



# ISCSI vs. Fibre Channel Explained: ISCSI takes its rightful place beside Fibre Channel

13 Jul 2007 | Stephen J. Bigelow

The landscape of storage area networks (SAN) is changing, and the established perceptions of Fiber Channel (FC) vs. iSCSI SANs are blurring rapidly. These days, you can build a Fibre Channel SAN for under \$10,000, but you can also spend \$1 million on an iSCSI SAN or run database applications from network attached storage (NAS). There are no absolute rules on what you can and can't do when it comes to Fibre Channel and iSCSI. However, there are some important pros and cons to each approach that you'll need to consider.

## **What are the technological differences between iSCSI and FC?**

All SANs connect servers to shared block storage arrays through a dedicated high-speed network of host bus adapters (HBA) and switches. The resulting storage network allows servers and storage to communicate, and multiple pathways can be established to ensure storage availability by enhancing redundancy and improving performance. The SAN is also configured to limit the visibility of storage to particular servers or applications.

SANs have traditionally employed the Fibre Channel protocol using the physical implementation and signaling detailed in ANSI standard X3.230-1994 (ISO 14165-1). Simply stated, Fibre Channel uses optical fiber, coaxial copper or twisted pair copper cabling to carry SAN data at speeds of 1 Gbps, 2 Gbps, 4 Gbps and more recently, 10 Gbps. Fibre channel can operate in point-to-point, switched and loop modes.

However, Fibre Channel has been widely criticized for its expense and complexity. For example, a specialized HBA card is needed for each server. Each HBA must then connect to a corresponding port on a Fibre Channel switch -- creating the SAN "fabric." Popular HBAs include the Atto Technology Inc. Celerity FC-44ES HBA, the Emulex LPe1150 PCI Express 4Gbps HBAs, the LSI Logic Corp. LSI7404EP-LC HBA and the QLogic Corp. QLA2462 4 Gbps Fibre Channel HBA. Brocade Communications Systems Inc. and Cisco Systems Inc. provide a range of high-performance intelligent SAN switches.

Every combination of HBA and switch port can cost thousands of dollars for the storage organization. Once LUNs are created in storage, they must be zoned and masked to ensure that they are only accessible to the proper servers or applications. These processes add complexity and costly management overhead to Fibre Channel SANs.

The cost and complexity of Fibre Channel has kept SAN deployment out of reach for small and mid-sized businesses until the introduction of Storage over IP (SoIP) SANs based on the iSCSI protocol ratified by the Internet Engineering Task Force (IETF) in 2003. There is nothing new about the idea of sending storage data over an IP network; the FCIP and iFCP protocols specify the means of sending Fiber Channel data over IP networks. But, iSCSI is the first protocol that allows native SCSI commands end-to-end over IP.

iSCSI emphasizes the idea of a "pervasive" Ethernet environment. That is, every organization from the smallest home network to the largest enterprise uses Ethernet LAN technology that is well understood and very inexpensive. In actual practice, an iSCSI SAN should employ good-quality network interface cards (NIC) and Ethernet switches in a segregated network. Some organizations try to improve iSCSI performance by deploying Ethernet NICs with TCP/IP offload engine (TOE) features to reduce the CPU demands for iSCSI command processing. But at the most basic level, an iSCSI SAN can be implemented using existing NICs and switches that are running on the LAN now. Today, iSCSI SANs operate at 1 Gbps Ethernet speeds, though this could increase to 10 Gigabit Ethernet (GigE) as NICs and switches are upgraded to accommodate 10 GigE.

The appeal of iSCSI is easy to understand. Instead of learning, building and managing two networks, an Ethernet LAN for user communication and a Fibre Channel SAN for storage, an organization can use its existing Ethernet knowledge for both LAN and SAN. "We believe that having an end-to-end Ethernet IP environment provides a lot of value in terms of skill sets, equipment, and commonality in services, capabilities and software tools," says Tony Asaro, senior analyst with the Enterprise Strategy Group in Milford, Mass.

It's important to understand that an iSCSI SAN is not the same as NAS, even though both use the same IP/Ethernet network. An iSCSI SAN provides block-level access to data (it gives you a disk drive) where NAS provides file-level access to data (it serves up a file). The choice of iSCSI or NAS will depend on the needs of applications that are accessing the storage.

