

USING iSCSI IN STORAGE INTENSIVE MAC OS X®
APPLICATIONS



By David J. Cuddihy
Principal Engineer

May, 2008

TABLE OF CONTENTS

ABSTRACT..... 3
USING MACS IN STORAGE INTENSIVE APPLICATIONS 4
 DAS Challenges 5
 SAN Solutions 5
 What about NAS? 6
USING MACS IN AN ENTERPRISE 7
THE ISCSI ANSWER 7
WHAT IS ISCSI? 9
 iSCSI Background 9
 How iSCSI Works 10
 iSCSI Security 12
 iSCSI Digests 13
SOME ISCSI EXAMPLES 13
 Audio Workstations..... 13
 Tiered SAN for Video Editing 14
 MacBook/MacBook Pro SAN Connectivity 15
ATTO: A LEADER IN ISCSI..... 16
CONCLUSION..... 17

ABSTRACT

Mac® OS X users who face growing storage requirements due to larger, more complex projects and ever-expanding file sizes often turn to Storage Area Networks (SANs) to help overcome their challenges. With benefits such as leveraging existing IT equipment and a lower cost-of-entry when compared to other storage networking protocols, many Mac users are turning to Internet SCSI (iSCSI) SANs as an easy, low-cost way to reap the benefits of a SAN in an OS X environment.

iSCSI is a protocol that transports SCSI commands and data across standard IP networks. Users can connect to an IP SAN via the standard Ethernet ports on their workstations, needing only an iSCSI initiator to access the shared storage. This effectively eliminates the traditional distance limitation associated with SCSI while introducing all the benefits of shared storage.

USING MACS IN STORAGE INTENSIVE APPLICATIONS

Mac OS X users often find themselves working in application environments with demanding storage requirements such as digital content creation. Uncompressed digital audio files, like those handled at digital audio workstations (DAW), use a significant amount of storage while also needing to be accessed quickly during editing sessions to avoid dropouts. Working with digital video requires a substantial increase in storage capacity, efficiency and speed during all phases of production. Digitizing and editing content, enhancing projects with graphics and effects, and color correction all require high-speed access to large amounts of storage.

Traditional digital video and audio production environments consist of a number of workstations, each connected to its own direct attached storage (DAS). DAS typically refers to storage connected to a single workstation using the workstation's built-in connectivity options, such as FireWire, or via host bus adapters (HBA) for protocols such as Fibre Channel and SAS/SATA.

Unfortunately, DAS has a number of limitations which prevents it from keeping up with ever-increasing capacity needs. As the number of people working on a project increases the problems inherent with DAS become more pronounced.

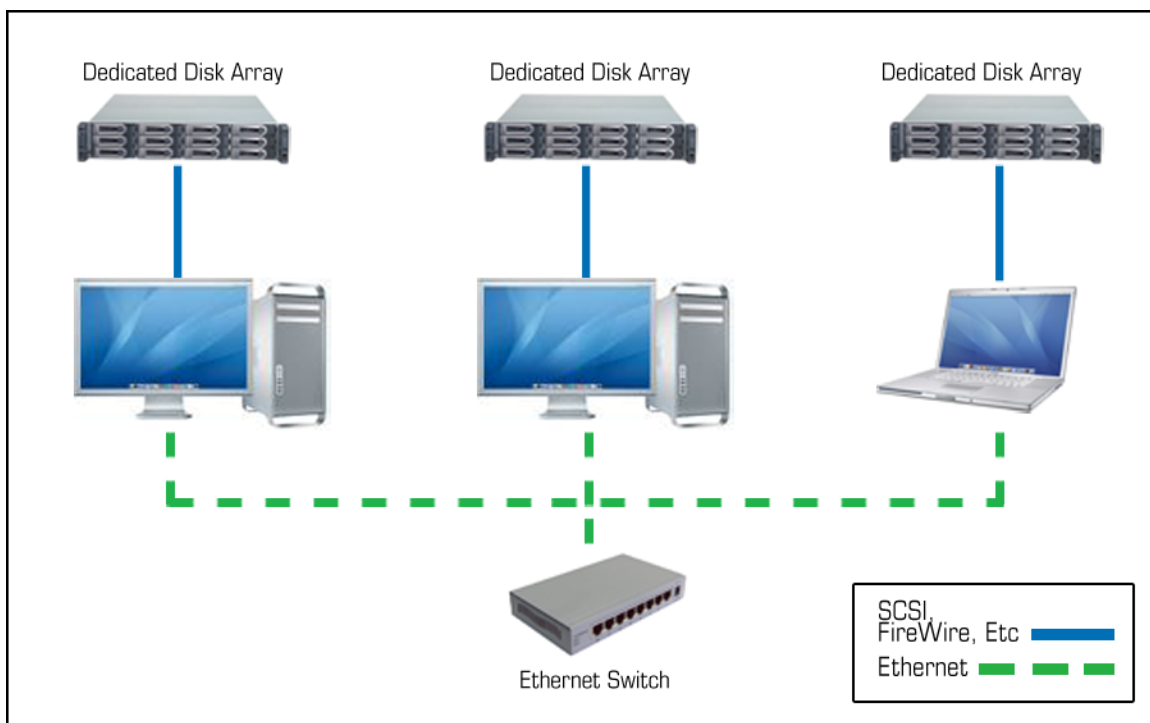


Figure 1: Direct Attached Storage (DAS)

DAS Challenges

- **Sharing Storage:** To access DAS from more than one workstation, the entire storage system needs to be directly connected to the workstation. This involves moving hardware, re-cabling, and possibly copying entire projects from one Mac to another. While this transfer is being conducted, the workstations are in-use, meaning no work can be done.
- **Backing-Up Data:** Backing-up a large number of Macs with DAS is a difficult and time consuming task for a production manager. Each piece of storage must be backed-up independently, which lead to gaps in the archive. A hardware failure may mean that work is lost forever.
- **Storage Failures:** Failures in local DAS cause the attached Mac workstation to be unusable until the storage can be repaired, resulting in unacceptable downtime.
- **Environmental Factors:** Locally attached disk arrays add heat and noise to the work environment, adding to user fatigue and making it difficult to conduct sensitive audio editing.
- **Storage Utilization:** Because of the nature of DAS, storage utilization can be exceedingly inefficient. One workstation may use only 50% of its available storage space while another workstation may be reaching its capacity limit. With DAS, the extra storage of the first user will remain unused, resulting in a significant amount of wasted storage space.

