

iSCSI SANs — Panacea or Placebo?

Analyst: David Reine

Management Summary

Today, access to the Internet has become the norm rather than the exception. Professionals in every industry use it to work remotely to enable a better life style. It is used to check our bank statements and to monitor our credit card activity. We try to guard against identity theft. Our children use the Internet continually. They use it to research their term papers, to “talk” to their friends, and to play games. **How businesses access the Internet, however, is another story.**

In the beginning, shortly before Al Gore invented the Internet, we were able to connect to a data center from a remote location via the telephone. Initially using an external acoustic coupler, and then evolving to higher speed, internal modems, we communicated over the same wire that our voices traversed. Unfortunately, when the computer was in use, the telephone was not available for voice use. It was either voice or data. This method of communication was also slow, painfully slow. From 110 bits per second (bps), we advanced to 19.2k bps, and later to 56k bps. Then we were provided with high-speed DSL modems to give us the flexibility to transmit data and voice simultaneously. It enabled us to virtualize the communications experience. We could talk on the phone and transmit email, and not necessarily to the person to whom we were talking. Furthermore, we gained access to high-speed transmissions, up to 80 times faster than over the phone line.

These same conditions existed in the enterprise data center. **The Information Technology (IT) staff was trying to manage the communication between servers over a local area network (LAN) at the same time that those servers were also sending information to, and receiving information from, a collection of storage devices.** For many larger organizations, it became impractical to attempt to transmit both server-to-server data and storage data through the same pipe. It was not fast enough or wide enough for efficient communications. Thus, the storage area network (SAN) was developed. SANs were installed in the data center on Fibre Channel (FC) networks to improve the speed of communications between server and storage and to enable multiple servers to share the same disk arrays and tape systems to improve the utilization of expensive enterprise resources. Unfortunately, SANs required the installation of an expensive infrastructure: a FC adapter in the server, two if high availability was desired, a FC switch, or two, to interconnect all of the devices, and FC-enabled arrays. This does not even address the cost of the human resource, the SAN administrator, high salaried and scarce. **For many smaller enterprises or branch offices, a FC SAN was too expensive and too complicated to implement.** Now, along comes the iSCSI SAN, which presents an interesting, FC-like alternative, without the cost and complication of Fibre Channel at each step along the path.

Using the TCP/IP protocol over faster IP lines, iSCSI can use existing networks and resources - but at what price and at what performance? To see if iSCSI can provide your data center with an effective storage solution without the cost and complexity of a FC SAN, please read on.

IN THIS ISSUE

➤ Challenges in the Enterprise Storage Environment.....	2
➤ The Current State of iSCSI.....	3
➤ The iSCSI vs. FC Dilemma	4
➤ Conclusion	6

Challenges in the Enterprise Storage Environment

With the remarkable growth of storage in data centers and remote offices, efficient block storage management and economic access to enterprise information has become both more important and more difficult to achieve. Enterprise storage is often doubling each year as businesses generate increasing volumes of business critical information:

- To support the rapid growth in existing applications, such as on-line transaction processing and corporate email;
- To comply with industry and government regulations concerning the preservation of enterprise financial data;
- To achieve a higher level of competitiveness through data warehousing;
- To implement processes for disaster recovery through data replication; and
- To rollout new e-commerce applications.

The overall requirement to store, backup, archive, and protect enterprise data has resulted in an era of unprecedented growth in networked storage while data center staffs strive to remain within the IT budget. The systems architecture of yesterday with large multi-processor servers and direct-attached storage (DAS) evolved into an environment focused upon distributed processing with multiple smaller servers, each running a different application and its own storage. Inefficiencies in the utilization of these servers led to server consolidation within the data center via larger scale-out configurations, implemented to make better use of enterprise resources and to lower the total cost of ownership (TCO). Concern over throughput on the LAN and inefficiencies in the provisioning and utilization of DAS storage led to the development of storage area networks (SANs) in the '90s. **SANs allow shared access to storage resources to map, and unmap virtual disks dynamically. This enables more efficient utilization of the resource and easier management through the development of storage networking technology that enables the sharing of network-attached storage by any server connected to the same network.** Specific objectives that drive the enterprise CIO to rollout a SAN configuration are listed in Exhibit 1, above.

