


# Storage Topologies



**With information growing at exponential rates, understanding how to cost-effectively deploy storage devices is becoming increasingly important.**

Robert  
Griswold  
Crossroads Systems

**O**ne meaning of *topology* is configuration—the “relative arrangement of parts or elements” or “something (as a figure, contour, pattern, or apparatus) that results from a particular arrangement of parts or components” (<http://www.m-w.com/cgi-bin/dictionary?book=Dictionary&va=configuration>). As applied to computer storage systems, a topology incorporates specific storage protocols, transport mechanisms, and physical connections usually associated with advanced storage practices including

- internal and external hard drives and their connections,
- laser- and optical-based rotating and fixed mediums,
- magnetic tapes,
- switched networking fabrics, and
- associated software techniques used to deploy these devices.

Not all computer storage designs or deployments are topologies. Declaring a networked storage environment to be a storage area network topology is misleading because a SAN—defined as “a network whose primary purpose is the transfer of data between computer systems and storage elements and among storage elements” ([http://www.snia.org/education/dictionary/s/#storage\\_area\\_network](http://www.snia.org/education/dictionary/s/#storage_area_network))—is not limited to a specific topology or its implementation. What makes the storage useful is what defines a particular topology or combination of topologies.

Whether they are internal to a desktop PC or deployed in a large enterprise IT department, most topologies have a common theme. Focusing on how storage topologies are deployed in end-user focused platforms provides a basis for understanding how

and why certain devices have migrated toward particular connections. In addition, clarifying the differences between directly attached storage and networked storage helps to explain why deploying large amounts of storage may not require an either-or decision.

## PC STORAGE

Until a few years ago, most hard drive manufacturers made drives using the same basic techniques for the base product but packaging different chips, firmware, and boards around the physical disks for different topologies and connections. Although enterprise storage topologies generally do not apply to personal computers, the need to easily incorporate the latest technology into PC systems has almost single-handedly driven development of hard drive and storage standards, resulting in new enterprise storage topologies.

## Parallel topologies

As the “Small Computer System Interface” sidebar indicates, this popular architecture supports most modern storage formats, regardless of the connection or transport topology used. Because systems can easily package and transport SCSI commands, delivered in command descriptor blocks (CDBs), across many mediums, they can link multiple devices to various types of buses.

SCSI’s versatility facilitated incorporation of new technologies such as CD-ROM drives, which stored data in standard 74-minute formats. In cooperation with the Small Form Factor Committee, Philips and Sony defined command sets based on SCSI definitions to attach CD-ROM drives to SCSI host bus adapters. However, the drives underutilized SCSI’s costly, high-end features, including command queuing and advanced error correction. This in turn led