SAN Solutions

SANs can help solve the problems inherent with DAS by allowing each workstation to connect to any storage device via a network. In this environment, storage is consolidated to a single, central location, greatly improving project workflow and providing other benefits.

- **Sharing Storage:** SANs eliminate the need to transfer large project files across workstations because all workstations have access to the same storage. Networked storage allows multiple users to edit simultaneously, reducing project completion time and simplifying project management.
- **Consolidated Backup:** Archiving and retrieving content in a SAN can be managed easily because the data is no longer spread out across multiple workstations. Additionally, the backup process is easier to manage for IT administrators and is much less time-consuming when compared to networked storage for backup.
- **Higher Availability:** The networked storage can be attached to a Mac cluster, eliminating downtime when a single cable, host bus adapter or Mac encounters hardware problems. If a workstation encounters a failure, users can continue editing at a spare workstation.
- **Better Work Environment:** By using a SAN, editing suites can be created without the noise and heat associated with large locally-attached disk arrays. iSCSI eliminates the distance limitations associated with DAS, allowing storage to be located away from the editing suite.
- **Pooled storage:** Adding additional storage to a SAN is a much simpler and less time-consuming process than in a DAS environment. Once storage becomes shared on a SAN, unused storage can be easily reallocated from one workstation to another resulting in less wasted capacity. Any increases in capacity are immediately available to all the workstations on the SAN. Finally, the amount of disk space per OS X workstation is no longer limited by

the capacity of the local drive array, eliminating the need to add new workstations to support increased storage requirements.

- **Lower Total Cost of Ownership (TCO):** SANs use resources more efficiently by recouping capacity in existing resources and streamlining the production workflow. This lowers the total cost of ownership.

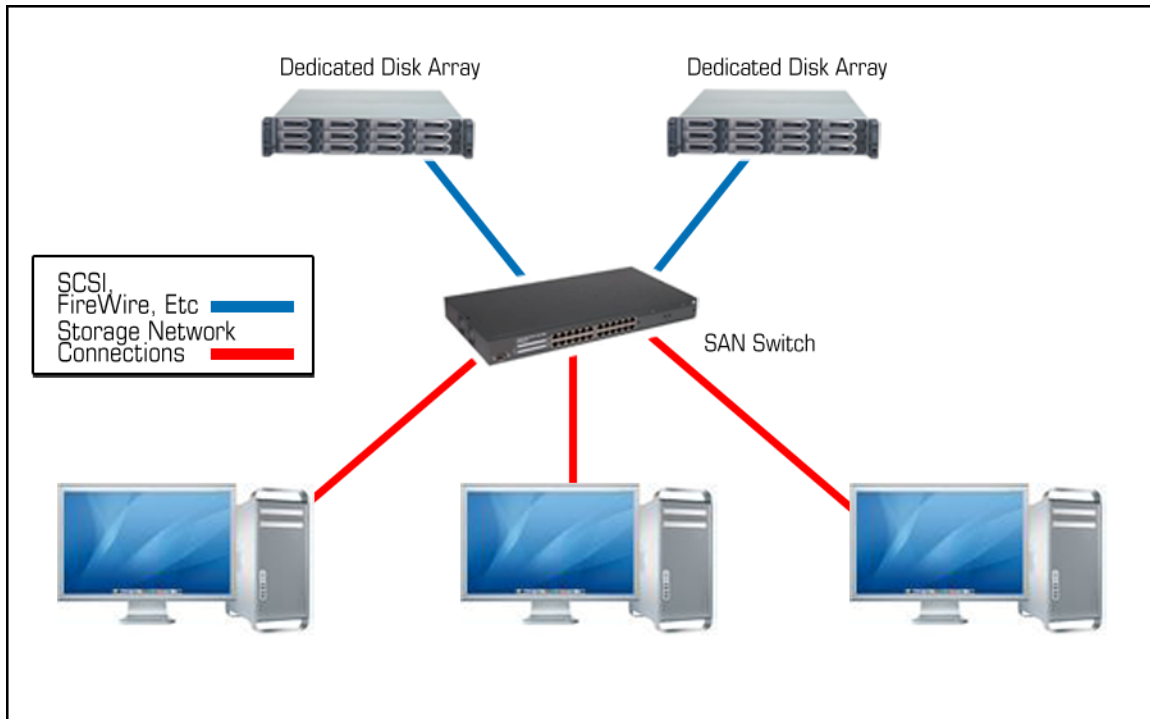


Figure 2: Storage Area Network (SAN)

Larger studios and enterprise environments typically install a SAN using the Fibre Channel protocol. Fibre Channel works extremely well in this regard and is especially well-suited for users who require the fastest available access to their SAN. Unfortunately, the high cost of deploying Fibre Channel often makes it suitable only for power users, effectively prohibiting the use of Fibre Channel SANs within smaller organizations. Many OS X users are excluded from Fibre Channel SANs entirely, since MacBooks, MacBook Airs and MacBook Pros cannot use Fibre Channel host adapters.

What about NAS?

Another type of networked storage, Network Attached Storage (NAS), uses custom hardware, often referred to as a "NAS server," which provides file-level sharing. This sharing happens via the NFS protocol or other standard Mac OS X file sharing protocols.

NAS differs from SAN in that NAS allows workstations to access data at the *file level*, not the *block level*, as is the case with SANs. This is a critical difference if a users' application needs optimized block-level data access, which is the case with many storage intensive applications. NAS storage appears as shared files to Mac OS X, while in a SAN the storage appears as if it were direct-attached.

While NAS security is typically implemented at the file-system level through traditional access-control lists, SANs offer a higher-level of security by utilizing features such as zoning. Finally, NAS typically is less reliable and available than SANs, as NAS generally features a single point of failure.

