Fibre Channel Basics

Fast, affordable, easy-to-configure storage networks.

Apple's Fibre Channel solution

Apple offers a powerful, scalable storage area network (SAN) solution that is both affordable and easy to deploy. At the core is Xsan, a 64-bit cluster file system that permits multiple Power Mac and Xserve systems to share Xserve RAID volumes over a high-speed Fibre Channel network. Designed for small and large computing environments that demand the highest level of data availability, Xsan provides the following benefits:

- **Storage consolidation.** Pool data across multiple RAID arrays for better performance and more efficient storage utilization.
- **High throughput.** Eliminate the bottlenecks of Ethernet-based networks.
- **Simultaneous read/write access.** Enable workgroups to collaborate easily and accomplish results faster than ever.
- **Easy volume scalability.** Plug in more RAID devices as storage requirements grow, expanding the shared volume easily.
- **Increased data availability.** Eliminate single points of failure through metadata controller failover and Fibre Channel multipathing.

Fibre Channel is a set of advanced data transport standards that allow large amounts of data to be moved reliably at multigigabit speeds between computers, servers, disk arrays, and other devices. Now Apple brings its trademark simple setup, ease of use, and affordability to Fibre Channel deployment—reducing its cost and complexity to the point where any growing business, agency, school, department, or other workgroup can enjoy the significant advantages it offers.

Apple chose Fibre Channel for high-end storage connectivity because of three main benefits: its high bandwidth, its proven reliability, and its guaranteed in-order delivery of data packets.

- **High speed.** In the technology's current generation, Fibre Channel provides consistent bandwidth in the range of 2 gigabits per second. This rate is expected to double in 2006, then double again within a few years to 8 gigabits per second, keeping up with the foreseeable needs of network users.
- **Proven reliability.** Fibre Channel was initially embraced by virtually all of the world's banks, Fortune 500 companies, and other institutions that required reliable, on-time data delivery at any cost. Today this same mission-critical data center reliability is available—much less expensively—to any business or organization.
- **Guaranteed in-order delivery.** Fibre Channel guarantees in-order delivery of raw block data. In-order delivery greatly boosts network efficiency and is required by some applications—video streaming and IP streaming, for example. Fibre Channel naturally streams video frames in order, eliminating reassembly bottlenecks that would otherwise noticeably degrade the video's required frames-per-second speed.

Fibre Channel Evolution

The development of the Fibre Channel family of interconnect standards serves as a model for the creation of modern data transfer technology. Unlike previous I/O architectures, Fibre Channel was standards based almost from its inception in the late 1980s. It has gone through a number of iterations, becoming continually more interoperable with other standard protocols and devices and winning American National Standards Institute (ANSI) approval in 1994.

"Fibre" vs. "fiber"
Contributing to its reputation for complexity are the confusing naming and spelling conventions that have stuck to Fibre Channel during its development. Note that the word "channel" in no way indicates a preference for channel protocols or environments. Also, although "fibre" and "fiber" are semantically equivalent ("fibre" is the international English spelling), "Fibre Channel" is the official spelling for the technology. "Fiber" has come to mean more specifically the optical (glass) media used for long-distance (up to 10 km) connections.

Fibre Channel first found use in banks, large corporations, and data centers. Installations tended to be complex, particularly when optical fiber was used as the transmission media. That has changed in recent years: Apple and several other vendors have pioneered broader Fibre Channel deployment by simplifying its installation and use while driving down prices. Today Fibre Channel is likely to be the most sensible choice for any organization or business with growing data storage needs.

Best channel and network traits

In data communications, channel environments—with direct-attached storage (DAS)—are structured and predictable. Disks and RAID arrays are connected directly to the server or host computer, and data transfer is handled almost entirely in hardware. This results in fast transmission and few errors, even when moving very large files. The downside is that configuration changes can be cumbersome, and scalability is limited to the number of devices that can be attached.

In unstructured networked storage environments, devices all intercommunicate freely, peer to peer, at very high speeds. With this kind of methodology, controlling the data transfer requires software support; but correspondingly it's very easy to add servers or disks and otherwise change configurations.