### **What are the performance differences with iSCSI and FC?**

From a practical standpoint, Fibre Channel and iSCSI are about equally able to handle storage applications, though experts agree that iSCSI may reveal performance limitations

with the most demanding applications. "Most people compare bandwidth with performance, which is not a direct correlation," Asaro says. "It's only in bandwidth-intensive applications where an Ethernet environment might run out of bandwidth and become a [performance] bottleneck." More specifically, extremely demanding applications, like OLTP handling a large number of small transactions, may be adversely affected by the packet overhead in an IP environment. Ultimately, a properly designed iSCSI infrastructure and storage platform should be able to handle any storage application currently handled by a Fibre Channel SAN.

In fact, concerns about iSCSI performance and reliability are largely misplaced. "I will tell you that the highest performing SAN that I ever saw was an iSCSI SAN, not a Fibre Channel SAN," says Stephen Foskett, director of data practice at Contoural Inc. in Mountain View, Calif. Users are also deploying iSCSI in core applications. Enterprise Strategy Group (ESG) research found that 50% of iSCSI early adopters are using iSCSI for mission critical applications, which is a strong statement of support for iSCSI reliability.

One potential problem with Ethernet performance is the common practice of oversubscription. Most Ethernet servers do not need high performance, so almost all Ethernet switches are oversubscribed. Experts note that a port may be oversubscribed as much as 10 to 1. When building a high-performance iSCSI SAN, those oversubscribed switches may not be able to handle the load, so select high-end Ethernet switches for deployment within an iSCSI SAN.

iSCSI performance may also be influenced by the choice of software-based iSCSI initiator or purpose-built iSCSI HBA, instead of a conventional NIC. Dedicated iSCSI HBAs include Alacritech Inc.'s SES2100 Accelerator card, the Magic 2028-4P 1 Gbit Copper TCP/IP Accelerated NIC from LeWiz Communications Inc. and the QLogic Corp. QLA4050C iSCSI HBA. Most software-based initiators perform well for general-purpose computing, so they are not essential now. But, the broad introduction of 10 GigE may require a shift toward hardware-based initiators. For now, it's important to use the best and most mature iSCSI initiator software available.

Another wrinkle with iSCSI initiators may appear in virtualization. Foskett points out that the VMware ESX version 3 initiator has received performance complaints, particularly in large-scale deployments, typically because VMware is extremely demanding of CPU and storage. "Once you have a dozen servers sitting on one VMware ESX box, you can 'really' use a lot of storage performance," Foskett says. The suggestion would then be to buy the dedicated iSCSI HBAs, but VMware does not currently support iSCSI HBAs, (though this is expected to change in the near future. In this case, your only real option to remedy performance problems would be to implement Fibre Channel

instead.

### **What are the implementation differences between iSCSI and FC?**

The simplicity of iSCSI, and its use of existing Ethernet components, makes iSCSI easier and faster to deploy. Asaro cites ESG research that suggests a 10%-to-30% savings in iSCSI capital expense ("capex") and ongoing operational costs over Fibre Channel. With recent advances, however, it is actually possible to spend more money on iSCSI than on Fibre Channel.

Foskett notes that the emerging crop of iSCSI storage arrays is implementing aggressive architectures that perform well and scale easily. This is heralding integrated features, like thin-provisioning, subdisk RAID and automated tiered storage, features that may not be readily available in Fibre Channel arrays. iSCSI arrays are also noted for their scalability, making it easy to buy and deploy additional iSCSI arrays over time with little (if any) direct management. "In practice, it doesn't tend to scale past six [iSCSI arrays], but the first few scale nicely while giving you a lower price point to get into it," Foskett says.

### **What are the security issues with iSCSI and FC?**

The biggest implementation differences involve security. Contrary to popular belief, Fibre Channel SANs are traditionally less secure than iSCSI. Experts note that the authentication protocols native to Fibre Channel are rarely used. Instead, most storage organizations rely on the fundamental differences in Fibre Channel fabrics and the complex nuances of LUN zoning and masking to keep SAN data secure. iSCSI actually has more security features than Fibre Channel. "From authentication to encryption, you find that iSCSI has many more options and that they're generally easier to use," Foskett says. "But nobody is using them anyway."

In a Fibre Channel SAN, you must establish logical relationships (zones) that connect servers and storage, then block (mask) all but the authorized volumes on any given disk. By comparison, iSCSI does not use zoning. This is often perceived as a security problem, but iSCSI deals with "targets," so it's only necessary to mask targets. This means that a Fibre Channel SAN can see multiple LUNs on any particular disk, but iSCSI can only deal with a disk target. Consequently, iSCSI authentication is very important, and iSCSI employs advanced authentication methods to establish security, such as Challenge-Handshake Authentication Protocol (CHAP). "They use CHAP, which is just a much more secure method and it's really super simple to set up because people have been using CHAP in the IP world for a decade," Foskett says. Further, Fibre Channel does not support native encryption over the wire, but iSCSI can utilize IPsec encryption to protect data in flight.