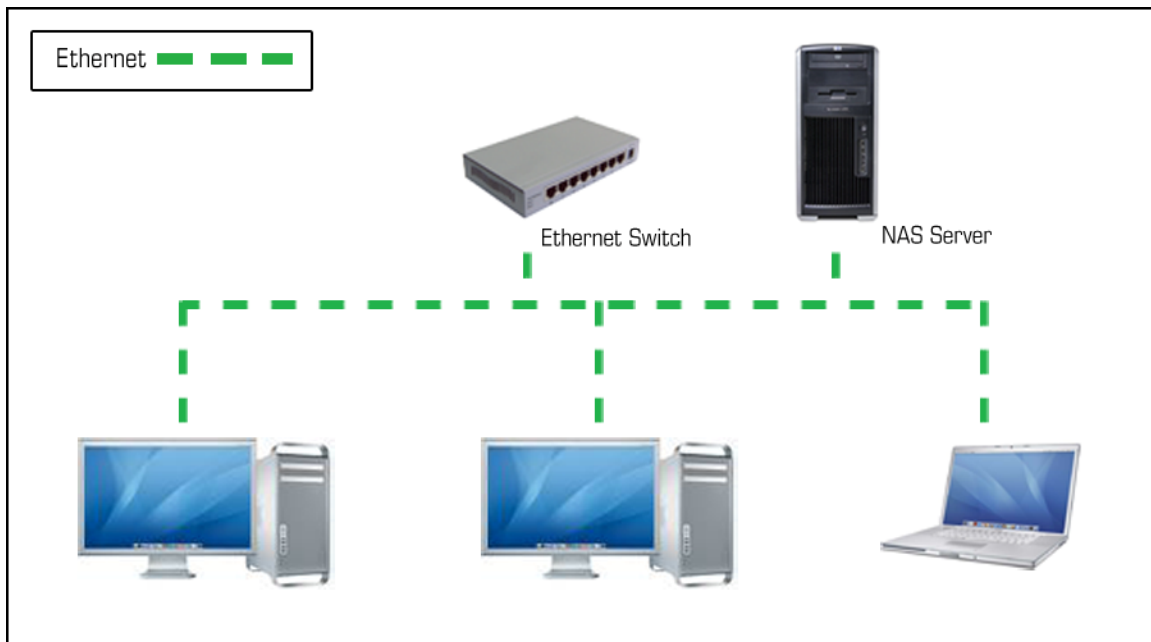


Figure 3: Network Attached Storage (NAS)

While NAS has many uses, most applications with strict storage requirements require block-level access and a high level of security and reliability, which cannot be adequately addressed with NAS storage.

USING MACS IN AN ENTERPRISE

Many OS X users work in an enterprise environment in which a Fibre Channel SAN already exists. If the Mac has an expansion slot available, a Fibre Channel HBA can be installed in the workstation to provide Fibre Channel connectivity. In some cases, though, this option is limited due to cost-constraints or lack of Fibre Channel switch ports. With OS X laptops such as the MacBook, MacBook Pro and MacBook Air, the option of adding Fibre Channel connectivity is limited because there is no physical expansion slot for the installation of a traditional Fibre Channel HBA.

THE iSCSI ANSWER

Implementing a low-cost, easy-to-manage shared storage environment which addresses the limitations of both DAS and NAS is made possible by using an iSCSI SAN. With iSCSI, Mac users can connect to an iSCSI SAN using only the built-in Gigabit Ethernet port standard on all iMacs, Mac Pros, MacBooks and MacBook Pros. Unlike other SAN protocols, iSCSI requires no dedicated cabling; iSCSI SANs can run over existing IP infrastructures.

By using existing IP equipment and expertise, and leveraging existing Gigabit Ethernet connectivity, iSCSI allows users to reap the benefits of a SAN at a lower entry-cost and in a less time-consuming fashion. Users can share files, share pooled storage and take advantage of consolidated backups and archives in an easy-to-manage iSCSI SAN environment.

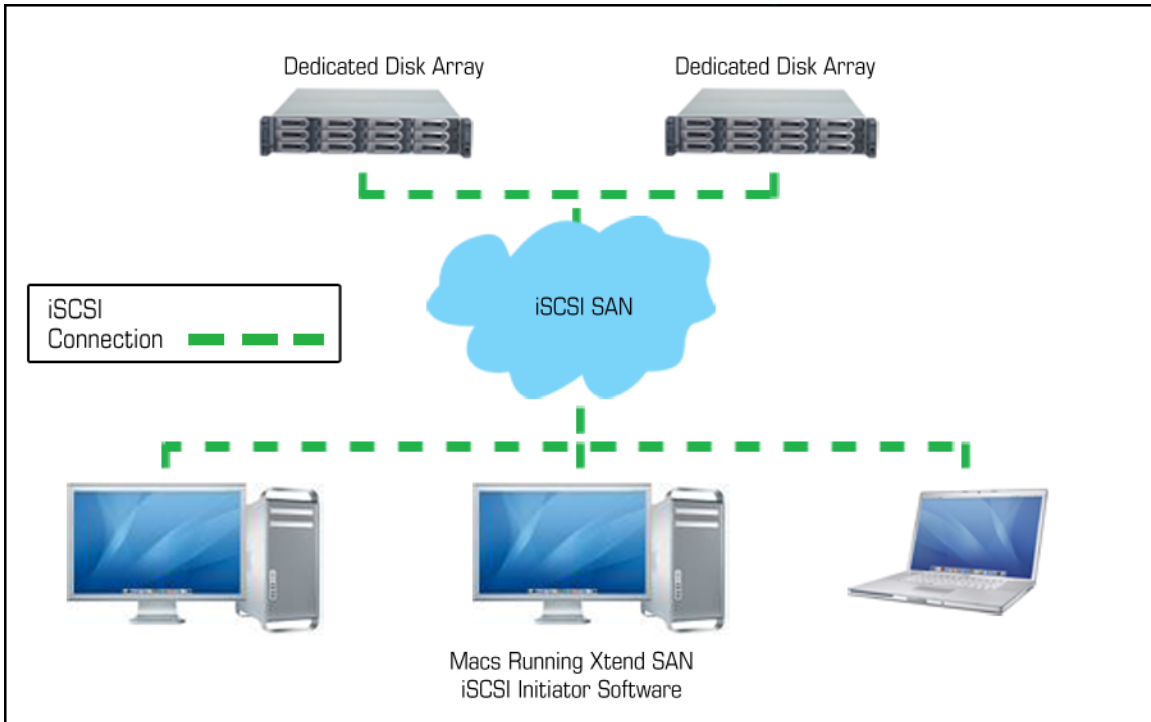


Figure 4: iSCSI SAN

iSCSI also creates an environment where users of MacBooks, MacBook Pros and iMacs can connect to existing Fibre Channel SANs through an iSCSI-to-Fibre Channel bridge. In this tiered storage environment, the Mac uses the iSCSI protocol to connect to the protocol bridge using standard Gigabit Ethernet. The bridge translates from iSCSI to Fibre Channel, allowing iSCSI users to connect to Fibre Channel SANs.