Fibre Channel combines the best of these two approaches, defining standards that support both channel and networked operations. For example, Fibre Channel supports DAS protocols, such as SCSI (Small Computer System Interface), at very high speeds. Fibre Channel can also simultaneously support a wide range of network protocols, from ATM (Asynchronous Transfer Mode) and IEEE 802 to IP (Internet Protocol) with all of its popular email, file transfer, and other services.

This remarkable Fibre Channel interoperability has made successful the ultimate storage environment: the storage area network, or SAN. In SANs, Fibre Channel has become the industry's de facto fast-switching-system standard for connecting client computers and servers to highly scalable volumes of data.

Storage Area Network

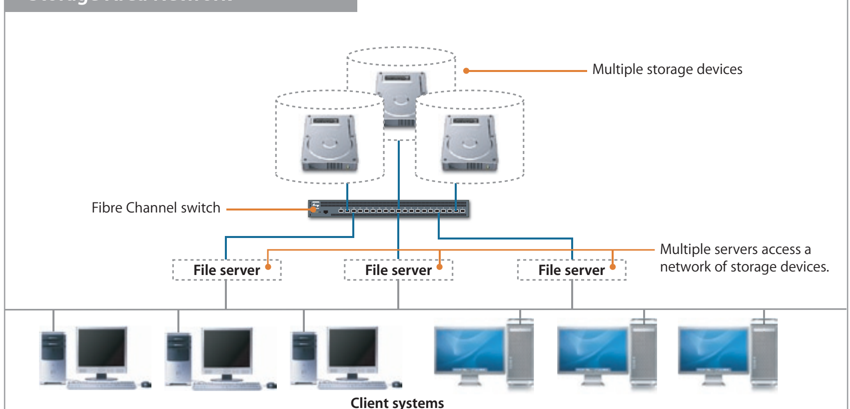

Many small businesses and organizations—from local government, real estate, and insurance agencies to school and university departments—require fast, frequent access to database files. Such workgroups would benefit greatly from the speed and reliability of a storage area network with Fibre Channel switching. With Apple Fibre Channel solutions, affordability and ease of setup and use are ensured.

Built to evolve gracefully

From early on, Fibre Channel technology was designed for efficient and economical upgrading. Users have always been able to migrate existing and scheduled installations, preserving their technology investments.

Fibre Channel network solutions operating at 1- and 2-gigabit data transfer rates have been in widespread use in SANs for some time, with 4-gigabit rates now becoming common. Already 8-gigabit Fibre Channel—the newest industry-approved standard—is being used in high-end deployments.

Fibre Channel and Protocols

What are protocols?
Protocols are sets of rules governing data during transmission between devices. The rules establish the timing, control, format, and other parameters.

Ethernet, SCSI, and other older input/output (I/O) protocols were typically constructed with specific rules, or sets of commands, governing server-to-server and server-to-storage system communications. Such I/O protocols dictate how the physical interface should look, giving rise to a profusion of cables and connectors.

Fibre Channel designers took a different approach. Segregating the functions of data transport from the functions that impart the meaning of the data, they fashioned Fibre Channel not just as a protocol, but as a carrier of protocols—including SCSI-3 command sets and IP command sets. One benefit is ready compatibility. No established command sets need to be re-created from scratch to work on a Fibre Channel network. Another benefit is that multiple protocols can be in use simultaneously, boosting network performance.

Comparing protocols

In a storage environment, there are performance, cost, and future-growth implications to consider when comparing data transfer protocols. Direct-attached and networked storage each have their own best-suited applications and specific limitations, but for transferring block data, Fibre Channel is the leading protocol choice for a number of reasons.

Block data versus file data
Block data is raw data that has been structured into groups called blocks, each a set number of bytes in size. Block data can be transferred much more efficiently than file data, because network devices can be designed to find the start of a structured block of data much more quickly than they can typically find an arbitrary point on storage media.

- **Fibre Channel and Ethernet.** The low cost and relatively easy setup of Ethernet make it one of the most common network protocols in use. Ethernet often provides the local area network connection between clients and servers, including SAN servers. In this role it works well. As a protocol, however, Ethernet is not designed for transferring block data in a networked storage environment.