## Small Computer System Interface

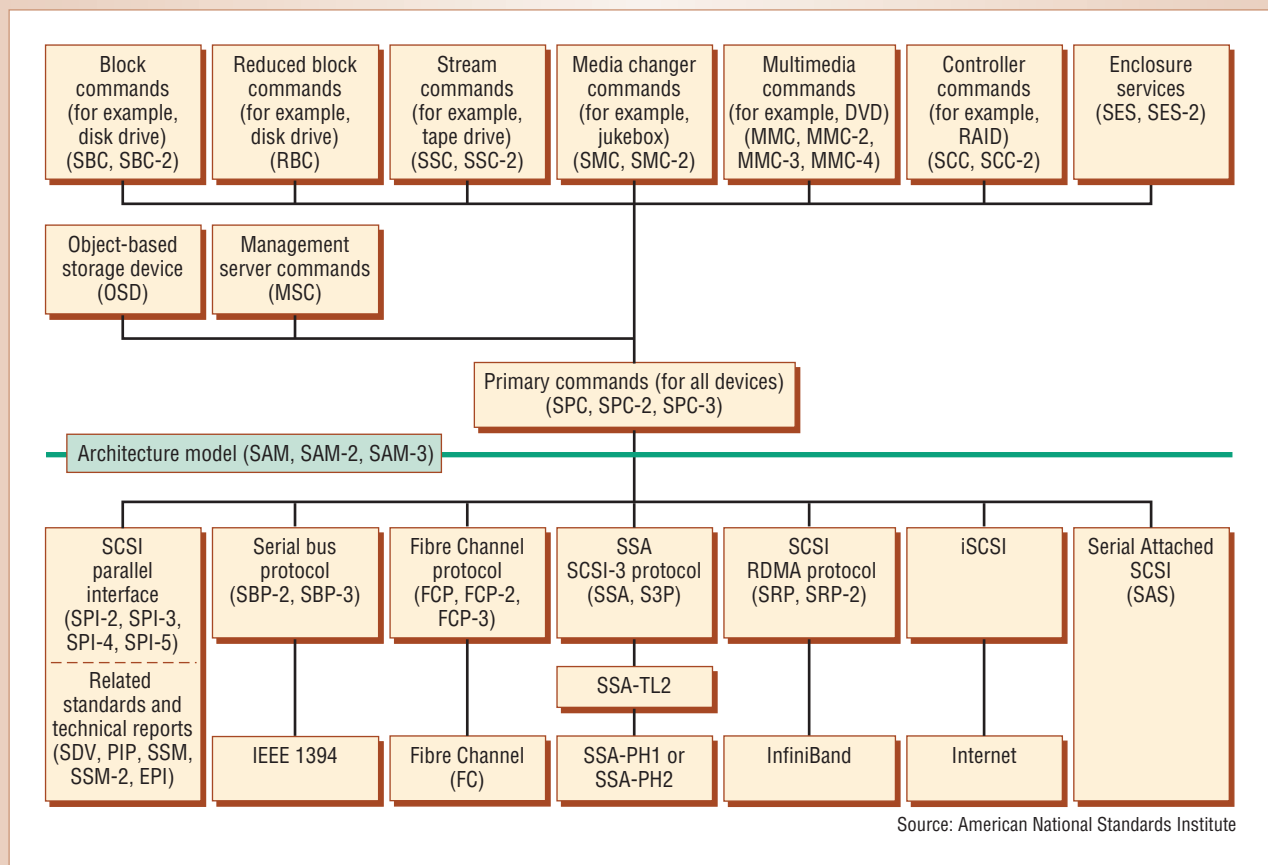
The advent of personal computers, mainly for the business community, and PC-based servers helped push development of a common interface for peripheral devices. IBM, Compaq, Leading Edge, Apple Computer, and other manufacturers began to ship standard hard drives with their systems in the mid-1980s. Largely based on Seagate Technologies' ST-506 and ST-412 interfaces, these drives offered limited interface logic—data went straight from the read and write channel to the host attached controller card for processing.

However, lack of interoperability among platforms frustrated early users who wanted to upgrade existing systems. In response to this problem, Seagate's Alan Shugart began developing the Shugart Associates System Interface. Presented to the American National Standards Institute in 1982, SASI enjoyed widespread use and industry support. ANSI changed the name to Small Computer System

Interface and in 1986 formalized what is now known as SCSI-1.

The original SCSI model incorporated only one command set, transport, and physical definition. However, as Figure A shows, it has evolved over the years to include numerous definitions and implementations. No other storage standard exists, or has existed in the past, with as much depth and breadth as SCSI. Depending on your view, this tenacity has either hamstrung the industry or is a testament to the model's success.

SCSI's primary asset is the ability to address physical blocks of data on a hard drive without knowing the drive's exact geometry. Advanced Technology Attachment Packet Interface, Universal Serial Bus, IEEE 1394 (FireWire), and Fibre Channel all enable and transport SCSI command sets, preserving upper-layer protocols in application clients and allowing new topologies without having to change the layers between applications and hardware.



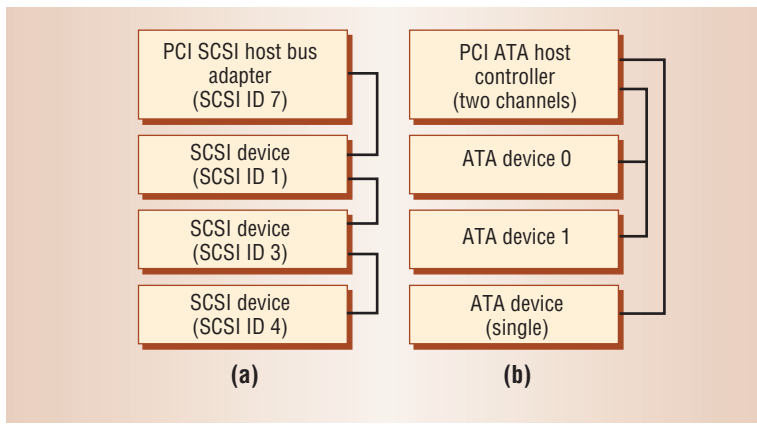
**Figure A. SCSI Architecture Roadmap, as of October 2002. Compared to the original SCSI model, which had only one command set, transport, and physical definition, the current storage standard offers unprecedented depth and breadth.**

to development of proprietary parallel interfaces, including Sony's well-known parallel CD-ROM interface, and other technologies for attaching CD-ROM drives to the ISA bus without SCSI.