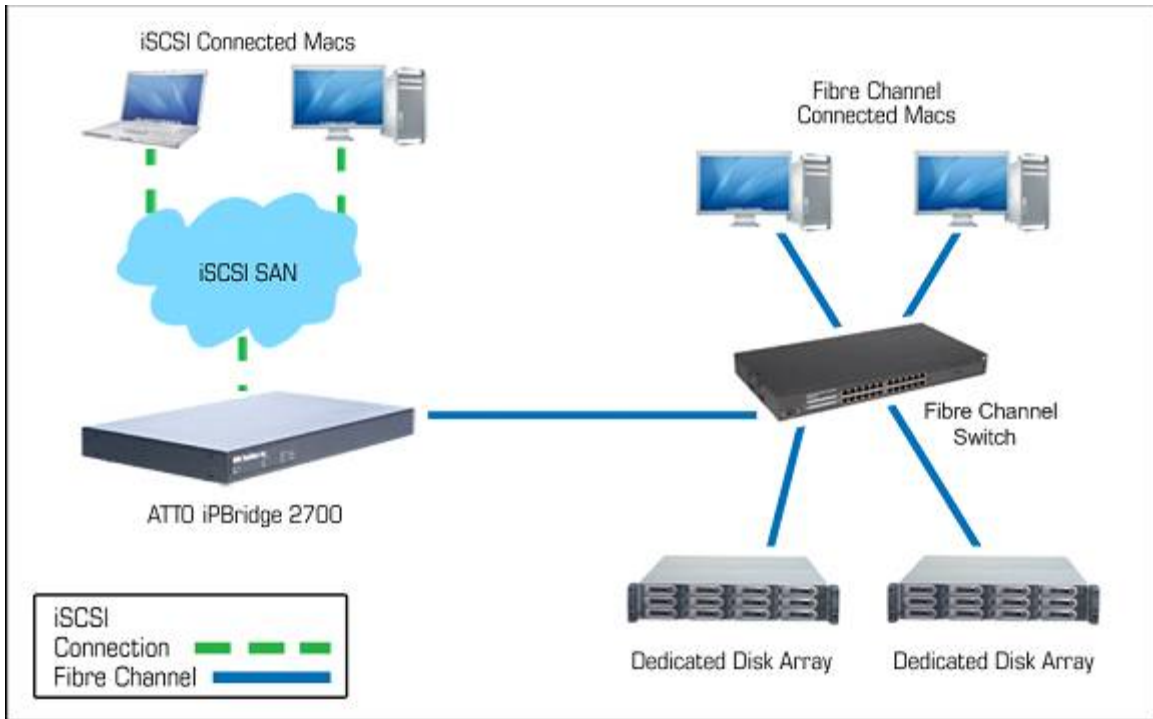


Figure 5: Tiered Storage

WHAT IS iSCSI?

iSCSI Background

The iSCSI specification (RFC 3720) was approved by the Internet Engineering Task Force (IETF) in 2003. Since its ratification, it has been reliably deployed in thousands of locations. As iSCSI has seen a steady growth rate, a large number of iSCSI infrastructure devices have become available. A larger number of iSCSI target devices are available such as iSCSI disk arrays and tape libraries. As of Q3 2007, iSCSI disk storage spending had posted a 57.2% revenue growth rate year-over-year, according to a study conducted by International Data Corporation (IDC).

In addition to the native iSCSI target devices available, iSCSI may also use existing DAS on their iSCSI SAN. Protocol bridges are available which add iSCSI connectivity via Gigabit Ethernet to legacy storage systems offering SCSI or Fibre Channel connectivity, providing iSCSI users with a myriad of options for iSCSI targets.

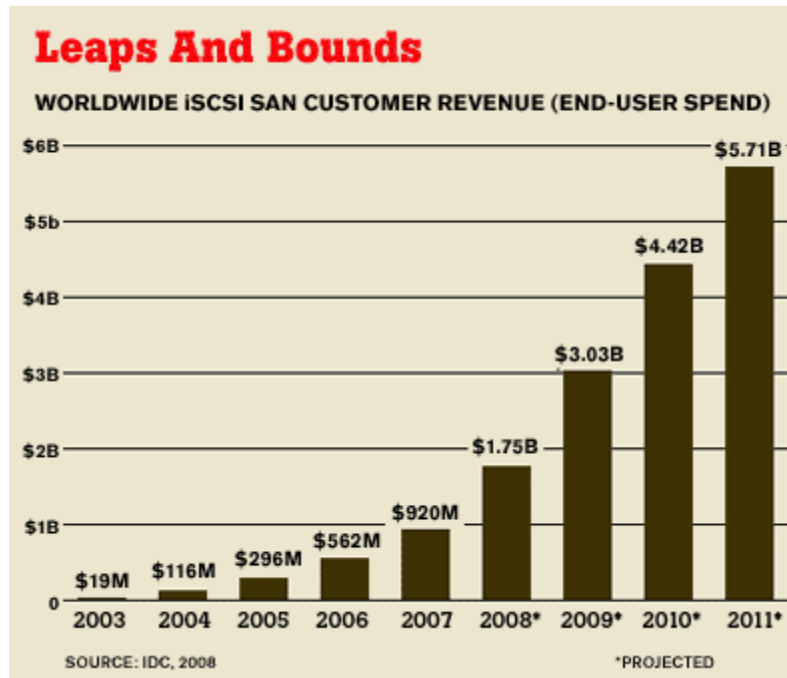


Chart 1: iSCSI Revenue

In the near future, iSCSI will benefit from the continuing market adoption of the 10-Gigabit Ethernet standard. iSCSI runs over 10-Gigabit Ethernet as it does with Gigabit Ethernet, but at a significantly higher speed. Additionally, 10-Gigabit Ethernet features a much lower latency than Gigabit Ethernet, making iSCSI over 10-Gigabit Ethernet more suitable for applications with stringent latency requirements such as 2K and 4K film editing.

How iSCSI Works

iSCSI uses the standard TCP/IP protocol over Ethernet (the same underlying protocols as your web browser) to pass block storage commands. Participants on an iSCSI SAN are classified as either a “target” or “initiator.” iSCSI uses the term “target” to refer to storage peripherals such as disk drives, tape libraries, bridges, etc. The term “initiator” refers to the workstation accessing or writing data to the iSCSI targets. The initiator runs software called an “iSCSI initiator” which translates block storage commands to iSCSI commands and allows the workstation to communicate with the iSCSI targets.

When a Mac attempts to connect to an iSCSI target, the iSCSI initiator software performs a “storage discovery,” during which the initiator software searches the network for available iSCSI targets. Once iSCSI storage is located, the initiator logs into that storage using a TCP connection over Ethernet, creating an iSCSI “session.”

