



IP Storage - Frequently Asked Questions

- Straight forward answers to IP Storage questions
- Discover which applications are optimal for IP Storage
- Learn the basic elements of an IP SAN
- Understand how simple IP Storage can be

The emergence of IP Storage solutions means that IT professionals now have a broader range of storage solutions to address a range of applications and environments. Storage consolidation, simplified data protection, affordable disaster recovery, and improved data management are examples of IP Storage deployments, and business applications running on midrange and smaller servers are today's sweet spot. The pace at which IP Storage is moving into the mainstream is swift and indicative of pent up demand for storage systems that are easy to deploy, highly scalable and affordable. And the pace will accelerate as the large storage vendors continue to broaden their IP Storage products families. The time is now right for all IT executives to take advantage of the new networked storage opportunities available for them, and find out how IP Storage can address some of their most critical data management pain points.

WHAT IS IP STORAGE?

In IP Storage, TCP is used to facilitate a storage interconnect to transfer block level data across an IP network. An Internet Engineering Task Force (IETF) working group, the IP Storage (IPS) working group, is developing a set of standards proposals that address the use of TCP/IP in transporting storage traffic. The iSCSI (Internet SCSI), iFCP, and FCIP protocols are examples of the protocol standards being developed by the IETF IPS working group.

WHAT IS THE DIFFERENCE BETWEEN IP STORAGE AND ISCSI?

IP Storage is a technology that encompasses block storage transfer over an IP network. iSCSI is a protocol standard for encapsulating SCSI in TCP/IP for transfer over an IP network.

WHAT ARE THE ADVANTAGES OF AN IP SAN OVER A FC SAN?

Costs can be significantly cheaper and leverages tools and infrastructure that many administrators are already familiar with.

I'M EXPLORING ISCSI AND IP SANS FOR MY IT GROUP. WHAT TYPES OF INFORMATION SHOULD I TRY TO OBTAIN FROM VENDORS WHEN INVESTIGATING THEIR PRODUCTS?

First, it helps to have a clear picture of the ways you plan to use a SAN. Some installations may focus more on centralized backup protection, restores and disaster recovery. Others may see a SAN as a natural component in their organization's larger server/storage consolidation project. Still others may be experiencing rapid storage growth in their database applications and need a SAN to more easily scale the associated, primary storage. Once you understand the ways you plan to use your SAN, you should then ask each prospective SAN vendor to walk you through the ways their product could be used for your needs. Some specific questions you may want to ask vendors include the following:

- How does your solution provision the storage needed for my applications?
- If I have to scale up my capacity and quickly add new storage to accommodate the needs of my business, how does your solution allow me to do that?
- What mechanisms or tools does your solution use to help manage the unified storage pool in the SAN?
- What features are available to ensure the SAN provides high levels of data availability (i.e., automatic failover, clustering, and replication)?
- What types of data protection and backup support features does your solution offer to protect my data in the event of a disaster (i.e., point-in-time or snapshot functionality, remote copy features, etc.)?

IN ITS EARLY DAYS, IP STORAGE WAS SAID TO SUFFER FROM PROBLEMS RELATING TO BOTH LATENCY AND DETERMINISM - HAVE THESE ISSUES NOW BEEN RESOLVED SUCCESSFULLY?

When the protocols were being defined (well before any products shipped), there was a certain amount of debate in the industry over whether storage over Ethernet was even practical. "Ethernet", in this context, was largely interpreted as "a wide area network comprising unknown numbers of components, switches and routers" - a very different environment from a serial Fibre Channel interconnect. The problems were theoretical, since no products were shipping, and besides, it was the wrong model. In practice, most IP-native SANs are private switched networks using Gigabit Ethernet switches and cabling (they are not mixing storage traffic with the corporate LAN or the WAN traffic). Latency and determinism are not issues. When using IP Storage to bridge geographically separated Fibre Channel SAN environments, the latency requirements of the application certainly have to be carefully considered - but that's all it is, a consideration, not an "issue".

DOES IP STORAGE HERALD THE CONVERGENCE OF THE LAN AND THE SAN? IF SO, OVER WHAT TIMEFRAME?

We will likely increasingly see use standard ubiquitous networking infrastructure to build networked storage

environments, but operational convergence is probably a long way off.

ARE THERE NOW NATIVE IP-ONLY SANS BEING DEPLOYED?

Yes, many customers are deploying native IP-only SANs to address a number of applications, with support for Microsoft applications, such as Exchange, SQL and basic file serving, being very popular. Increasingly they are being deployed in Linux and UNIX environments.

WHAT DOES THE ROADMAP LOOK LIKE FOR IP STORAGE?

We are now starting to see second and third-generation iSCSI products hit the market. First-generation iSCSI storage products targeted their feature set at shared storage for modest numbers of smaller servers, delivering the performance and availability requirements of mid-range business applications.