  The problem lies in the way that Ethernet handles data collisions—the errors that result when more than one computer on a network tries to transmit data simultaneously. As network usage increases, the number of collisions in an Ethernet network can grow dramatically, eventually consuming all available bandwidth. In contrast, the asynchronous protocol design of Fibre Channel ensures that even when network loading is very heavy, collisions are handled efficiently and maximum throughput is maintained.

- **Fibre Channel and iSCSI.** iSCSI is one of the most talked-about new storage delivery protocols. There have been very few large-scale iSCSI deployments, so it has yet to be "data center hardened." iSCSI proponents claim that it is less expensive than Fibre Channel, but with Apple and other vendors driving down Fibre Channel prices, any such advantage is eroding. Like Fibre Channel, iSCSI is asynchronous in design.

  iSCSI is built on an underlying TCP/IP protocol, which is broadly supported and well understood by the IT community. However, TCP/IP has saddled iSCSI with a software-centric orientation that hasn't always worked to its advantage. Because the iSCSI command set has been implemented only in server-based software, it incurs significant processing overhead. In contrast, the Fibre Channel command set has been encoded in silicon since early in the protocol's development, making it much more efficient than iSCSI at any given level of processing power—and more suitable for the data transfer loading of a networked storage environment.

Network elements
Nodes are individual Fibre Channel devices, including computers, servers, and disk drives. Each node has at least one access point, or port. Ports are physically connected to one another by links, which are pairs (one receiving, one sending) of copper or fiber-optic cables.

- **Fibre Channel and Infiniband.** Another new storage delivery protocol is Infiniband, which uses its own, channel-based approach to distribute I/O processing. Infiniband is a transport network that can support iSCSI and Fibre Channel protocols. Like Fibre Channel and iSCSI, it is both asynchronous and switched fabric. Participating nodes can be concurrently interconnected rather than having to share media, as does Ethernet, or to share bandwidth, as does token ring (another network protocol).

  Infiniband is still a very new technology. Its visibility is limited to research computing, and it is likely to be years before it gains broad use in data centers.

Fibre Channel and Networks

Though it has many features of a network, Fibre Channel is less a network than a high-speed switching system that interconnects relatively local devices. With its high bandwidth and ability to support multiple protocols simultaneously, Fibre Channel enables near-instant access to massive amounts of data in SANs and other modern computing environments.

Network limits and characteristics
Collision-based Ethernet networks are ubiquitous, largely because they allow multiple individual clients to share retrieved data in a very simple and economical way. Such networks are most successful when supporting front-end functions. However, they're too inefficient to be used in block-level storage environments, such as those found in data centers.

For throughput, scalability, and attainable network lengths, Fibre Channel is far superior to Ethernet.

- **Data throughput.** With the currently available 2Gb-rated Fibre Channel in the network, data transfer rates are very close to 200MB/s, as expected. In a Gigabit Ethernet network, however, collision management claims so much bandwidth that even 1Gb rates are difficult to achieve consistently.
- **Scalability.** Whether device connections consist of a single point-to-point link or involve hundreds of integrated, enterprisewide servers, Fibre Channel networks perform with equal reliability, high rates, and flexible configuration, achieving scalable densities up to thousands of ports. Although IP-based storage networks theoretically can scale to hundreds of ports, there is no widespread use to demonstrate this capability.
- **Network lengths.** With Fibre Channel, the network's physical media—the switches and cables that carry the data—can be either copper or optical fiber. Performance is the same, though copper is limited in length to less than 3 meters. In comparison, standard Ethernet networks are copper based and use twisted-pair cabling. Without the benefit of repeaters, long-haul copper Ethernet networks are limited to 100–200 meters, depending on the Ethernet protocol version. Currently, the maximum theoretical distance for long-haul Fibre Channel networks using fiber-optic links is 10 kilometers. Apple Fibre Channel PCI cards, when equipped with long-haul transceivers, permit use of fiber-optic cables up to 500 meters in length.

Fibre Channel Topologies

What are topologies?
The collection of components that connect ports is called a topology. Topologies are also shorthand descriptions of the physical layouts, or shapes, of networks. In effect, they define different aspects of device connection or configuration—including the kinds of devices that can be configured, the number of devices, and the way they can be configured.