In larger enterprises, the acknowledged

Exhibit 1 – SAN Objectives

- Implement Storage Consolidation for Heterogeneous Application Servers;
- Improve Storage Scalability;
- Enable a Flexible Utilization of Storage Resources to Improve Data and Application Availability and Reliability;
- Centralize and Improve the Performance of Backup and Recovery;
- Create a Storage Pool to be Accessed via SAN (FC or iSCSI) or NAS (via Virtualization);
- Reduce the Cost and Complexity of Storage Management;
- Improve Employee Productivity; and
- Increase Customer Satisfaction and Enterprise Reputation.

standard for storage networking is via FC SANs, a specialized technology designed for the high-speed exchange of data between high-performance disk arrays and open system servers, with those storage devices connected via a FC network.

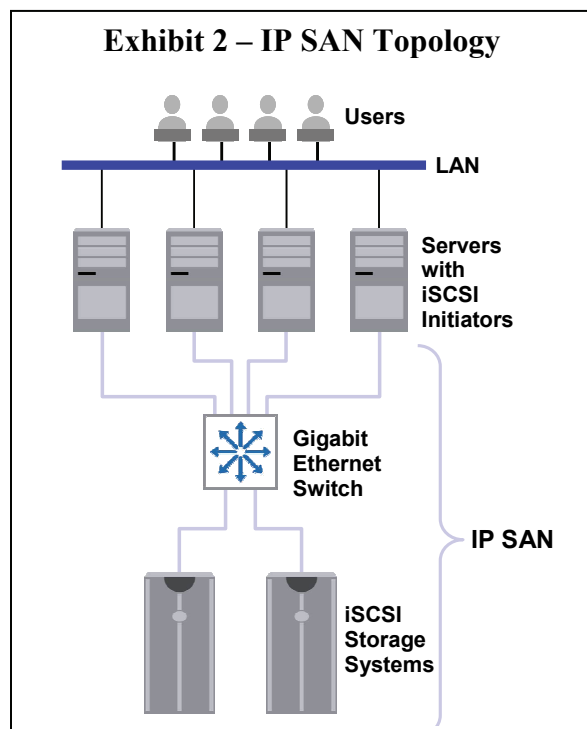
Unfortunately, to achieve the speed and functionality provided by Fibre Channel, the enterprise often had to create a complex, hard to maintain storage infrastructure, resulting in relatively high costs through the acquisition of FC adapters within the servers, FC storage devices, specialized FC switches on the network, and, experienced technicians and administrators to manage the whole network and keep it viable¹. All of those costs are justifiable for a primary, high-speed network within an information lifecycle management (ILM) environment where rapid access to the most mission-critical data is mandatory². In non-mission-critical environments, where speed is not an issue and cost is, alternatives can be examined to complement FC as a secondary network, or replace it in less demanding environments as a primary network, remote offices for example.

¹ Recently, FC vendors have started to reduce the cost of this infrastructure.

² See the issue of **Clipper Notes** dated February 28, 2007, entitled *Tiered Storage Classes Save Money — Getting the Most Out Of Your Storage Infrastructure* and available at <http://www.clipper.com/research/TCG2007032.pdf>.

In an *e-era* where a lost file can cost your Internet business thousands of dollars for every minute waiting for retrieval, doing disk-to-disk backups over a FC SAN may have been the only logical alternative to tape to enable rapid recovery. However, some data does not require instant access, and the higher associated costs. That secondary tier of data can be backed up on a slower, less expensive network, such as iSCSI, and interface with a variety of devices, not only Fibre Channel, but also ATA and SATA. Combining slower, less-costly devices such as SATA with iSCSI enhances the savings for the data center, yet keeps this data on-line.