Second and third-generation iSCSI solutions are now expanding this "sweet spot", adding high-availability support - with multi-pathing, failover, and host cluster services; SAN boot support for dense server environments; support for virtual server environments; and performance enhancements – in initiators, targets, and in the I/O path with multi-connection sessions. In addition, we have 10GbE on the horizon.

WHAT NEW FEATURES WILL DRIVE BROADENED ADOPTION OF ISCSI IN 2007?

The key drivers for next-generation IP SANs includes features such as Multi-Connection Sessions and ErrorRecoverLevel>0. These features enable iSCSI based-IP SANs to make use of the built-in parallelism and fault tolerance functions of TCP/IP and standard Ethernet ports for higher performance and higher availability. Higher performance is achieved with sessions that span multiple 1 Gb paths; and simple active/active high availability is accomplished without the expense and complexity of special HBAs and host multi-pathing software. These enhancements will drive IP SANs into more demanding application environments that require higher performance workloads and higher levels of data availability.

WHAT NEW TRENDS FOR DO YOU EXPECT IN 2007? WHAT WAS SIGNIFICANT FOR 2005 AND 2006?

2005 saw iSCSI storage solution become mainstream in Windows server environments. In 2006, there was a growing trend too support virtual server environments. In 2007 we will see rapid expansion into departmental Linux and UNIX environments, particularly for blade servers and small-to-midrange hosts

IN 2007 WHAT TYPES OF COMPANIES WILL DEPLOY IP STORAGE SOLUTIONS?

Both large and mid-size companies will utilize IP SANs for servers that are not currently supported by FC SANs or in locations where FC SAN expertise is not available due to cost and complexity. This will include data centers, departments, branch office and remote offices.

IS IT FAIR TO SAY THAT IP SANS ARE 'LOWEST COMMON DENOMINATOR' SOLUTIONS WHEN COMPARED TO FC SANS - AGAIN, DOES THIS MATTER?

No - in fact, today the opposite is true. Many entry-level Fibre Channel arrays lack the sophisticated data management capabilities of their mid-range and high-end cousins (such as snapshots, clones, remote mirroring). However, almost all the iSCSI-native arrays shipping today include these capabilities. And in fact, that's what customers are buying. It's not just about inexpensive storage

WHAT IS FIBRE CHANNEL OVER ETHERNET (FCOE)?

This is a proposal for a new protocol specification which has been submitted to the INCITS T11 Committee as a possible work item. It appears to be a means of transmitting Fibre Channel commands natively over Ethernet by using an optional transport mechanism instead of TCP/IP, while maintaining backward compatibility with existing Fibre Channel endpoint infrastructure. It is targeted at data center SANs using a "converged" (10 Gigabit) Ethernet network infrastructure. If it gains traction in T11, there is a good chance that it will emerge as a new ANSI standard in the 2010-2011 timeframe.

HOW DOES FCOE DIFFER FROM ISCSI?

iSCSI is a standard SAN protocol ratified by the IETF in 2003, which transports SCSI block storage commands over Ethernet using TCP/IP. iSCSI was designed to take advantage of all the ease-of-use, availability and guaranteed delivery mechanisms provided by TCP/IP, while providing a seamless path from 1 Gigabit Ethernet to 10 Gigabit Ethernet and 100 Gigabit Ethernet. iSCSI is being used as a standard SAN storage protocol in tens of thousands of IT organizations around

the world today for both Gigabit Ethernet and 10 Gigabit Ethernet connections. Market growth is expected to continue to exceed 100% in the coming years.

WHAT ARE THE IMPLICATIONS OF THIS PROPOSED NEW FCOE STANDARD?

The most obvious implication is the inevitability of 10Gigabit Ethernet as a data center SAN interconnect. We are already seeing interest as leading-edge IT organizations are starting to implement iSCSI over 10Gigabit for port aggregation in large-scale data center environments. You can have it today with iSCSI, or wait 4 years and maybe get it with your legacy Fibre Channel equipment.

WHAT IS iFCP?

The Internet Fibre Channel Protocol (iFCP) supports Fibre Channel layer 4 FCP over TCP/IP. It is a gateway-to-gateway protocol where TCP/IP switching and routing components complement or replace the Fibre Channel fabric.

WHAT IS FCIP?

The Fibre Channel over IP protocol provides a mechanism to "tunnel" Fibre Channel over IP- based networks. This enables the interconnection of Fibre Channel SANs, with TCP/IP used as the underlying wide-area transport to provide congestion control and in-order delivery of data. Standard Fibre Channel fabric services are used for switching and routing.

FCIP is a point-to-point tunneling protocol for sending FCP (Fibre Channel Protocol) data over IP networks (typically a WAN). It is primarily targeted at connecting geographically separated FCP environments (for instance, two FC SAN islands in separate data centers). This is also complementary to iSCSI.