Apple Copper Fibre Channel Cables
With fiber-optic links, Fibre Channel offers long-haul network capability—invaluable, for example, for college or corporate campuswide networking. But optical media is expensive, complicated to install, and unnecessary for many purposes. Fibre Channel links are often used entirely within racks, over short distances where optical fiber provides no advantage over more durable, less expensive, easier-to-install copper cables. High-quality Apple Copper Fibre Channel Cables, 2.9 meters long, simplify the plugging together of Fibre Channel switches, cards, and other components.

The Fibre Channel standard defines three distinct topologies, or network shapes. These topologies satisfy different application and installation requirements, exhibit different performance characteristics, and are subject to different scalability limits.

- **Point-to-point.** The most basic Fibre Channel topology, point-to-point refers to two devices that are directly connected by a Fibre Channel cable. Addressing is simple and device availability is complete. Point-to-point Fibre Channel is not a common topology today, but it has been used to connect RAID and other storage subsystems to servers in server-centric computing environments.
- **Arbitrated loop.** This topology makes it possible for individual transmission paths, or loops, to be time-shared by up to 126 devices. Each time a loop is available for use, the devices arbitrate to determine which device gets to send data or commands next. Providing economical interconnection, Fibre Channel arbitrated loop is used to connect disk drives to RAID controllers or host bus adapters.
- **Switched fabric.** Modern SANs depend on a powerful Fibre Channel topology called switched fabric. A fabric essentially consists of one or more high-speed interconnection devices, called switches, controlling a large number of port-to-port transfers between nodes. (These transfers of data and commands are also known as frames.)

  Fabrics function like telephone systems, which carry a number of calls simultaneously over the same wires. Within a fabric, multiple interconnections happen concurrently—and all frames are routed to their proper destinations. Systemwide bandwidth can be as much as an order of magnitude greater than the speed of any single Fibre Channel link. With such high transfer rates, many users in a workflow can work with the same data at the same time, facilitating collaboration and increasing productivity.

Management Software

Fibre Channel SANs make use of software tools and techniques to manage the switched fabric. SAN file systems such as Xsan, for example, provide security, scalability, and access control. SANs may also leverage techniques for fabric management. These include zoning, LUN masking and mapping, and multipathing.

- **Zoning.** In switched fabric Fibre Channel topologies, zoning is usually performed within the switch. It involves partitioning a fabric into smaller subsets for the purposes of restricting interference and increasing security—for example, when establishing an IP firewall. Some types of zoning restrict the name services of a device, preventing servers looking at that fabric from seeing it if they are not permitted to do so. Other more robust types of zoning restrict the actual communication across a fabric.
- **LUN masking and mapping.** Logical unit numbers, or LUNs, are the addresses of individual storage devices such as disks, RAID sets, or even slices of RAID sets. Aggregating LUNs into a single volume, a SAN file system like Xsan provides the access control necessary for the same data to be used simultaneously by multiple hosts. Lacking a file system, a less sophisticated SAN must "map" each LUN to a single host for its exclusive use and "mask" it from any others. Such LUN masking and mapping techniques help simple SANs avoid the data loss that would result if more than one host wrote to an array.
- **Multipathing.** This fault tolerance technique describes how, in SANs with file systems like Xsan, multiport connections make possible more than one physical path to mass storage devices. This redundancy protects against a potential single point of failure at the cabling layer. Another benefit is that multipath software layers can leverage the redundancy to enhance dynamic load balancing, traffic shaping, and other features.

Deploying Fibre Channel

As with any network implementation, successful deployment requires making a careful determination of needs, understanding technical requirements and trade-offs, and choosing the optimal configurations and components for your situation.

To begin with, decide which type of Fibre Channel network will best suit your needs by answering a few questions. For example:

- Is it a new network or an addition to an existing one?
- What will be the network's total physical length?
- How many devices will be networked?

Detailed answers to questions like these will help determine the optimal physical layout and specifications of your network, including the type of cabling and the type and number of components and connectors.

Apple Fibre Channel host bus adapters
Apple Fibre Channel host bus adapters are multiport cards that provide every Power Mac or Xserve system with two connections to the SAN. Fibre Channel multipathing takes advantage of this dual connection in two ways. First, in the event of a cable failure, Xsan can use the remaining path for storage access. Second, multipathing can increase performance by dynamically load balancing the traffic across both connections.