The Advanced Technology Attachment (ATA) disk interface did not require advanced SCSI features—its firmware was simpler, and its application-specific integrated circuits were smaller and easier to build. With the advent of the Peripheral Component Interconnect (PCI) interface, ATA tech-

nology became the motherboard standard for PC hard drives and CD-ROM drives.

The ATA Packet Interface (ATAPI) is a method for wrapping SCSI CDBs and related data around standard ATA task delivery mechanisms. With ATAPI, CD-ROMs and DVDs of all types can operate on a non-SCSI bus. The interface also delivers data for use in other devices such as Iomega's Zip drives, magneto-optical disks, and other removable media drives.



**Figure 1. Common parallel (a) Small Computer System Interface (SCSI) and (b) Advanced Technology Attachment (ATA) cable plants. Both technologies use signal sinking or termination theories to minimize transmission bounce and crosstalk, along with signal and ground pairs for higher-speed cables.**

Parallel ATA cannot operate outside the computer, so it is not a viable solution for nonstorage SCSI devices such as scanners and printers. Universal Serial Bus and IEEE 1394, or FireWire, are physical and transport mechanisms that use SCSI commands, whether wrapped in ATAPI packets or delivered through direct memory access (DMA).

The basic SCSI command delivery topology lends itself easily to a dedicated DMA controller, which can use the ISA, PCI, or IEEE 1394 bus, making the data transfer phase extremely fast. However, because ATA was an extension to the ISA interface, it originally did not use even built-in secondary DMA features. When the ATA community realized they were in a battle with SCSI promoters for PC adoption, they began to advocate newer methods that offered faster data transfer than the original programmed I/O. This led to the introduction of PIO modes 2 and 3, multiword DMA modes 0 through 2, and Ultra DMA modes 0 through 4, leading finally to UDMA 133.

Because physics limit speed improvements based on theoretical maximums of cable topology, SCSI and ATA are moving toward standards-based serial implementations. Figure 1 depicts the basic cable plants of parallel SCSI and ATA deployments. Both use signal sinking or termination theories to minimize transmission bounce and crosstalk, along with signal and ground pairs for higher-speed cables.

Parallel SCSI defined ground and signal pairs for transmission length and stability, but many manufacturers ignored these specifications for lower-speed implementations. Most notably, Apple Computer shipped systems with SCSI connectors missing more than half of the required ground pairs. Apple's commitment to SCSI led the PC industry in adoption of SCSI devices, just as its adoption of USB drove USB device availability; however, the company eventually moved away from SCSI to the less expensive ATA storage standard as a cost-saving measure.

The ATA topology allows for up to two devices per channel on a typical two-channel controller, whereas the basic 8-bit SCSI topology permits up to seven devices besides the host on one bus, each

with its own addressable ID. Each physical ATA device shares its channel with another device, allowing only one device access to the channel with an active command sequence at any given time.

Parallel SCSI topologies, however, allow for an active initiator (host), target (device), logical unit (LUN), and queue (commands) set at any one time. Because this I\_T\_L\_Q nexus is not limited to one active command sequence for each device, compliant devices can have multiple outstanding command queues active at once. However, even during DVD playback on an ATAPI-based system, the overall system speed of today's PCs is sufficient to stream these high bit rates with no requirement to overlap commands.

As storage topologies began evolving from parallel to serial or fabric interfaces, manufacturers incorporated proven storage basics to increase distance and speed. Since the early 1990s, new topologies such as Fibre Channel, serial storage architecture (SSA), USB, and IEEE 1394 have had to leverage the work done for SCSI or ATA in some way or face slow and painful adoption.

Just as Apple lost its commitment to SCSI, the company is no longer standing behind high-speed USB, or USB 2.0, because it believes the high-end transport space belongs to FireWire. USB 2.0, however, has huge adoption rates in mainstream PC platforms, and Apple will eventually remain the only PC supplier that ships FireWire on every platform without USB 2.0 support.