HOW DOES ISCSI RELATE TO IP STORAGE PROTOCOLS SUCH AS iFCP?

iSCSI is an IP-native storage protocol. Internet Fibre Channel Protocol (iFCP), in contrast, is a protocol for running Fibre Channel traffic over a TCP/IP network through gateways. iFCP is complementary to iSCSI in that it can also be used to connect iSCSI hosts to existing Fibre Channel SANs, which in some environments might be an appropriate evolutionary step to native IP storage.

WHERE IS THE BEST PLACE TO OBTAIN INFORMATION CONCERNING THE DEVELOPMENT OF THE ISCSI PROTOCOL?

All of the specifications concerning the iSCSI protocol are on the IETF web site (www.ietf.org). The Storage Networking Industry Association IP Storage Forum (www.ipstorage.org) web site is another good place for IP Storage information. There are texts such as *How IP SANs: A Guide to iSCSI, iFCP, and FCIP Protocols for Storage Area Networks*, by Tom Clark, Addison Wesley, 2001. Also, the Storage Networking Industry Association has a list of many good reference articles and books about IP Storage and storage networking in their resource center (www.snia.org).

MOVING ONTO THE SNIA IP STORAGE FORUM, PLEASE GIVE SOME BACKGROUND AS TO HOW LONG IT HAS BEEN GOING, WHAT IT HAS ACHIEVED TO DATE - AND IS IT GENUINELY TECHNOLOGY AND VENDOR AGNOSTIC - OR DO SOME VENDORS, AND ISCSI, TEND TO DOMINATE?

The IP Storage Forum was formed 5 years ago to promote the adoption of the IP Storage protocols (FCIP, iFCP and iSCSI). For the first few years, the Forum focused on education about the new protocols as they were being developed in the standards community. Over the past two years, the focus has moved to promoting the adoption of products based on these (now) standard protocols. As with other SNIA education activities we have a rigorous approvals process to ensure vendor neutrality, and an outreach agenda, which represents all three protocols, and the interests of all our member companies. Feedback from IT end users attending the seminar series we conducted in North America during 2004 was that we have been very successful in our neutrality goals.

SNIA Europe is an independent regional affiliate of the Storage Networking Industry Association, with separate governance and membership dues, it made sense to create a SNIA-Europe interest group focused on IP Storage - the SNIA Europe IP Storage Initiative (IPSI). You can think of it as a sister organization to the IP Storage Forum, focused on localizing the educational and promotional materials and initiatives created within the Forum, and delivering them within Europe.

WHAT IS ISCSI AND HOW DOES IS WORK?

iSCSI (Internet Small Computer System Interface) is a storage networking protocol based on the Internet Protocol (IP) that was developed by the Internet Engineering Task Force (IETF). iSCSI encapsulates SCSI block storage commands into

Ethernet packets for transport by TCP over IP networks. This enables servers to communicate with shared storage devices over standard IP infrastructure using standard SCSI storage commands. The iSCSI protocol specification is an IETF standard (RFC 3720) that was ratified in February 2003.

The iSCSI protocol defines the rules and processes to transmit and receive block storage data over TCP/IP networks. iSCSI replaces the parallel SCSI direct cabling scheme with a network fabric. At the physical layer, iSCSI is media independent and will support any media that supports TCP/IP. Servers and storage devices that support iSCSI interfaces can be directly connected to an existing IP switch and router infrastructure. Most first commercially available iSCSI implementations have implemented iSCSI over Ethernet, the most widely deployed networking infrastructure. The iSCSI protocol itself, resides above the physical and data link layers of the network model and interfaces to the operating system's standard SCSI Access Method command set. iSCSI enables SCSI-3 commands to be encapsulated in TCP packets and delivered reliably over IP networks.

Unlike direct-attached storage, IP SAN (iSCSI) storage solutions enable businesses to share storage resources across servers and to easily expand storage capacity. Because Ethernet is widely deployed and understood, iSCSI enables businesses to tap the data management and flexibility benefits of networked storage without the high costs and complexity typically associated with traditional Fibre Channel SANs. With iSCSI, companies with distributed environments can use existing network investments and expertise to cost-effectively deploy data center storage with sophisticated management capabilities.

WHY IS ISCSI IMPORTANT?

Although just a data transport protocol, iSCSI opens up a new world of networked storage for many IT groups who were previously unable to afford a SAN or who were concerned about the complexity associated with Fibre Channel. iSCSI was officially ratified as a protocol by the Internet Engineering Task Force (IETF) in 2003. In that same year, Microsoft gave iSCSI a formal nod of approval with the release of its free iSCSI initiator/device driver. By 2003, a host of other vendors had also released iSCSI-enabled products.

WHAT WAS THE REASON FOR SUCH GROWTH IN ISCSI ACCEPTANCE AND ACTIVITY?

Many IT managers and vendors started to see iSCSI as the first standards-based, low-cost and easy to deploy technology available to help them build storage area networks on top of their existing Ethernet networks.