Copper or fiber optic?

Choosing the type of overall network interconnection, copper or fiber optic, is crucial. The decision usually depends on the distances between the Fibre Channel devices being connected.

Copper cabling performs as well as fiber optic when the distance between devices is 2.9 meters or less. This is typical in point-to-point and other topologies when devices are mounted in the same rack or are located near each other in the same room. Easy to work with, copper cables are durable and can withstand being stepped on or pulled. They're also inexpensive to purchase, install, and maintain.

Fiber-optic-based Fibre Channel networks cost more and are more complicated to deploy than copper-based ones. But they are necessary to achieve the high data rates of Fibre Channel when servers and other network devices are located more than 2.9 meters apart—for example, in different buildings or on different floors of a building.

It's not just the cabling that contributes to the cost and complexity of fiber-optic installations. Cable runs require transceivers at each end, and the brand and type must be specified correctly. Some are short-haul, some are long-haul (over 150 meters), and none are cheap.

Fiber-optic cable

Fiber-optic cables are immune to the electrical resistance and electromagnetic interference (EMI) problems that affect signals carried over a copper network. However, dirt, dust, and moisture can significantly diminish signal strength in fiber-optic cabling, as will material defects in the glass fiber or damage to it caused by inexpert handling and installation.

Fiber-optic cables come in two types: single-mode fiber, at 9 microns in diameter; and multimode fiber, at diameters of 50 or 62.5 microns. Your Fibre Channel integrator or installer can tell you which type is best for your network. After installation, make sure your installer tests and verifies that the cabling is not broken and that it delivers maximum performance. Variation in cable performance can have a significant impact on the performance of the entire network.

Maximizing throughput

Optical fiber transmission media is naturally fragile. Curl a loop of fiber-optic cable and the glass core will likely fracture microscopically. Such fractures—as well as any scratches, nicks, or imperfections in the glass itself—disperse some of the light passing through the core, thereby reducing the strength of the signal carrying the data. This loss in signal strength, called attenuation, has to be minimized throughout the network.

Fiber-optic cable installation requires great care, special tools, and expertise. To ensure maximum network performance, engage a certified fiber-optic installer to do the work. Make sure the installer has a cable tester that will check for high levels of throughput in each section of the network.

Critical connectors

To select the correct fiber-optic connectors, you need to know what the cables are connecting to. For example, nearly all Fibre Channel switches require SFP (small form factor pluggable) interconnect cables. In 2Gb and 4Gb networks, the SFPs use a connector technology designated "LC"; in older 1Gb networks, they use a connector technology designated "SC."

Connectors

Duplex SC

SC Optical GBIC Transceiver

SC

LC

LC Optical SFP Transceiver

LC Optical SFF
Board-Mount Transceiver

Installation methods vary, but all share certain cautions. For example, when stripping the cable's outer jacket and the fiber coating, the installer must use special techniques and tools to avoid damaging the glass fiber. To ensure long-term, consistent performance and reliability, it's also important to avoid touching the end face of the fiber after the connector is attached.

All fiber-optic connectors have to be carefully tested after installation. Whenever possible, it's best to buy certified, preterminated and tested cables.

The Apple Fibre Channel advantage

Apple emphasizes ease of use and lower installation cost in providing customers with Fibre Channel networking solutions. The result is copper interconnect because it's easier to install, cables that come preterminated, and SFPs that work best with Apple products.

Apple Fibre Channel solutions simplify and reduce the costs of installation and management. Most implementations using Apple and Apple-specified components require no special expertise to deploy and are surprisingly affordable. Even the more robust deployments of Apple Fibre Channel technology cost much less than competing solutions—and prices, as well as installation complexity, are steadily decreasing.

Deployment Scenarios

Use the following scenarios to visualize the type of Apple Fibre Channel implementation that best fits your data storage requirements.

Single- and dual-host access to RAID storage
These basic deployments of Fibre Channel technology take advantage of the two ports on Apple's Fibre Channel PCI interface. When one Xserve host connects to an Apple Xserve RAID system, both ports on the Fibre Channel PCI card can be used for the full available bandwidth and full storage capacity. When two hosts connect, each independent RAID controller serves seven drives to a separate Xserve or Power Mac. The independent controllers protect and isolate data between servers. Because the Apple Fibre Channel PCI card comes with two copper-based cables, you can immediately connect Xserve RAID to two servers.

