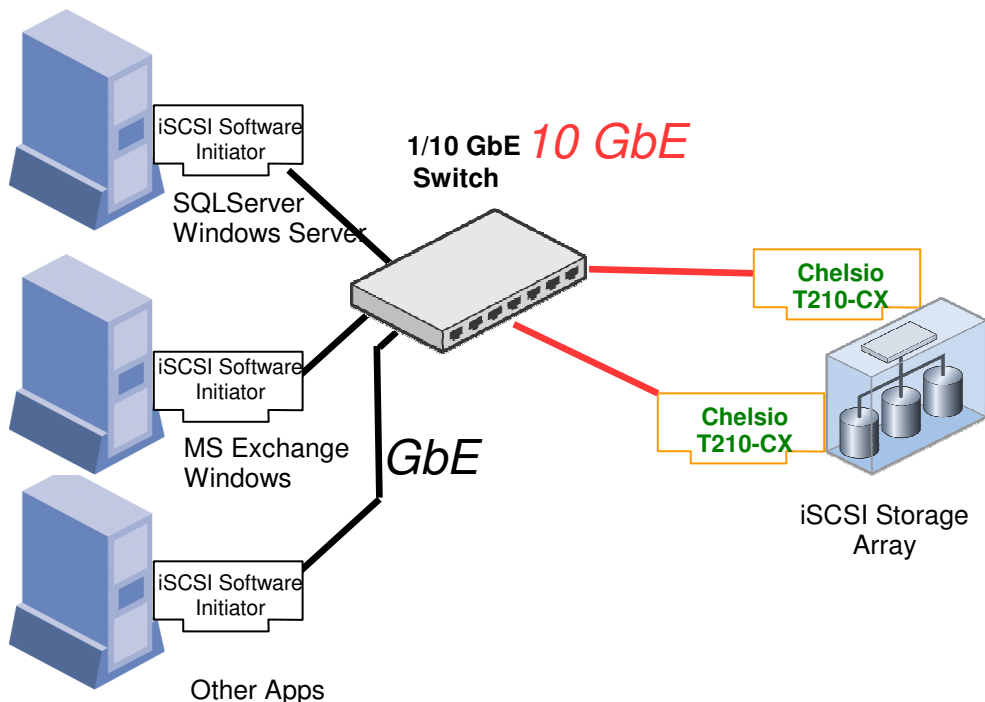


Figure 3: Typical Data Center Installation

The assumptions behind this topology are:

1. Servers typically have multiple motherboard-based Gigabit Ethernet ports as a default for enabling data and storage networking. Thus, GbE-based server connections to an iSCSI SAN can be considered to be “free”.
2. GbE iSCSI initiator connections provide enough bandwidth for typical applications (e.g. SQL, Exchange, etc) to run on each server.
3. iSCSI software initiators are built-in to all volume operating systems including Linux, Windows, Mac OS, Solaris, etc., so IP SAN-attachment for servers will also be “free”.
4. Server to storage SAN communications, including iSCSI deployments, will typically use a physically (or virtually using VLANs) segregated network for reasons of security and minimizing traffic contention, although with IPSec, SAN and LAN traffic could be safely mixed on the same fabric.

5. Two 10GbE connections will be made to the storage arrays, in order to provide enough storage bandwidth to support applications running simultaneously on all 48 servers, with active-active failover, if one connection is lost

The following table can now be constructed which compares the costs of using iSCSI, Fibre Channel and InfiniBand fabrics to implement such a SAN topology in late 2006:

	10GbE	FC	IB	Assumptions
Cost of SAN Attachment for 48 servers	Free	\$17,328	\$26,400	1. 2G FC HBAs @ \$361 and IB HCAs @ \$550. 2. Cable costs for 1GbE, FC, and IB are negligible, \$45, and \$100 respectively.
Storage Array Attachment delivering aggregate 20Gbps I/O capacity	\$3,400	\$2,225	\$1,300	1. Each 10GbE target port costs \$1,600. Each 4Gbps FC target port costs \$500 per port and each 10Gbps IB target port costs \$550. 2. Cable costs for 10GbE, FC, and IB are \$100, \$45, and \$100 respectively.
Switching Capacity for 48 SAN-attached servers and 20 Gbps capacity for storage target attachment	\$5,000	\$21,164	\$12,500	Switching capacity cost for GbE assumes 48 GbE ports and two 10GbE ports using HP Procurve 3400cl as an example at projected 2006 prices  FC switch pricing is \$368 per 2G port and \$700 per 4G port, and IB switch pricing is \$250 per port.
Storage Area Network Acquisition Cost	<b>\$8,400</b>	<b>\$40,717</b>	<b>\$40,200</b>	48 server initiators with 20Gbps of storage array bandwidth

Note: FC and IB cost data are based on Gartner (2004) and ESG (2005) projections for 2006 respectively.

Thus, even without taking into account the considerable costs of specialized IT management, hardware, and driver software, a 10GbE iSCSI based SAN target solution costs about 20% of a Fibre Channel or InfiniBand SAN for similar or better storage performance. This mostly comes about because the scalability on the server (initiator) side is essentially free.

Thus, iSCSI has the advantage of enabling migration to a unified data/storage networking fabric, driven by the growth in 1GbE and 10GbE networking, providing a compelling value proposition in terms of cost, ease-of-use and performance.

## Chelsio iSCSI Solution Overview

While Chelsio is committed to a full range of product offerings to serve the emerging iSCSI market, we have initially started with the iSCSI target space since software initiators dominate the initiator side. Since Gigabit Ethernet connectivity comes free with the server, as does the software initiator, the near-term growth potential - and ROI - is huge because sites can scale the population of SAN-attached servers for free.

## **Chelsio's Current Products**

- T210 is the only 10GbE TOE/iSCSI offload adapter available in the market today. It has available models for long-range and short-range fiber as well as a CX4 copper version.
- T204 uses the same silicon engine to deliver 4Gbps in a quad-port GbE TOE/iSCSI offload adapter; it therefore integrates easily into the existing GbE infrastructure to provide a sizeable boost to server connectivity.
- N210 is a 10GbE NIC, a stateless offload product for low-end applications.

To learn about the requirements for iSCSI target systems, download our white paper, *Building High-Performance IP SANs with Chelsio's 10GbE iSCSI Solution* from [www.chelsio.com/solutions/whitepapers.php](http://www.chelsio.com/solutions/whitepapers.php).

## **Chelsio's Value for Storage**

Chelsio delivers the best-in-class, highest performance iSCSI target solution in the industry today, including:

- Full TCP and iSCSI CRC offload which reduces CPU utilization and maximizes server performance.
- A unique architecture which provides consistent performance scaling from one to tens of thousands of connections, with a performance profile allowing equal and stable bandwidth per connection.
- Programmable architecture that allows for field upgradeability.
- An enterprise ready product.

It is known that products in the iSCSI and TOE space, especially 10 GbE, are easy to demonstrate but difficult to ship. Customers are often dismayed by the gap between the performance obtained in real deployment scenarios and the claimed benchmark numbers, which are collected in staged demos. Chelsio's protocol offload engine has successfully cleared this hurdle, with independently proven performance, and is currently shipping its 2nd generation 10GbE silicon. Furthermore, Chelsio has a commitment to standard manufacturing processes and was recently certified for ISO 9001, a critical milestone for an enterprise supplier.

The overarching goal behind our rich roadmap is to deliver an additional offloaded protocol at half the price every twelve months. Our intention is to drive the premium for TOE technology and iSCSI technology to zero within the next twelve months while providing Ethernet connectivity at all speeds, with all protocols using all host busses. This will provide the lowest total cost of ownership for Ethernet connectivity for the iSCSI space.

## **Conclusion**

10GbE iSCSI is ready for mainstream deployment and offers the lowest total cost of ownership among the available SAN alternatives. TCP/IP and iSCSI protocol offload is a critical feature of 10GbE iSCSI infrastructure, delivering optimal use of both the 10 Gigabit network capacity and host CPU cycles.