WHAT TYPES OF STORAGE TRAFFIC DOES ISCSI MOVE?

iSCSI supports the movement of block-level storage data over IP networks. File-level access is achieved through the use of file servers or NAS products connected to the IP network.

WHAT ARE THE PRIMARY BENEFITS OF ISCSI?

There are a number of benefits that iSCSI offers customers just by being a storage area network technology. For example, storage area networking technologies greatly increase the user's ability to utilize more storage resources, as opposed to having storage on a single system that cannot be accessed by other systems or remote users. Customers can enjoy the benefits of storage consolidation (better disk capacity utilization, fewer points of storage management, consolidated backup/restore). iSCSI enables customers to cost-effectively build and manage storage area networks based on technologies that they already support and understand (i.e., SCSI, IP, Ethernet, SNMP). The physical IP infrastructure, management tools, and support staff that installs and maintains the IP infrastructure can be leveraged by using iSCSI for storage networking. iSCSI can also take advantage of the inherent distance advantages that IP networks provides to LAN, MAN, and WAN environments.

WHAT ARE SOME TYPICAL ISCSI SUPPORTED APPLICATIONS?

Typical applications tend to be business-critical enterprise applications at the departmental level, though you do see mission-critical applications for smaller enterprises. The most popular applications in these environments include email, file serving, small DBMS systems, and business applications running on top of smaller databases. Messaging, Web, ecommerce, and some technical applications are also proving to be a particularly good fit with iSCSI environments. Finally, customers with home-grown applications running on smaller servers with internal or direct-attached storage are finding iSCSI a great storage consolidation solution, since the migration to an IP SAN is affordable and relatively painless.

iSCSI also enables truly affordable disaster recovery, backup, and secondary storage solutions.

WHAT CUSTOMER PROBLEMS DOES ISCSI SOLVE?

Many companies have a compelling need for low-cost and easy to deploy SAN solutions to consolidate the data in their proliferating Windows® server environments and to provide centralized data management for their Linux® server farms, their stranded UNIX® servers, and their Novell environments. This situation is exacerbated by the fact that in many cases the applications running on these servers (e-mail, small databases, CRM, and analytic) are now considered business critical. iSCSI provides an excellent SAN solution for these customers.

iSCSI enables affordable storage consolidation solutions, particularly in environments populated with low-cost servers, where simplicity, flexibility, and price/performance are critical IT decision factors.

IS THIS NEW TECHNOLOGY READY FOR REAL-WORLD DEPLOYMENTS?

Current deployments of iSCSI span a broad range of industries and IT environments. Not only is it ready, but iSCSI-based SANs are now considered a mainstream storage networking solution by the leading storage analysts, who track iSCSI storage system shipments. NetApp alone has more than 4,000 customers using IP SAN solutions in production IT environments.

WHAT DOES AN ISCSI SOLUTION CONSIST OF?

iSCSI solutions consist of iSCSI "targets" (native iSCSI storage devices) connected to iSCSI "initiators" in the host systems via standard Ethernet infrastructure (switches and cables). iSCSI initiators can be software drivers that come with the host operating system and use standard Ethernet network interface cards (NICs) for the physical network connection, or they can be iSCSI host bus adapter cards or TNICs (TCIP Off-Load NIC) that provide the network connection and offload iSCSI and TCP/IP processing from the host CPU.

HOW IS ISCSI MULTI-VENDOR INTEROPERABILITY BEING ADDRESSED?

Both the SNIA and UNH Interoperability Lab host a number of multi-vendor interoperability demonstrations and tests. The Storage Networking World Interoperability Lab, sponsored by SNIA is another good place to see iSCSI products working together.

IF ISCSI IS AN ENABLER, WHAT'S DRIVING THE TRANSITION FROM DIRECT-ATTACHED STORAGE?

The driving force is often high data growth and the need for cost control together with improved operations agility and efficiency. In these circumstances the legacy direct-attached storage environment becomes increasingly complex; backup/restore operations become increasingly unreliable, and the storage environment is unable to support the demands of the business. On top of all that, regulatory compliance pressures, the need to integrate geographically-dispersed data assets, and the availability of effective ILM solutions also feature as major motivators for SAN installation.

HOW MANY, AND WHAT TYPES OF ORGANIZATIONS HAVE ADOPTED ISCSI TO DATE?

More than 6000 enterprises worldwide are estimated to be using iSCSI-based SAN storage in IT production environments today. These are mostly mid-sized and large enterprises wanting to introduce SAN storage in environments dominated by mid-range and low-end servers.

I'VE HEARD YOU NEED A LOT OF SPECIAL HARDWARE TO RUN THE ISCSI PROTOCOL ON YOUR EXISTING NETWORK. IS THAT TRUE?