This login process is exceedingly important for iSCSI as security parameters, including access controls and masked LUNs, are typically established at login. After the login process has completed, the application can send block-level commands to the storage just as it would in a direct-attached environment. This means that the application does not have to do anything differently in order to read or write to storage.

When a Mac application needs to access a disk on the iSCSI SAN, the read (or write) command is issued as a SCSI command. The iSCSI initiator software encapsulates the SCSI command and data into iSCSI

commands. Additional SCSI information, including routing, sequencing and error detection, are added to the iSCSI headers and the command is passed to the TCP layer. The iSCSI commands, in turn, are split into TCP packets and sent over the Ethernet network. The use of standard SCSI commands gives the storage system a high level of efficiency, and iSCSI's use of TCP over Ethernet guarantees that routers will forward the data to the proper destination.

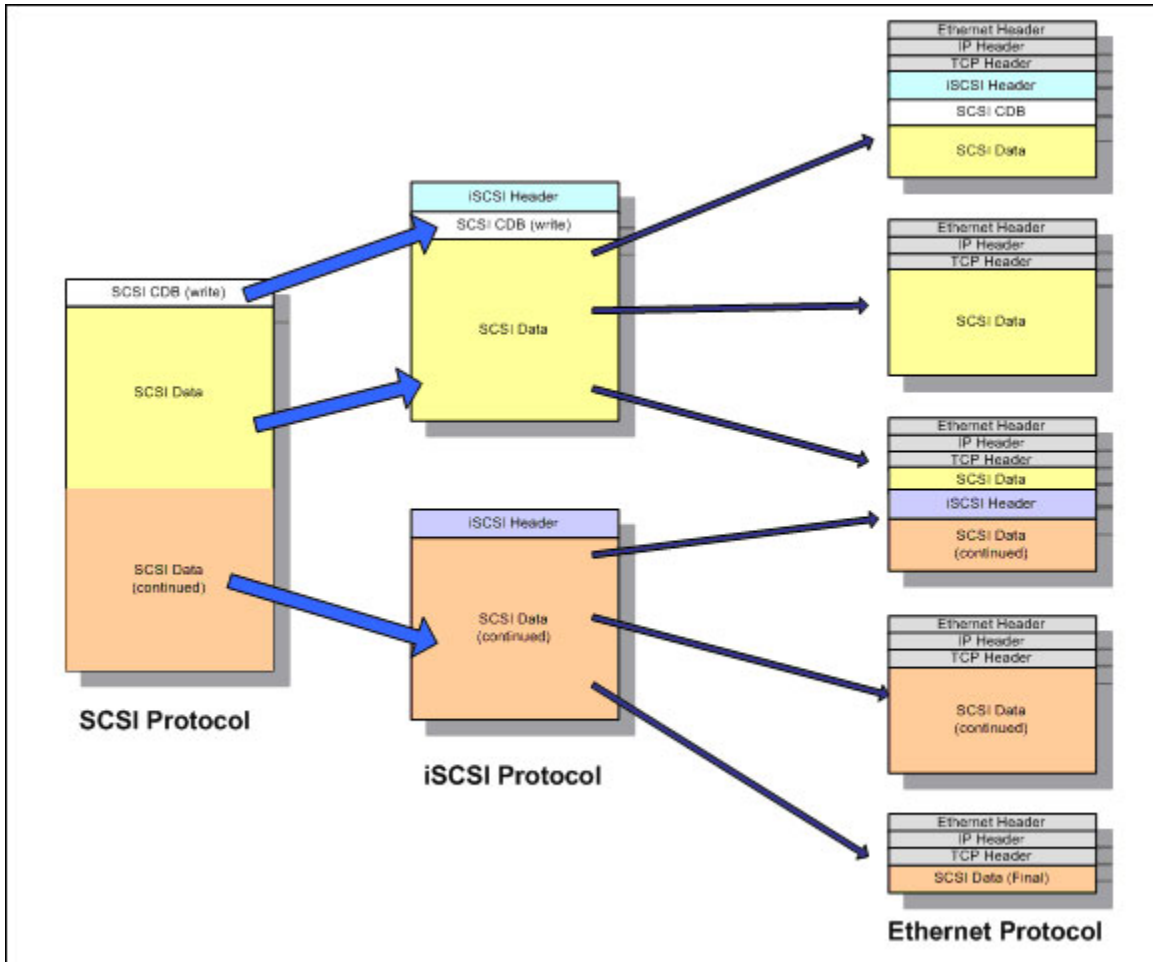


Figure 6: iSCSI Encapsulation

When the iSCSI target receives the commands, the underlying SCSI command encapsulated in the iSCSI command is decoded and the actual read or write is performed. Responses are created by the target, which then places a standard SCSI response into an iSCSI command. The response, having been created using the routing information contained in the command's iSCSI header, is then returned to the host on the original TCP connection.

In iSCSI, multiple ethernet ports can be shared by a single iSCSI session. By grouping ethernet connections together, clients can add the bandwidth together to achieve higher data rates, or the connections can be used for seamless failover: If one connection fails, the iSCSI session can automatically fail over to the other connection. Commands outstanding on the broken connection can be reassigned to the active connection in order to complete the commands without loss of data.

iSCSI Security

iSCSI mitigates the security concerns of networked storage by using familiar networking techniques such as Access Control Lists (ACLs) and Virtual LANs (VLANs). These techniques provide functionality and security similar to Fibre Channel zoning.

ACLs limit the initiators that can login to a certain target by allowing the user to define a list of Ethernet addresses which can access to the target. Only the hardware listed in the ACL will be allowed access to the storage.

VLANs enhance iSCSI security by allowing users to virtually segregate a portion of the LAN from other traffic. In a VLAN, a virtual network connection is configured using a special VLAN tag on the Ethernet frame. VLAN-enabled network switches will recognize the defined VLAN, and will not route VLAN traffic to machines using the wrong VLAN tag. This improves iSCSI security by hiding the iSCSI data from prying eyes on the general network.

Additional security features built into iSCSI revolve around iSCSI logins. iSCSI logins can pass through a “security phase,” where password-hashed challenge phrases are exchanged between the initiator and the target. This is accomplished via the Challenge Handshake Authentication Protocol (CHAP), the primary authentication mechanism of iSCSI. CHAP allows the target to verify that an unauthorized initiator is not pretending to be a valid user, and at the same time protects authorized initiators from connecting to fake targets.