Advancements within Ethernet technology, enabling high-speed, secure Internet communications, raises the possibility of the use of TCP/IP for SAN communication using the same protocol that supports communication between the servers. (See Exhibit 2, below.) Thus, an IP SAN could significantly lower the TCO for the storage network through the sharing of the IP infrastructure without affecting business operations or losing the benefits of the SAN, while taking advantage of the availability of a larger pool of highly-skilled technicians and administrators. The total cost of networking, whether new or consumed from existing resources, needs to be included in any TCO comparison. This presents an interesting challenge for the IT staff in the larger enterprises trying to reduce their TCO.



- How do you take advantage of the advancements in iSCSI technology while preserving application performance and protecting the information assets of the enterprise?
- In smaller operations or departments, can the staff take advantage of iSCSI to consolidate the small or mid-sized office? Do they replace or complement the existing architecture?

As enterprise departments and smaller businesses, with a significant infrastructure of x86 servers running Windows or Linux, continue to deploy low cost IT solutions for e-mail, database and web applications, the costs associated with FC SANs will undermine their ROI. Let's look at iSCSI and see what it is and how it compares to FC.

Current State of iSCSI

Currently, iSCSI is enjoying an initial surge in popularity as major vendors and lesser-known rivals present their iSCSI solutions to an IT community looking to take advantage of the business benefits that iSCSI can provide. Two years ago, Microsoft established credibility for iSCSI by endorsing it with the inclusion of an iSCSI initiator in Windows. This validation enabled hardware vendors, large and small, to join into what has become a mature ecosystem at thousands of sites, with packaged solutions that bundle servers with iSCSI initiators to facilitate the deployment of iSCSI using other operating systems. This has enabled low cost, entry-level storage arrays that can connect to iSCSI and FC SANs, as well as NAS - arrays that come with a mix of SATA, ATA, SAS, and FC devices. However, iSCSI is more than just a poor man's SAN. **Flexibility and ease of use are the keywords, as Tier 1 storage suppliers enable the overnight deployment and installation of iSCSI arrays rather than the days, or weeks, necessary to deploy a FC solution.** In response, many enterprises are asking for a definition of the business benefits for their organization before committing to any expenditure, even the relatively light one for iSCSI. Here are some of iSCSI's benefits.

- **Enables an economical secondary tier for the deployment of an ILM strategy;**
- **Facilitates the remote replication of storage for branch offices;**
- **Simplifies the network through ease of installation and deployment;**

- Eliminates the requirement for specialized administrators and technicians;
- Reduces total cost of ownership by taking advantage of existing infrastructure and low-cost drives such as SATA; and
- Improves enterprise profitability without affecting the enterprise business model.


The greatest potential for iSCSI could be realized in a remote office or branch of an enterprise, or a smaller business, where the greatest growth in storage is occurring and economy is the watchword. This is especially true for government agencies and schools where TCP/IP networks already exist. In this environment, iSCSI can be used as the primary network with applications being deployed to run email solutions, databases, and web services. iSCSI also enables an inexpensive remote replication of data for disaster recovery (see Exhibit 3, above) and facilitates the replacement of tape in the smaller office.

iSCSI is an ideal enabler of server storage consolidation. A secondary tier of storage, or secondary network, can save the enterprise a significant amount in money and resources. The application set in the data center usually contains a mix of applications, some not as mission-critical as others. In fact, high-speed FC communication may be overkill for on-line transaction processing, while it may be ideal for enterprises that transmit large multimedia files. Establishing iSCSI as a tool in the storage consolidation toolbox as a complement to FC can contribute to the bottom line in an ILM architecture, using iSCSI in remote locations and as an iSCSI gateway or a secondary tier in the data center. iSCSI has also proven to be an effective tool as a gateway in the rollout of NAS servers.

Recently, a broad range of iSCSI storage arrays has been introduced from some of the leading players in the IT arena, legitimizing iSCSI and pulling it into the mainstream for storage networking.