No. Most initial iSCSI installations can use much of your existing Ethernet infrastructure: Basic network interface cards (NICs), existing servers and existing network cabling. You need the iSCSI device driver on your server, but this is available as a free software download from companies like Microsoft, Novell and Cisco Systems. You also need one or more iSCSI-enabled storage devices, such as storage arrays, tape libraries, etc. And, if you want a robust IP SAN, you also need to invest in SAN management services and software that comes with many SAN solutions. Some customers have found it advantageous to install a TNIC or iSCSI HBA on the host to both increase transfer rates and increase efficiency of the host CPUs.

ISN'T ISCSI REALLY JUST A WINDOWS PHENOMENON?

It's true that the vast majority of the deployments today are on Windows. However, Novell NetWare and Linux

environments also becoming popular, and solutions for small RISC servers running UNIX are just now starting to emerge. The UNIX vendors were initially slow to embrace iSCSI, but that's changing fast.

IS ISCSI PERFORMANCE GOOD ENOUGH FOR ENTERPRISE APPLICATIONS?

Many customers are pleasantly surprised by performance of iSCSI solutions – even software-only solutions. iSCSI provides perfectly adequate performance levels for the vast majority of business application workloads. The important fact on performance is that iSCSI is almost always significantly faster than direct-attached storage when benchmarked with real-world workloads.

DOES ISCSI IMPOSE LOTS OF CPU OVERHEAD?

CPU overhead depends on the implementation of the iSCSI initiator. A software-only iSCSI initiator uses a standard Ethernet NIC and relies on the host CPU to process the iSCSI commands and the TCP/IP stack. This could consume up to 1GHz of CPU for 1Gb Ethernet. Clearly, this is much more of an issue with older CPUs than with newer ones. With a hardware-assisted iSCSI initiator, the iSCSI and TCP/IP processing is off-loaded to an iSCSI HBA. With a TNIC, the software initiator is utilized, sitting on top of the TCP stack as an application, and like an iSCSI HBA, TCP/IP processing is off-loaded resulting in greatly reduced CPU overhead, comparable to that of a Fibre Channel HBA.

WHAT ARE VENDORS DOING TO ENSURE INTEROPERABILITY OF THEIR ISCSI SOLUTIONS?

Operating system and application vendors often have a catalog of qualified hardware solutions. The Microsoft Windows Catalog, for example, lists qualified iSCSI hardware devices. The leading IP storage vendors work closely with all the OS vendors to ensure continuing compatibility with their software.

CAN THE SCSI PROTOCOL SUPPORT MULTIPLE SPEEDS OF ETHERNET?

Software iSCSI runs on any standard Ethernet connection. 10 Gigabit Ethernet connections are expected to potentially come onto the market in volume in the late 2008 timeframe.

WHAT TYPE OF NETWORK CABLING CAN ISCSI SUPPORT?

iSCSI products support both copper and optical fiber cabling schemes.

WHAT TYPE OF ISCSI PRODUCTS ARE ON THE MARKET TODAY?

There are a variety of products/solutions that are based on various drafts of the iSCSI specifications on the market today. There are iSCSI targets such as disk arrays, storage routers, storage switches and iSCSI initiators. On the host, customers can use standard NICs, iSCSI HBAs or TNICs (NICs featuring TCP/IP offload). More products will reach the market, as the iSCSI specification becomes a standard. iSCSI vendors can be found on the www.snia.org website on the IP Storage Forum page

WILL ISCSI/IP SAN TRAFFIC PUT TOO MUCH OF A LOAD ON MY EXISTING NETWORK TRAFFIC AND BANDWIDTH?

Most networks operate at 30 percent or below of their network bandwidth potential. With common iSCSI implementations, network bandwidth may climb to about 50 percent utilization. In addition, many other factors can impact performance and data throughput rates. The network is seldom the bottleneck with iSCSI implementations. Often, the bottleneck originates either from the application itself, from the process of getting the data on or off a physical disk drive or rare instances where the CPU on the server is close to saturation because of excessive loads (where an iSCSI HBA or TNIC can often help). For installations where IT groups are trying to consolidate storage or centralize backup operations with an IP SAN, many also implement IP SANs in a separate subnet or private network.

IF I WANT TO USE ISCSI, CAN'T I JUST PUT A TAPE LIBRARY SOMEWHERE OUT ON MY NETWORK, INSTALL THE ISCSI SOFTWARE ON MY SERVERS AND BE READY TO BACK THEM UP?

Yes, you can do this. However, this scenario is not a storage area network or SAN, but an iSCSI/network derivative of the traditional direct-attached storage (DAS) model. With iSCSI, you can decouple storage devices from their servers and place them somewhere else on the network. But, you won't gain the benefits of a SAN this way since you can't easily share the decoupled storage across multiple servers. You also can't easily scale the storage independently of the server. Some solutions offer a more basic iSCSI implementation, but without the SAN management tools and SAN ecosystem, you won't be able to take advantage of centralized management and advanced data protection features inherent in today's

robust SANs.

WHAT IS AN INITIATOR?