### Serial storage topologies

Serial transports have provided access to storage on servers since the advent of networking. Although *network-attached storage* (NAS) is a fairly new term, most major PC and server operating systems have provided shared data storage for quite some time. Emerging from the need to separate this function from general-purpose or networked operating systems, NAS has fueled the growth of an industry that was worth US\$1.84 billion in 2001. As a storage topology, it is unique in that the decision to purchase NAS is likewise a decision not to buy a Novell- or Windows-based file server.

Serial storage topologies make it easier to attach dedicated block-based storage to smaller and mobile PCs, eliminating bulky SCSI cables. Apple developed the IEEE 1394 for attaching high-speed serial devices such as scanners, video equipment, and printers, but FireWire also quickly became the focus of storage engineers as a replacement for external SCSI connections.

**Serial ATA.** Recently adopted by the industry and

proposed as an addition to the current T13 ATA specification, Serial ATA will likely become the standard for PC storage, surpassing parallel ATA shipments by 2004. This implementation differs from parallel ATA by offering a star topology, shown in Figure 2, in which each device has dedicated channel bandwidth with direct connection to the host controller.

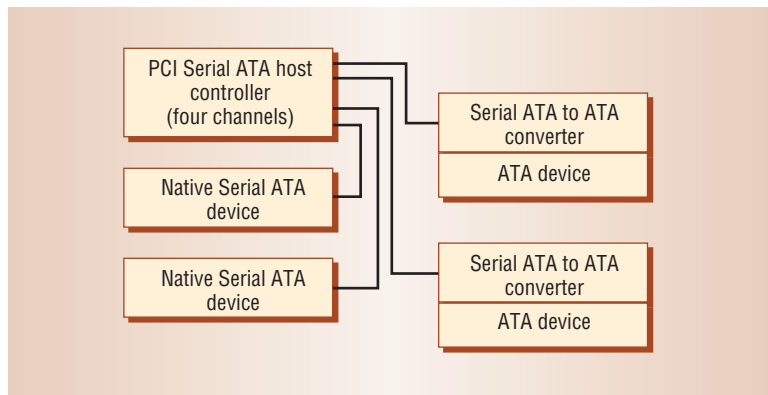
Initial PCI-based host controllers that ship in support of Serial ATA will be backward compatible with parallel ATA, so that existing operating systems can function. Ensuring that customers can adopt the technology without forcing change in the entire system is a key point of adoption and should guarantee its success.

**Serial Attached SCSI.** While not as far in development as Serial ATA, Serial Attached SCSI is the newest solution to the parallel bus problem for SCSI proponents. Borrowing from the physical and some of the transport work done for Serial ATA, SAS offers a way to encapsulate SCSI CDBs into packets that route natively over the same physical wires as Serial ATA. It also defines *expander devices*, downstream protocol converters that live as new definitions in the SCSI delivery subsystem.

SAS coexists with Serial ATA by defining methods for SAS-compliant infrastructures to include Serial ATA host and target ports, delivering their uniquely formatted command packets simultaneously over the SAS topologies. This unique approach validates Serial ATA in both physical wire definition and basic transport technology, while suggesting a better use for the topology.

The industry is now less fractious than it once was—most companies with ATA positions want SCSI to succeed and thus all topologies to continue. SAS's use of Serial ATA also lets companies charge more for devices that benefit from the work done for the lower-cost interface, without the initial upfront effort of reengineering a complete solution. SAS promoters have submitted a draft to the T10 SCSI committee proposing the technology as the next logical standards effort for SCSI. Ultra 320 SCSI is expected to be the last viable parallel SCSI (SPI-4) definition, leaving proponents of Ultra 640 (SPI-5) or higher parallel clocking rates looking for dedicated niche markets.

**Serial storage architecture.** Developed by IBM in the early 1990s, the serial storage architecture has been limited almost entirely to the company's own products. IBM officially submitted SSA's specifications to the T10 committee in hopes of standardization and industry adoption, but its lock on the interface and the topology's interoperability dis-



**Figure 2. Expected common Serial ATA topology. Each device has dedicated channel bandwidth with a direct connection to the host controller.**

sued most mainstream vendors from using it. IBM continues to sell limited quantities of SSA in support of its SSA-enabled server products.

## DEVICE TOPOLOGIES

Hard drives, CD-ROMs, tape drives, and other fixed-connectivity devices have little life past their original topology. The storage industry has expended much effort ensuring that customers who have purchased a device cannot easily move it to a new deployment when technology changes.