Apple servers and storage products
Apple's best-in-class servers and storage products provide the ideal infrastructure for deploying a SAN. Qualified for use in heterogeneous networks, Xserve RAID offers RAID protection, Fibre Channel connectivity, and fast, reliable Ultra ATA hard drives—providing up to 7TB of storage in a 3U system. Xsan is a true SAN file system, allowing multiple servers and clients to share a common storage pool. Add Xserve servers to provide infrastructure services and file sharing to Mac, Linux, and Windows clients—or for storage management as an Xsan metadata controller. Xserve systems are also ideal NAS replacements, providing single-point management and easy scalability using Xsan.

Single Xserve Host

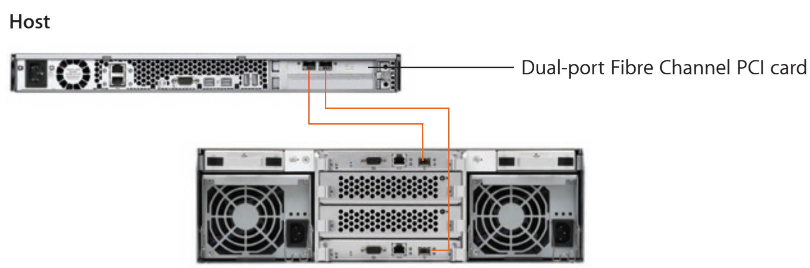

A single Xserve host connects both ports on its Apple Fibre Channel card to an Xserve RAID storage system.

Two Xserve Hosts

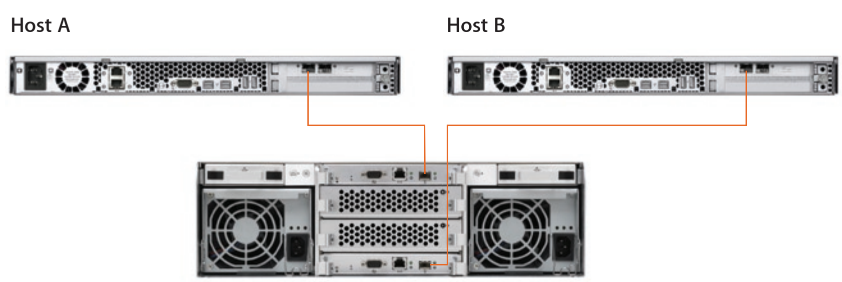

Two Xserve hosts connect to an Xserve RAID storage system.

Fibre Channel infrastructure integration
Apple has qualified several of the most popular Fibre Channel switches for use with Xsan, Xserve RAID, and Apple Fibre Channel cards, including switches from QLogic, Cisco, Emulex, McDATA, and Brocade.

RAID storage for multiple hosts via Fibre Channel switch

Adding a Fibre Channel switch results in the most flexibility for medium-size network deployments using RAID storage. Simply plug several servers into the switch, using either copper or fiber-optic cabling. Each server can access part of the overall storage on the Xserve RAID. The switch provides full bandwidth to each server.

Four Xserve Hosts

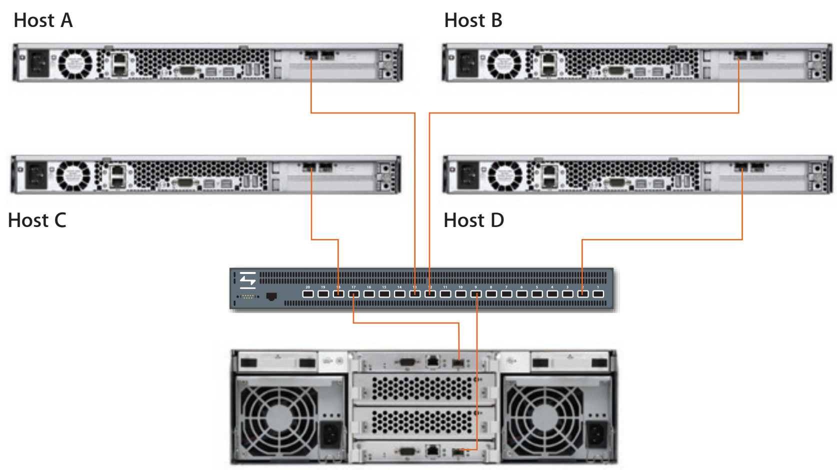

Four Xserve hosts access Xserve RAID storage via a Fibre Channel switch.