An initiator is typically a server hosting an application, where the application makes periodic requests for data to a related storage device. Initiators are also referred to as servers or host computers. The iSCSI device driver that resides on the server may also be called an initiator. Initiators “initiate” (or begin) iSCSI data transport transactions by making an application request to send/receive data either to or from one or more storage devices. The application request is immediately converted into SCSI commands, and then encapsulated into iSCSI where a packet and header are added for transport via TCP/IP over either the Internet or traditional Ethernet networks.

WHAT IS AN ISCSI INITIATOR?

An iSCSI initiator is a system (server, workstation, storage device, router, switch) that starts (initiates) the transfer of information to/from an iSCSI target.

WHAT IS AN ISCSI TARGET?

Targets are one or more storage devices that reside on the network. Targets receive iSCSI commands from various initiators (or servers) on the network. On the target’s side, these commands are then broken down into their original SCSI format to allow block data to be transported between the initiator and the storage device. The target will respond to a server’s data request by sending SCSI commands back to that server. These commands are again encapsulated via iSCSI for transport over the Ethernet. Targets can be any type of storage device, such as a storage array that is part of a larger IP-based storage area network, or IP SAN. They could also be a separate tape library residing on either the SAN or elsewhere on the network.

SHOULD I USE A HARDWARE OR SOFTWARE-BASED INITIATOR?

The decision to use a hardware or software iSCSI initiator depends on several factors, including budget, performance requirements, and server workload. Software iSCSI initiators provide the lowest-cost iSCSI solution. A software-only iSCSI initiator uses either a standard Ethernet NIC and relies on the host CPU to process the iSCSI commands and the TCP/IP stack or a TNIC which does provide TCP/IP offload processing. While current-generation servers with 2+ GHz CPUs, iSCSI protocol processing does not impose a significant overhead for many customer workloads, some workloads, such as file serving, can benefit tremendously from TCP/IP offload.

If your server has a Gigabit Ethernet NIC, there is little risk in evaluating a software initiator, because initiators are available at no charge for most popular operating systems.

With older CPUs and heavily loaded servers, a hardware-assisted iSCSI initiator, such as an iSCSI HBA, or a TNIC, making use of the software iSCSI initiator, may be more appropriate. Both choices offload iSCSI and TCP/IP processing. The result is greatly reduced CPU overhead, comparable to that of a Fibre Channel HBA. A TNIC, unlike an iSCSI or Fibre Channel HBA, is capable of also offloading TCP/IP processing for all network traffic

HOW DO HARDWARE AND SOFTWARE INITIATORS IMPACT APPLICATION PERFORMANCE?

For many application workloads on a lightly loaded server, even iSCSI overhead with a software initiator would not noticeably affect application performance. With more heavily loaded servers or demanding workloads, an iSCSI HBA or TNIC would be recommended. In many environments, transport bandwidth and host CPU overhead are not issues—application performance problems are much more likely to be due to the fact that the storage is spindle-limited (too few disk drives for the workload). This is not a protocol or transport issue. It is, however, a common problem with direct-attached storage—a problem that can be easily solved with an iSCSI-based SAN solution. Independent lab tests have shown that iSCSI solutions using software initiators outperform DAS with up to 41% faster throughput and 63% better response time using consolidated real-world and simulated OLTP workloads.

HOW DOES ISCSI INTEROPERABILITY COMPARE TO EARLY FIBRE CHANNEL INTEROPERABILITY?

Fibre Channel interoperability problems were primarily due to two issues. First, the server vendors implemented the SCSI-3 command set differently. Over the past eight years, however, vendors have resolved many of the problems related to this issue.

The second issue is that Fibre Channel lacks built-in networking capabilities, and this has resulted in significant complexity due to proprietary features. An iSCSI environment is inherently less complex than a Fibre Channel environment since it is

built on the standards-based networking capabilities of Ethernet. For these reasons, we have seen none of the interoperability issues that plagued Fibre Channel. Its sweet spot today, however, is with mid-range business applications - particularly those generating significant data growth. For these applications direct-attached storage is increasingly unworkable, but a Fibre Channel SAN is likely to be overkill from a performance, cost or complexity perspective.

FIBRE CHANNEL IS ALREADY DESIGNED TO SEND SCSI COMMANDS OVER A SWITCHED NETWORK FABRIC; WHY BOTHER WITH ISCSI?

Fibre Channel SAN solutions have become the de-facto standard storage networking technology for high-performance, mission-critical applications in large data center environments. However, Fibre Channel is a complex and expensive environment to deploy and support, which has inhibited deployment in distributed, departmental, and workgroup environments. Ethernet-connected IP SANs offer a simpler, more affordable solution for those environments.

HOW IS ISCSI DIFFERENT FROM FIBRE CHANNEL?

Fibre Channel is a high-speed serial interconnect, not a network. Fibre Channel technology has been historically expensive, with no built-in capabilities to do routing and node failover and only primitive address management and security capabilities (although this is slowly changing). In a Fibre Channel SAN, all these capabilities have to be operator configured and host managed. The result is an expensive, complex environment. Since iSCSI is an IP-native SCSI protocol, it can leverage the standard networking capabilities and ubiquity of Ethernet. The result: a simpler, less expensive SAN.