Experts agree that security is vastly improved by blocking off the SAN from the outside

world. This was a natural element of Fibre Channel, but presents a challenge for Ethernet-based SANs -- you don't want iSCSI SAN data "leaking" out over the user LAN. It is possible to build a different LAN and use it as a dedicated SAN, and that may be the preferred tactic when iSCSI performance must be optimized. However, it is far more common to establish an iSCSI SAN using a virtual LAN (VLAN) that carves up the physical LAN into a logical portion that is used exclusively by the SAN, allowing administrators to tightly regulate and guard the traffic that the VLAN carries.

### **How can a company make iSCSI and FC work as a mix?**

For many organizations, the choice is not Fibre Channel or iSCSI, but rather a mix of the two. A mixed SAN infrastructure has become a popular choice because it preserves any existing Fibre Channel infrastructure while supporting the introduction and expansion of iSCSI in the enterprise. One popular example of this trend is "SAN inclusion," where secondary applications and servers that may have been too costly to place on the Fibre Channel SAN can now be interconnected into an iSCSI SAN. "They [IT staff] might go back and say: Look, we paid \$50,000 for our Unix server, and we're paying \$5,000 for Linux servers -- why don't we use iSCSI in that [low-cost] environment?" Asaro says.

SANs can also be interconnected using an iSCSI gateway, a Fibre Channel switch with iSCSI support, intelligent storage switches and gateways, and multiprotocol storage arrays. iSCSI gateways are simple and unobtrusive (though they can be expensive). Gateways perform all of the translations between iSCSI and Fibre Channel. Examples of these iSCSI gateways include Brocade Communications Systems Inc.'s iSCSI Gateway, Cisco Systems Inc.'s MDS 9216i, Emulex Corp.'s 725/735 iSCSI Storage Routers and QLogic Corp.'s SANbox 6140 Intelligent Storage Router.

By including iSCSI support in the Fibre Channel switch, it's easy to add intelligent features, like the Virtual Router Redundancy Protocol (VRRP) or iSCSI Server Load Balancing (iSLB). Such integration also offers a single management console with the redundancy and performance of an intelligent switch. For example, Brocade offers the SilkWorm FC4-16IP iSCSI blade for its SilkWorm 48000 Director. Cisco provides the IP Storage Services Module and the Multiprotocol Services Module for its MDS 9200 Series Multilayer Fabric Switches and MDS 9500 Series Multilayer Directors.

Intelligent storage switches and gateways add advanced storage services, like virtualization, snapshots, replication and mirroring. Network Appliance Inc.'s (NetApp) V-Series gateways and the Sanrad iSCSI V-Switch are two popular examples of intelligent storage controllers with no storage attached, allowing for iSCSI, Fibre Channel and NAS connectivity to a storage pool. Multiprotocol arrays can also offer the same features but include storage in the same box. For example, EMC Corp. offers mixed

protocol support in its Clariion CX3-20 and CX3-40 arrays, while Hewlett-Packard Co. (HP) supports iSCSI in its StorageWorks XP and EVA arrays.

Regardless of the way you choose to merge Fibre Channel and iSCSI SANs, there should be no performance penalty in either side. But, experts stress that there are no performance guarantee, particularly in the iSCSI deployment. For example, iSCSI target drivers can vary a great deal in their implementation, so some optimization may be required. The IT staff can help to analyze and optimize network performance for iSCSI.

### **Examples of companies making choices regarding iSCSI vs. FC**

Financial institutions, like Texas Trust Credit Union in Grand Prairie, Texas, face a dual challenge of accommodating spiraling storage demands, especially with document images, while meeting retention and regulatory requirements. The core business relies on an IBM P-series database and application server running Unix with hundreds of gigabytes of internal storage. But there are other storage hogs to contend with. Microsoft Exchange for email and EMC Legato software for document management have swelled the total storage demand to about 2.5 terabytes (TB). "We see our storage needs in that area increasing for the next several years; rather dramatically," says Boyce Crownover, system administrator at Texas Trust. Today, those applications are implemented on a modular [Cutting Edge ISS](#) NAS/iSCSI storage device running [the custom Cutting Edge OS](#).

While the choice of iSCSI involved many parameters, Crownover cited compatibility and performance as the two most important criteria. Ethernet copper cabling was already installed in the infrastructure, greatly simplifying the installation and supporting a wide range of Ethernet-based storage systems. The choice of iSCSI also meant a more substantial role for network personnel, rather than strictly storage professionals. Perhaps even more important, iSCSI presented adequate performance for running database and other applications. "iSCSI has the same availability as any NAS," Crownover says. "If we're using databases, which we do quite a bit, then we expect the iSCSI to pay off significantly over other file-level network types [NAS]."