The iSCSI vs. FC Dilemma

iSCSI is a rapidly growing, industry-standard communication protocol that enables SCSI block I/O commands, sequences and attributes to




be sent over an IP network via unsolicited packets to provide transparent support for storage from cross-platform (Linux, Windows, UNIX, etc.) applications. (See Exhibit 4, on the next page.) This is similar to the manner in which SCSI commands are transmitted over Fibre Channel and *UltraSCSI*³ media. iSCSI uses TCP flow control and congestion control as it builds upon IP addressing and discovery procedures. With 128-bit addressing, iSCSI is highly scalable, running over a one gigabit per second (Gbps) network, with ten Gbps speeds planned. The discovery process used by iSCSI enables a node to hard code an address, query a storage name server, or send a multicast message inquiring which devices it can access. For data transmission reliability, iSCSI uses a 16-bit CRC⁴ for checksum. With commodity network interface control (NIC) adapters, it may be necessary to do the CRC calculation in memory using up important execution cycles, so specialized adapters may be preferred. iSCSI can be implemented either as a combination of a commodity gigabit Ethernet NIC with the TCP/IP and SCSI layers in software or as a specialized TOE adapter, built with Ethernet, IP, TCP and SCSI layers⁵.

³ UltraSCSI will not be addressed in this paper because its scalability limitations, in terms of distance and connectivity, inhibit its acceptance as a networking technology.

⁴ Cyclic Redundancy Code is used within error detection mechanisms.

⁵ TOE adapters maximize system performance for general networking by providing complete protocol offload for TCP/IP traffic.

Exhibit 4 – Heterogeneous iSCSI Environment



TCP/IP has been used in the data center for decades to communicate between servers over a LAN. Although not a serious problem for servers, the TCP flow control mechanism allows packets to be dropped during periods of congestion, which can cause a performance problem for LAN based storage communications. iSCSI tries to leverage the higher speeds of the existing TCP/IP over gigabit Ethernet infrastructure to minimize this problem. With gigabit ports integrated into commodity servers, and the cost of external adapters continuing to fall, iSCSI provides significant cost savings over FC.

Unlike iSCSI, which is software-based, Fibre Channel is a hardware-based protocol and is designed for higher performance by including the physical, link, network, and transport layers of the OSI network stack. The FC protocol contains a definition for SCSI over Fibre Channel called *FCP* and transmits via solicited packets. Like iSCSI, FCP was designed as a secondary network; it does not compete with internal protocols such as PCI. FCP currently supports network speeds up to 4 Gbps, four times that of a gigabit Ethernet network, with 8 Gbps performance planned for the 2006-2007 timeframe, and 16 Gbps planned. However, ten gigabit Ethernet has arrived. In fact, a new 10-Gbps Ethernet chip for servers and switches has already been announced to provide full duplex operation over a single unshielded twisted pair cable to improve throughput for storage require-

ments in excess of 50TB. With 64-bit addressing, FC is significantly more scalable than UltraSCSI, which is limited in the number of target devices, 15, and the total length of any SCSI bus, but not as scalable as iSCSI. With reduced scalability, there is less of a chance for congestion and the flow control mechanism for FC does not allow for packet loss. The CRC process for FC uses a 32-bit calculation, more robust than the 16-bit calculation of iSCSI, and is always done on the NIC as the data is moved through, preserving the processing cycles for the mission-critical applications. In FC, the fabric manager, or switch, controls the discovery process on the SAN, although a FC storage device can be connected directly to the server as a DAS. When a new node is activated, it notifies the switch to inform all the devices that have registered earlier with a request to be notified. Fibre Channel is noted for high availability, high performance, and the

reliability needed for connection from a data center full of servers to the shared storage array. Microsoft distributed revision 2 of their iSCSI initiator for Windows in September 2006, including multi-path I/O (MPIO) support for iSCSI, providing a transparent failover capability for high availability, along with load balancing, to improve performance and robustness from Windows.⁶

iSCSI enables the IT staff with the capability to replace, or complement, their FC SANs with an iSCSI SAN in environments where the requirements are less and the inefficiencies tolerated. In effect, with sufficient throughput, this could mean the consolidation of the server and storage networks onto a single LAN, uniting networks that were architected with different design considerations and different performance requirements. In addition, the IT staff could implement an iSCSI SAN to complement the FC SAN within an ILM framework, to create a multi-tiered storage environment that would improve the throughput of the FC SAN while reducing the TCO for storage of secondary data.