IS ISCSI GOING TO REPLACE FIBRE CHANNEL?

Fibre Channel is a well-established, widely deployed technology with a proven track record, particularly in high-performance, mission-critical data center environments. Many in the industry view Fibre Channel and iSCSI as complementary technologies that will coexist for the foreseeable future. Already it is clear that iSCSI is accelerating the transition to networked storage by providing a compelling alternative to direct-attached storage in distributed, divisional, and workgroup data centers in large companies, and in the core data center for medium-sized enterprises which lack FC SAN expertise.

WHEN SHOULD I DEPLOY ISCSI INSTEAD OF A FIBRE CHANNEL SAN?

For environments consisting of many midrange or low-end servers, an IP SAN (iSCSI) solution often delivers the most appropriate price/performance. The performance of iSCSI and FCP is roughly comparable (both are much faster than direct-attached storage), and for typical business applications most end users won't notice the difference.. Cost, management overhead, and application workload are all considerations that must be taken into account.

HOW DO THE FEATURES OF ISCSI STORAGE ARRAYS COMPARE TO FIBRE CHANNEL ARRAYS?

One early assumption was that iSCSI SAN solutions would simply be cheap SAN storage with limited functionality. This seemed like a reasonable assumption since entry-level Fibre Channel arrays historically have lacked the sophisticated data management capabilities of their mid-range and high-end cousins (such as snapshots, clones, remote copy, mirroring). Interestingly, however, almost all first-generation iSCSI-native arrays provided sophisticated data protection and even disaster recovery capabilities - including snapshots, mirroring, and distance replication.

WHAT IS AN ISCSI-FIBRE CHANNEL STORAGE ROUTER OR STORAGE SWITCH?

iSCSI storage routers or storage switches are gateway devices that support interconnecting other storage topologies (Fibre Channel, SCSI, Infiniband, etc.) to IP networks and enable block level data to pass between the two networks. IP storage switches may also provide switching among the protocols such as Fibre Channel switching, IP/Ethernet switching, or Infiniband switching.

WHAT IS THE DIFFERENCE BETWEEN DR AND BACKUP? HOW ARE THEY IMPLEMENTED?

Both DR and back-up provide data protection. Disaster Recovery deals generally with replication (copying data between sites) and building systems that are capable of bringing up the data and executing transactions at remote site in the event of failure at primary site. Backing-up is generally to protect against data loss, generally focused on a given site. Backups have historically been made to tape, but increasingly, customers are backing up to disks, providing higher speeds within shrinking back up windows and faster recovery times. (Tapes and disks could be removed and placed a different site for safekeeping. In the event that data is lost, a tape, for example, can be retrieved and data reconstructed.)

HOW DOES ISCSI COMPARE TO FIBRE CHANNEL IN TERMS OF SECURITY?

Fibre Channel environments are perceived to be secure since they are controlled private networks. iSCSI is perceived to be less secure since it is Ethernet-based. However, Fibre Channel has essentially no security capabilities, while iSCSI has very extensive security features. The iSCSI spec covers initiator and target authentication (using CHAP, SRP, Kerberos, and SPKM) to prevent unauthorized access and permit only trustworthy nodes. IPsec prevents insertion, modification, and deletion, and IPsec encryption provides privacy and prevents eavesdropping. Furthermore, all the other Ethernet-based security mechanisms (VLANs, firewalls, etc) can be applied.

CAN I DEPLOY A SECURE ISCSI SOLUTION TODAY?

Yes. iSCSI implemented over a private IP storage network (separate from the data communications network) offers security features similar to those of Fibre Channel today. For companies who do not wish to deploy a separate network for their SAN, but want to leverage their existing Gigabit Ethernet infrastructure. Similar levels of security can be achieved by creating a subnet and using VLANs to protect server-to-storage traffic.

AT THE MOMENT, COST OF IP STORAGE MUCH CHEAPER THAN FIBRE CHANNEL - WILL THAT PRICE ADVANTAGE REMAIN?

Storage costs come in a number of categories; acquisition cost for storage, acquisition cost for infrastructure, and ongoing operations costs. The infrastructure costs for IP Storage will likely remain somewhat lower than Fibre Channel (though the advent of IP Storage did cause a reduction in Fibre Channel infrastructure costs) simply because of the ubiquity of the technology. The really big cost savings, though, are in the operations costs and IP Storage will probably continue to be much more economical because it's less complex, easier to use and based on well-understood networking tools and capabilities (every server user knows how to configure Ethernet networks).

I'VE HEARD YOU HAVE TO REALLY WATCH SECURITY ISSUES WITH IP SANS. IS THAT TRUE?