Workgroup SAN

Consider this type of deployment when you need to manage many large files and rapidly growing amounts of data—video and audio editing and storage, for example. The RAID storage pool allows quick, reliable scaling of storage and backup capabilities. Fibre Channel uniquely provides in-order delivery of data, necessary for efficient access to media files. In addition to the Fibre Channel switch, dedicated metadata controllers help mediate traffic for maximum data transfer rates.

Workgroup SAN

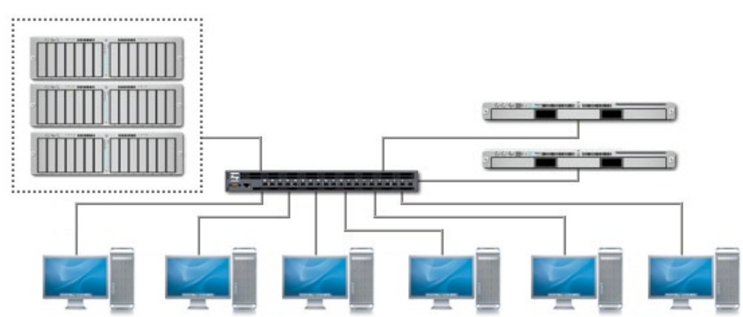

Workstations access multiple Xserve RAID systems in a storage pool via a Fibre Channel switch.

Enterprise-class SAN

This enterprise-class SAN deployment is especially suited to providing data and services to multiple clients via HTTP, FTP, or file sharing. Because Xsan is completely interoperable with the StorNext file system, direct Fibre Channel block-level access to storage pool data can be provided to Windows, Linux, and UNIX platforms. IT administrators can consolidate data, manage it centrally, and allocate access to any application or user anywhere on the network. Attached Xserve hosts, functioning as "NAS heads," provide scalable, high-performance file services for Mac, Windows, and Linux clients on the network. A backup metadata controller assures high availability.

Enterprise-class SAN

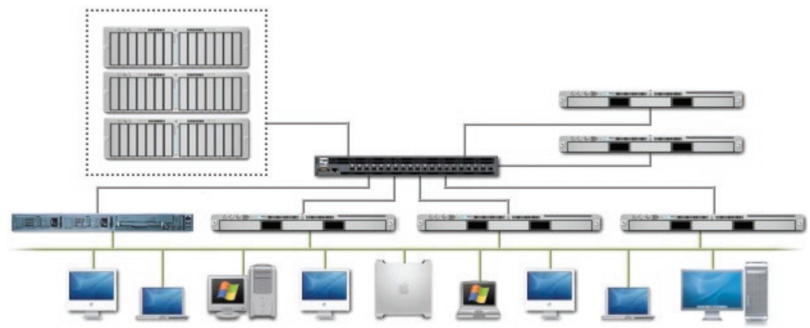

An Apple Xsan solution serves many clients, with multiple Xserve hosts networked to a RAID storage pool via a Fibre Channel switch.

Server and Storage Solutions from Apple

With tight integration between hardware and software, Apple's server and storage products allow you to extend your computing infrastructure, while lowering your management and maintenance costs. Combining the latest open source technologies with Mac ease of use, Mac OS X Server unleashes the power of Xserve, Apple's rack-optimized server hardware—resulting in an unparalleled server solution for businesses, schools, and research centers. For massive storage capacity, organizations can add Xserve RAID storage systems—and share terabytes of data over a high-performance Fibre Channel network with Xsan, Apple's SAN file system.

To learn more about Xserve, Xserve RAID, Mac OS X Server, and Xsan, visit www.apple.com/server. Download user's guides, technology briefs, and customer profiles; and find out about training and support opportunities designed for IT managers and technical professionals.