In addition to the savings from the use of an IP infrastructure, the data center can also exper-

⁶ It should be noted that the leading suppliers of iSCSI devices are the same vendors who provide the IT community with FC and, therefore, control the advances in iSCSI functionality.

ience significant savings with the implementation of SAS⁷ and SATA drives, in conjunction with iSCSI, to lower further the cost of the SAN. Although not as performant as Fibre Channel drives, they are significantly less expensive and have become widely accepted. Due to the limited on-board capabilities of iSCSI, many of the features executed within the layers of a standard FC adapter stack have to be executed in memory, taking extra execution cycles from the mission-critical applications⁸, or requiring the acquisition of a specialized iSCSI server adapter as part of a total iSCSI/SATA solution. This certainly could lower the TCO for an online archiving environment, for example, where SATA response would be significantly faster than tape and not as expensive as FC. Another solution that is becoming popular for the iSCSI/SATA combination is D2D (disk-to-disk) backup or D2D2T where the SATA drives serve as an intermediate stop, on the way to tape, for short-term recovery requirements. iSCSI also provides an ideal environment to access information directly from the data center, possibly eliminating the requirement for storage in the branch. In lieu of that, iSCSI can be used for backup requirements in a branch office where TCP/IP is used over the Internet for recovery from the data center.

Conclusion

Much like the pioneers who challenged the existing trends and placed Linux in their mission-critical operating system environment, early adopters of iSCSI have put their collective toes into the water to test this newest “commodity” architecture in order to reduce the cost and complexity for a variety of infrastructure applications, including databases, email, file storage, video processing, and backup/recovery. Like Linux today, **you can expect iSCSI to become ubiquitous throughout the enterprise community due to economies available by using their existing network infrastructure and the improving functionality available.** The largest enterprises may not use iSCSI exclusively, but as its speed improves to ten Gbps, you will find it in every data center

whether in SANs or a NAS environment. Many smaller enterprises, without high performance requirements, will take advantage of the TCO gains and implement a single networking infrastructure, iSCSI. Others will wait for the higher speeds before proceeding. It should be noted, though, that the network is rarely the cause of system bottlenecks. **Every smaller business or department in the larger enterprise needs to consider iSCSI as a primary network to access enterprise storage at the right price for their email, database, and web services, as well as other mission-critical business applications.**

Fibre Channel will live on today, and tomorrow, in the data center for high performance, mission-critical applications, especially where significant investments in FC architecture have already been made. There are hundreds of thousands of SANs installed in the enterprise today. The vast majority of these are Fibre Channel, even though they all do not have to be. iSCSI has been accepted. It has made a significant penetration in thousands of high-end data centers that are deeply committed to ILM and have added iSCSI to their portfolio of tools being used to reduce the TCO of the storage network.

Every enterprise needs to consider a primary and secondary storage network to implement an ILM environment. For some enterprises, iSCSI can be their single storage networking solution.



⁷ Serial Attached SCSI. See the issue of *Clipper Notes* dated February 28, 2007, entitled *Breaking the I/O Paradigm - SAS Enters the Nearline Storage Race*, available at <http://www.clipper.com/research/TCG2007033.pdf>.

⁸ This may not be so significant, as the average commodity server has a utilization rate of around 35% for its CPU cycles.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

David Reine is Director, Enterprise Systems for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial “123” when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, and “clipper.com” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, *“Navigating Information Technology Horizons”*, and *“teraproductivity”* are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.