Much like Fibre Channel LUN masking, iSCSI target logins can be used to grant specialized access to specific initiators. For example, imagine there are two projects occurring simultaneously. A production manager needs access to all projects, but members from Team A and Team B only need access to their own projects. An IT administrator can configure three separate targets in iSCSI: Target A exposes only Project A, Target B exposes only Project B, and the Manager Target provides access to both projects.

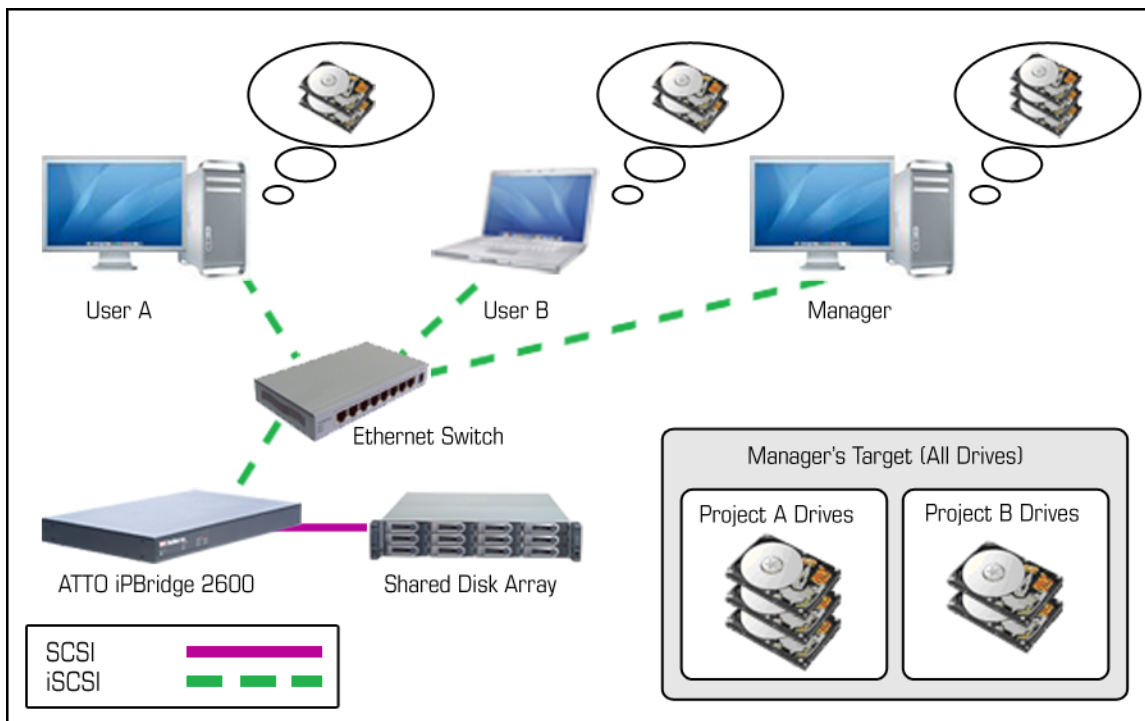


Figure 7: iSCSI Target Mapping

iSCSI Digests

Users requiring a high level of error detection can use iSCSI digests. An iSCSI digest, often referred to as a Cyclic Redundancy Check (CRC), is a unique number used to detect accidental alteration of the data being transmitted. If digests are enabled, the digest is calculated by the sender and passed along in the iSCSI command. The receiver of an iSCSI message uses the same calculation and compares the result with the digest transmitted. If the digests do not match, an error has occurred, and the command is automatically retried. Digests can be enabled on the iSCSI header, the iSCSI data, or both.

What is a TOE?

A TOE (TCP Offload Engine) is an ethernet controller with specialized hardware to accelerate TCP handling. Early adopters of iSCSI (around the year 2000) needed to use a TOE because their workstations were not fast enough to handle iSCSI traffic. With advances in processors in the years since iSCSI's invention, the use of TOE's has declined to all but a few specialty applications. In general, Mac users on an iSCSI network will not need a TOE to get gigabit speeds in an iSCSI network. In fact, Mac users will see close to full line rate transfers using iSCSI, achieving 98% of the available 118 MB/s, without a TOE.

SOME iSCSI EXAMPLES

Audio Workstations

In professional audio environments, space and cooling are at a premium. Noisy, hot disk arrays placed near a DAW exacerbate those concerns. iSCSI solves this problem by using standard Ethernet to connect disk drives to the DAW, letting users place the disk drives out of the listening space for greater sound isolation.

Many recording studios with multiple control rooms still use "sneaker-net," moving disk drives from room to room by hand, or burning the audio to a DVD-R and physically carrying the media from place to place. Without iSCSI, the investment in a SAN between control rooms is often too costly to consider. With iSCSI, the rooms can be linked via inexpensive Gigabit Ethernet, creating an improved workflow and better use of time and space.

Consider a studio with a large live area and a smaller solo booth, each with its own control room. A third control room is used for production and/or mastering (Figure8).