When IP SANs are part of a company's Ethernet network or corporate LAN, IT staff should take the same security precautions they would take with other areas of their network. iSCSI comes with common, IP-based authentication and verification functionality. Likewise, IP SANs based on iSCSI are equipped to work with common security methods and protocols such as IPsec. Security features are discussed further in the previous section of this document. The good news is that IP is a mature protocol with significant security tools and capabilities that have been made available over the years. Many security issues can also be avoided as well by implementing IP SANs on their own subnet or private network.

IF THE PRICE GAP DOES CLOSE OR DISAPPEAR, WILL FIBRE CHANNEL ALWAYS BE A PREFERABLE SOLUTION TO IP IN THE STORAGE NETWORKING FIELD?

Fibre Channel is a great solution for those applications that require the highest performance and availability. However, most business applications do not fall into this category, and IP Storage is perfectly adequate for them. When flexibility, ease of use and availability of human resources are important considerations, IP Storage is clearly a preferable solution.

WHAT'S THE STORY WITH ISCSI BACKUP?

Since snapshot functionality is so common with iSCSI storage, most customers do their backups from a recent point-in-time copy, eliminating the traditional backup window. Data recovery is also extremely fast, when recovery can come from the disk-based point-in-time copy. The savings in administrative overhead and end-user downtime from these features alone can, and often do, justify the entire IP SAN investment.

In addition, the wide availability of remote copy technology is resulting in the widespread use of disk-to-disk backup – with tape becoming an archive medium in a centralized facility. This makes it possible to automate the backup of remote data assets in satellite data centers and remote offices.

WHAT ABOUT IP STORAGE AND DISASTER RECOVERY?

The broad availability of distance replication and mirroring technology with iSCSI storage systems has enabled even small and medium-sized enterprises to implement a practical disaster recovery plan. This usually involves taking snapshots at regular intervals during the day, and using remote copy and/or asynchronous mirroring to copy the changes to a remote system. Since everything is IP-native, there is no need for the expense of gateways and protocol conversions, and the solution can be designed to deliver the required service level of the applications at that particular facility. In addition a primary application of the other IP Storage protocols (FCIP and iFCP) is disaster recovery between separate Fibre Channel SAN environments.

WHAT IS THE DIFFERENCE BETWEEN SOFTWARE ISCSI ON A GIGABIT NETWORK INTERFACE CARD

(NIC) AND USING ISCSI HBAS OR TNICS?

Software iSCSI protocols run in the operating system and allow standard Ethernet NIC to be used for iSCSI traffic. iSCSI adapters (e.g., iSCSI HBAs and TNICS) run optimized logic on the adapter for TCP processing, enabling higher performance data throughput with lower CPU utilization. iSCSI HBAs focus on TCP processing for iSCSI while TNICs are capable of TCP processing for all network traffic in addition to providing standard NIC functionality.

WHAT ARE THE DIFFERENCES BETWEEN ISCSI HBAS, AND TNICS?

For the most part, all iSCSI HBAs and TNICs attempt to optimize the handling of block level storage traffic by offloading the system from doing that CPU intensive protocol processing. Many vendors achieve this by putting the iSCSI protocol processing on the HBA. However, where vendors and industry analysts draw a distinction between iSCSI HBAs and TNICs is that TNICs are dual-purpose, handling of both standard networking traffic and storage traffic acceleration.. Users looking to support both storage and networking traffic from a single adapter should ask their iSCSI adapter vendor how these cards handle the protocol off-load for each type of traffic

WHAT IS TCP PROTOCOL OFF-LOAD OR TOE?

TCP off-load engine is logic that is typically embedded on the network adapter and is used for processing the TCP protocol stack. By off-loading the processing of the TCP protocol stack from the operating system, the system's overhead is greatly reduced, enabling more CPU cycles for processing applications and/or supporting more clients. Modern approaches make use of what is termed partial or dynamic off-load, such as Microsoft's TCP Chimney. In this approach, "stateful" adapters, making use of the OS's TCP protocol stack, are transferred "connection control" to expedite data transfer between the network and an application's memory buffer. This exchange of connection control, to and from a server's CPU and OS and the adapter, executes dynamically based upon TCP header details- pointing to application data packets, or adapter memory buffer conditions.

WHAT IS THE ISCSI API?

The iSCSI API defines an application-programming interface for the management of iSCSI target and initiator devices.

WHAT IS IPSEC?

IP Security (IPSec) is an architecture for device authentication and data encryption over IP networks.

WHAT ARE ISNS AND SLP?

iSNS is a common discovery, naming, and resource management service for all of the IP storage protocols. SLP is the Service Location Protocol for querying lists of network resources.

We hope that you have found this document useful. For more information visit our website at www.snia.org/ipstorage. If you have any questions about the IP Storage Forum or this document, feel free to send an email to ipmarketing@snia.org.

IP STORAGE FORUM MISSION

To further drive the adoption of IP-based Storage Area Network (SAN) solutions through global marketing initiatives and end-user education.