There is always an element of uncertainty in any new technology deployment. Testing resources were limited, and Crownover's team visited sites with Fibre Channel, iSCSI and mixed infrastructures to determine the best fabric for its estimated throughput. "Honestly, if we had to put our primary database system on some sort of SAN, we might have gone Fibre anyway," Crownover says, noting that the disk storage already attached to the P-series provided adequate I/O capability. "We didn't have a need for anything [performance] that would exceed what we expected to get from iSCSI."

With only about three months in live production, Crownover says that there are no lingering issues. Any early concerns over the longevity of iSCSI have vanished as the

industry continues to support iSCSI growth. "That was one of the factors for considering options besides Fibre," he says. "The difference between TCP/IP vs. Fibre technology improvements over the last few years seems to have dramatically favored TCP/IP." The possibility of moving to 10 GigE in the future is another powerful advantage that weighed into iSCSI deployment.

Beyond Exchange servers and other everyday applications, civil engineering firms also face the challenge of storing particularly large, data-rich files and databases, yet still ensure adequate performance for application users. As an example, a typical AutoCad project might use files that are hundreds of megabytes in size, and this requires speedy but cost-effective storage. For the Timmons Group, a civil engineering and environmental design firm based in Richmond, Va., the answer came in an EMC Clariion CX3-20 -- this facilitates a mixed environment that supports 3 TB of corporate storage across both Fibre Channel and iSCSI SANs.

The introduction of iSCSI into the environment has brought significant flexibility while maintaining cost effectiveness. "iSCSI allows us to be dynamic, because there's no cost for Fibre Channel HBAs, it's just a NIC," says Bryan Moore, IT infrastructure manager with Timmons Group. "It's very quick. We just carve out a LUN on the back-end SAN and attach it via iSCSI. We can even get redundant with the iSCSI." That level of flexibility allows the organization to change as its client's needs change.

Testing was limited, though Moore's team was able to do some iSCSI pilots to stress features like speed, redundancy and backup. "We found out that you can't do Fibre [Channel] and iSCSI on the same host," Moore says. "I'd say it was a good week's worth of piloting and testing with users." The testing process involved two dedicated engineers, but the mixed storage infrastructure was ultimately deployed to everyone's satisfaction.

Moore notes that Fibre Channel deployment was not terribly difficult, but ensuring the right firmware, hardware drivers, HBA emulator versions and other details could complicate HBA installation. Fibre Channel HBAs also need downtime for installation and configuration. By comparison, iSCSI deployment proved much easier. "Install Microsoft initiator tools, set some IPs with your NICs, carve out a LUN, and you're ready to go," Moore says. Into the future, Moore looks forward to migrating the iSCSI side of storage to 10 GigE. NIC and switch migration should be straightforward, as long as EMC follows through to provide a 10 Gbit iSCSI module for the CX3-20, which is currently expected.

### **What does the future look like for iSCSI and FC?**

There is absolutely no doubt that iSCSI has become the SAN of choice for mid-sized organizations. Even large enterprises are testing the waters with iSCSI deployments in

workgroups or remote offices to gain a cost advantage. Asaro notes that international IT professionals with little, if any, Fibre Channel experience may leapfrog right over Fibre Channel to adopt iSCSI. However, nobody sees the end of Fibre Channel anytime soon. Few organizations are willing to discard their existing Fibre Channel infrastructure investment and experience base. The shift is happening now and will continue into the future. "We expect to see more and more of a coexistence between them," Asaro says. "We expect that over time iSCSI will be the dominant SAN protocol, but it's probably going to take another three-to-five years for that to occur."

In the near future, Asaro notes that virtual server adoption should be an important catalyst for pervasive iSCSI storage networking in large and midsized organizations. The use of virtualization products, like VMware, allow for massive server consolidation, using a handful of physical servers to host many more virtual servers. "Once I do that, I'm going to put all of my VMware images, applications and data onto a storage network," Asaro says, citing a drive toward universal storage adoption with iSCSI between the virtual server systems and storage.

While both technologies are advancing together, experts expect the appearance of 10 GigE to have little direct impact on iSCSI adoption. Some Fibre Channel devices already offer speeds to 10 Gbit, so the users that embrace iSCSI at 1 Gbit today will probably not be significantly more inclined toward iSCSI adoption. "The two technologies seem to be maintaining parity for the time being," says Phil Goodwin, president of Diogenes Analytical Laboratories in Erie, Colo. "They're continuing to advance abreast of one another."