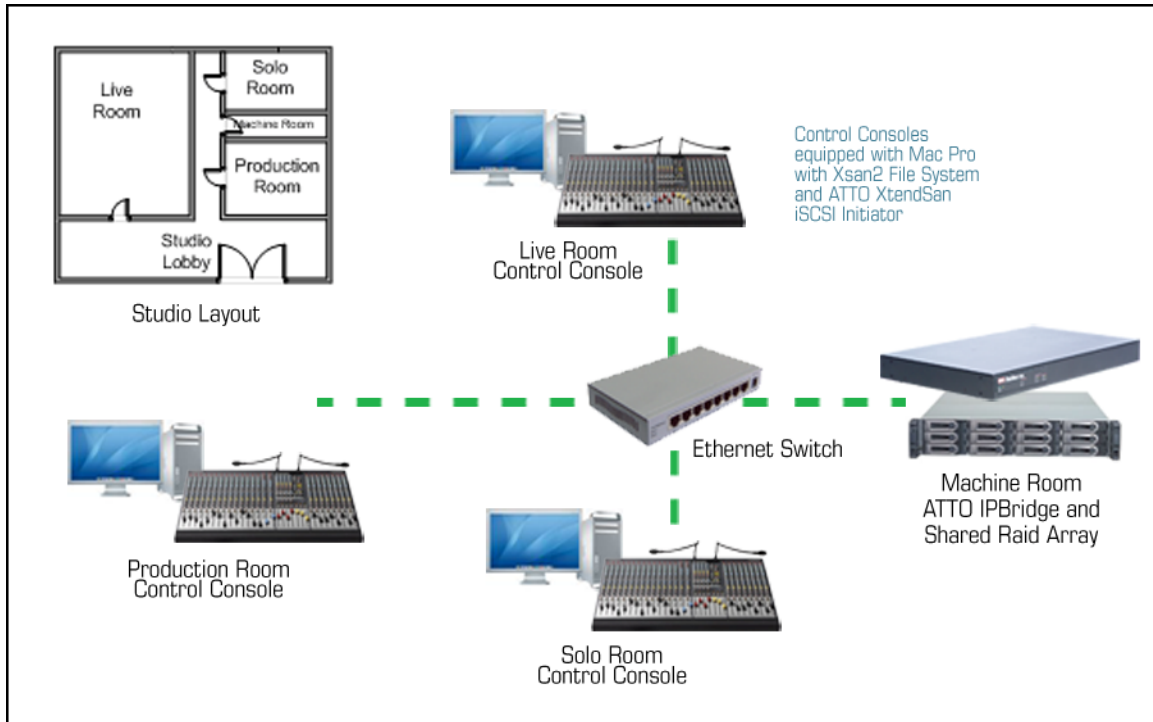


Figure8 : Audio Recording Example

In this example, the Macs in the large room, solo room, and production room are connected to the same RAID drive array using iSCSI. In addition to the appropriate audio editing suite, each Mac is running the Apple Xsan 2 file system to allow seamless SAN file sharing, along with the ATTO Xtend SAN iSCSI Initiator. By using an iSCSI SAN, all three areas can be used simultaneously without waiting to transfer data.

Tiered SAN for Video Editing

Many digital video editing environments have a Fibre Channel SAN implemented for their extremely high-bandwidth users. A tiered SAN incorporates iSCSI to inexpensively connect other workstations to the existing Fibre Channel network (Figure 9).

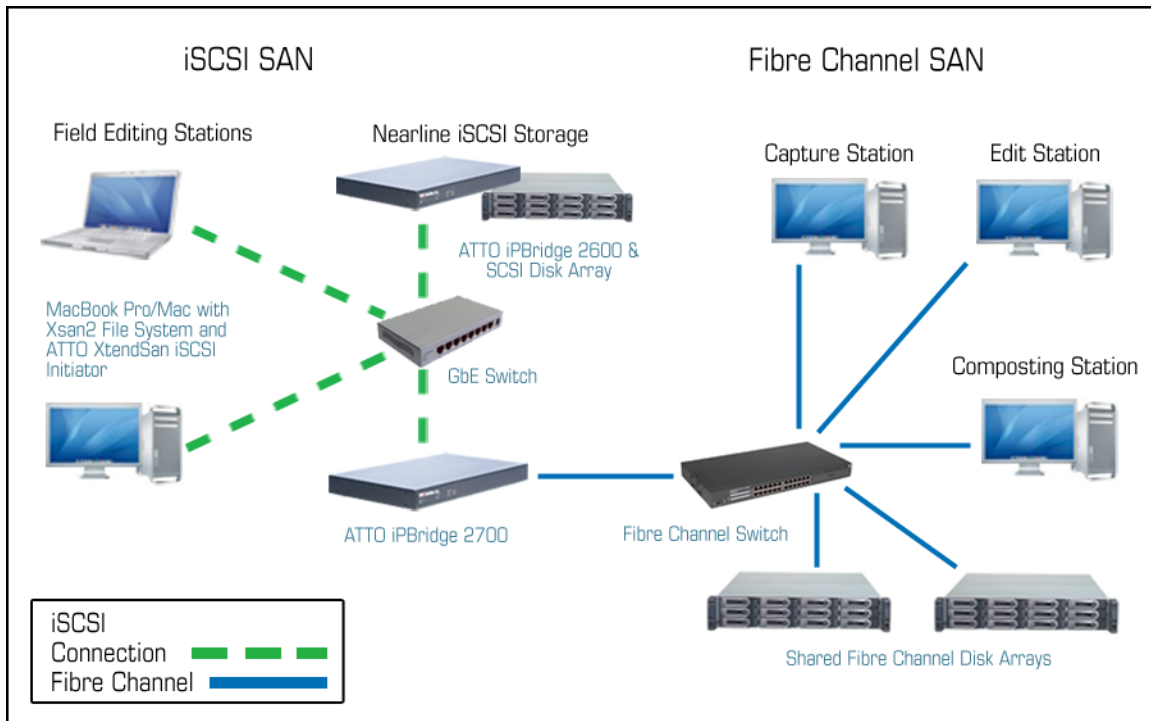


Figure 9: Tiered Storage Video Example

In this example, power users are connected to a Fibre Channel SAN for video capture, composition and video editing. Various stations have access to a pool of shared storage on the Fibre Channel SAN.

The field edit stations, again running Apple Xsan 2, use the ATTO Xtend SAN iSCSI initiator and connect through their Gigabit Ethernet port to an ATTO iPBridge™ 2700 (an iSCSI-to-Fibre Channel bridge), which enables them to access the shared Fibre Channel disk arrays. By implementing this tiered SAN setup, users can inexpensively add new field edit workstations to access that Fibre Channel SAN. iSCSI SANs also allow the field edit stations to have their own inexpensive near-line storage.

MacBook/MacBook Pro SAN Connectivity

Since iSCSI uses Gigabit Ethernet, MacBooks, MacBook Pros and even the MacBook Air (via wireless connection) are all physically equipped to support the standard. Any OS X laptop user looking to use iSCSI needs only to install Xtend SAN to connect to the iSCSI SAN via their standard Gigabit Ethernet or Wi-Fi connections.

While OS X laptops can connect to iSCSI SANs via their Gigabit Ethernet ports, using ATTO iSCSI products can allow those laptops to connect to Fibre Channel SANs. Since OS X laptops lack the expansion space to accommodate a Fibre Channel host adapter, their options are limited when connecting to a Fibre Channel SAN. Using ATTO's iPBridge 2700 in concert with Xtend SAN allows users of OS X laptops to use iSCSI to achieve that connectivity.

This connectivity can also be achieved by using Wi-Fi. Although wireless bandwidth may limit performance in these applications, iSCSI opens the window of SAN connectivity to wireless MacBook users.

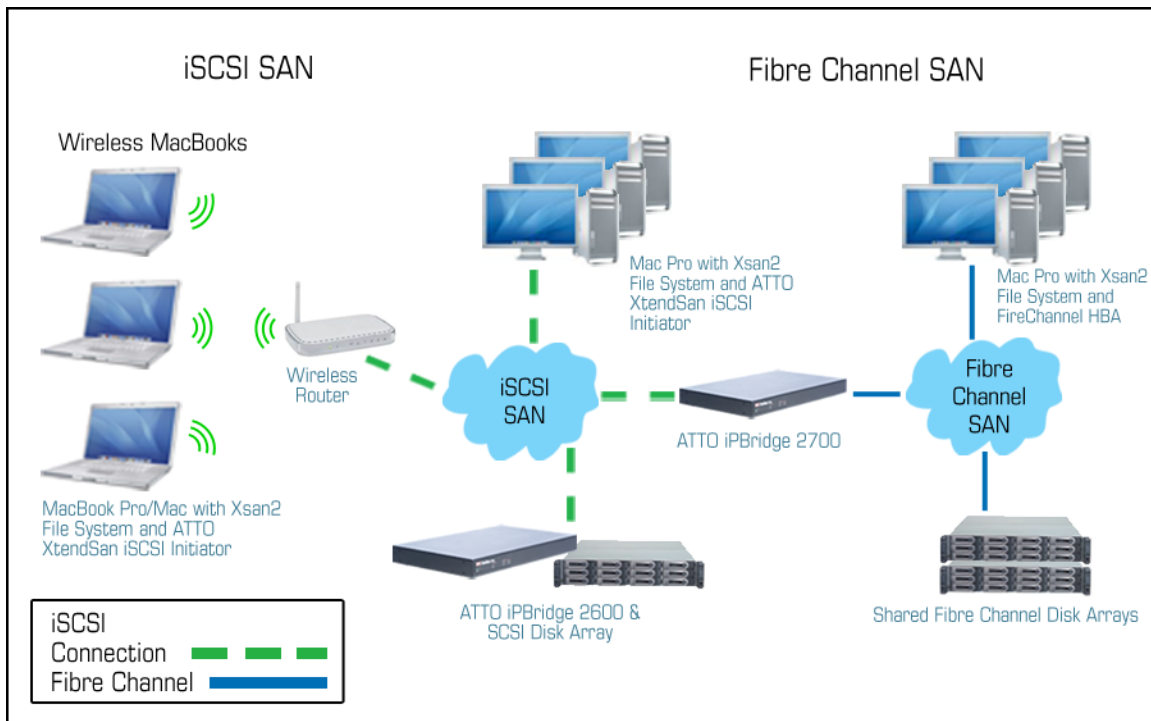


Figure 10: MacBook Wireless iSCSI

In this example, MacBook users can use Xtend SAN over wireless connections to share both the near-line iSCSI storage and Fibre Channel disk arrays.

ATTO: A LEADER IN iSCSI

ATTO Technology, Inc. is a leading provider of iSCSI solutions for Mac OS X. ATTO has been designing and manufacturing storage networking products for over 20 years, and has been working with iSCSI technology since its inception. As a result of this combined experience in storage networking in general, and iSCSI specifically, ATTO has acquired a wealth of knowledge and expertise to help customers meet their iSCSI storage networking needs.

ATTO's iSCSI solutions include the Xtend SAN iSCSI initiator for OS X and a full line of iPbridges offering iSCSI-to-SCSI, iSCSI-to-SAS, and iSCSI-to-Fibre Channel connectivity.

Xtend SAN was the first commercially available iSCSI initiator software available for OS X. Unlike other operating systems, such as Microsoft Windows, OS X does not offer a built-in iSCSI initiator. ATTO developed Xtend SAN to provide OS X users with the tools necessary to employ the benefits of a iSCSI. Xtend SAN sports advanced features including Challenge-Handshake Authentication Protocol (CHAP) support, Internet Storage Name Server (iSNS) client support, asynchronous logout and login redirect functionality.

ATTO's iSCSI-to-SCSI bridges include the iPBridge 1550E/D and the iPBridge 2600R/D. Both products add iSCSI connectivity to SCSI storage, effectively extending the life of legacy storage, and providing users with a cost-effective solution to implementing iSCSI storage without purchasing new storage.

ATTO's iPBridge 2700C/R/D is an iSCSI-to-Fibre Channel bridge. The iPBridge 2700 can be used to connect Fibre Channel storage to iSCSI SANs, but also can act as an interconnect between Fibre Channel and iSCSI SANs, enabling a tiered SAN solution as shown in an earlier example. When used with Xtend SAN, the iPBridge 2700 can enable compelling iSCSI applications including remote backups, laptop connectivity to Fibre Channel SANs, and implementing a multi-protocol SAN environment.

All ATTO iPbridges offer ExpressNAV™, ATTO's browser-based GUI for remote administration and diagnostics. ExpressNAV simplifies both the initial configuration and continuing management of the iPBridge and the storage connected to the bridge by allowing users to work within an everyday web browser. ExpressNAV also features ATTO's ExpressWizard configuration assistant which helps optimize the bridge for specific applications by asking a series of basic questions during initial setup, and configuring the bridge to optimally perform while in that setup.

ATTO iPBridge products also offer a wealth of additional features to aid the user in their iSCSI SAN setup and management. SpeedWrite™ increases write performance by efficiently managing outstanding write commands in a tape backup application, while Virtual Drive Response reduces the number of backups that fail due to busy or slow-to-respond tape drives. E-mail notification of errors and SNMP support reduce downtime and increase productivity, while the ever-increasing security concerns are addressed through support for both ACLs and VLANs.

CONCLUSION

While DAS has been prevalent for a number of years, the increasing requirements of today's storage applications are rendering it unfit for usage in certain instances. The drawbacks of DAS, which include storage that is not shared across multiple hosts, lengthy back-ups, storage failures, and inefficient storage utilization, are addressed by utilizing a SAN. Another type of networked storage, NAS, is also used in a number of applications, but SANs provide a number of benefits in storage intensive applications when compared to NAS, such as block level data access and higher reliability and availability of storage.

iSCSI is a low-cost, easy-to-manage shared storage environment which utilizes standard Gigabit Ethernet. iSCSI addresses the limitations of DAS and NAS by providing benefits such as increased storage utilization and collaborative workflows. Since iSCSI runs over Ethernet, users can leverage existing investments in IP equipment and expertise, making it an easier-to-manage and cheaper solution when compared to Fibre Channel SANs.

iSCSI is an important enabler for a number of applications from both an IT and a digital video/audio perspective; for example, iSCSI provides easy methods to attach laptops to SANs, and also to create collaborative workflow environments.

OS X users operating in storage intensive applications will find that iSCSI helps solve many limitations associated with using DAS or NAS in IT or digital video/audio environments, and in a fashion which is a perfect complement to the ease and simplicity associated with OS X.