### Hard drives

ATA drives, currently only available in parallel form, do not allow attachment outside the servers or boxes where the connections exist. This limitation forces any ATA drive deployment into a bridged, or protocol-converted, operation for use as an external storage topology—an increasingly common practice in departmental and enterprise storage deployments.

NAS applications first used ATA drives for external storage, with pioneers such as Snap developing topologies based on Ethernet connections using shared file systems. High-end NAS vendors such as Network Appliance have only recently begun substituting ATA drives for SCSI or Fibre Channel drives in their lower-end applications.

This lack of focus on ATA as a viable topology for enterprise storage has motivated those companies promoting Serial ATA to begin defining Serial ATA II, providing new opportunities for manufacturers eager to bring this technology into the high-end storage space. Serial ATA II will offer built-in features for cyclical redundancy check (CRC) error correction through the entire system, overlapped commands similar to SCSI's queued commands, and multiple connections for high availability.

Despite its superior performance as a networked storage topology in SANs and storage fabric applications, Fibre Channel has not seen wide adoption as a native hard drive interface. The same issue underlying the debate between ATA and SCSI—cost versus market penetration—has also forced leading hard drive manufacturers to slowly adopt native Fibre Channel interfaces. Adoption rates

**The desire to share expensive devices drives protocol bridging and routing in the enterprise space.**

have been historically low because native parallel SCSI hard drive interfaces, bridged by RAID (redundant array of independent disks) or aggregation methods, are less expensive than deploying Fibre Channel connections natively on the hard drive. However, this trend is changing as more vendors ship complete Fibre Channel storage systems and more hard drive manufacturers ship native Fibre Channel drives.

Hard drive configurations in which the disks are present as drive sets rather than single entities are also topologies. Just a Bunch of Disks (JBOD) arrays usually deploy for the simplistic localization of the resource but are a shared resource if aggregation devices or software exist in the topology. This aggregation can be as simple as coordinated software running on the attached hosts, limiting access to individual logical units of the JBOD—typically deployed as single SCSI or Fibre Channel entities, with subservient LUNs—or as advanced as intelligent storage routers that map the resources.

RAID subsystems that allow the definition of LUNs by disk mirroring for data security, disk striping for data performance and security, or a combination of both present those LUNs to attached hosts for use as virtually defined disks. These subsystems act on the disks' set, enabling easier management and sharing.

### **Tape**

For backup and archiving, no technology is more cost-effective and flexible than tape drives. Although their popularity among PC users declined with the advent of CD-RW and other removable media formats, the 11 September 2001 terrorist attacks galvanized enterprises to invest in such devices to protect their installed information base for future restoration. Thus, tape-drive manufacturers have continued developing new technologies despite the recent downturn in IT spending and sliding overall revenues.

Tape drives are inherently lower-performance devices than direct-access hard drives or optical devices. Most attach directly to the hosts they support, even when the device is a large library supporting a media changer and multiple drives and media slots. Although they offer unlimited, off-site archival storage, tape libraries are also expensive—IT managers cannot justify buying one library for each host, thus forcing sharing of the system. Together, these limitations drive tape drive deployment in network configurations.

At its core, SCSI does not prohibit using multiple

initiators—or application clients—on one bus; rather, the operating systems or applications running on the hosts cannot respond correctly when additional initiators exist on the same bus. The RESERVE and RELEASE commands are the basis for sharing, but command interoperability, host behavior with reserves, and lack of support when LUNs reset make this feature unusable in direct attached SCSI environments.

Hard drives have the same limitation; however, hosts usually see only those hard drives that they need access to, either through virtualization of the LUNs with RAID devices or fabric zoning, masking these interoperability issues. Because IT departments prefer to buy one large tape library and share it among different hosts, this problem occurs almost exclusively with libraries.

### **Bridging devices**

For PCs, protocol conversion aims to expand the base of existing products for particular manufacturers rather than address problems with device sharing. No practical applications exist for sharing storage devices on PC platforms, as the entry cost for individual storage devices is low enough to allow a separate purchase decision. However, the desire to share expensive devices—either to extend their useful life or as an initially shared application—drives protocol bridging and routing in the enterprise space.

Storage routing—terminating one storage protocol for redirection to another—leads to new topologies. Accessing parallel SCSI devices as actual fabric devices, for example, lets IT managers deploy one topology without replacing devices or application clients. Intelligent storage routers use the appropriate SCSI commands to either map individual logical units to attached application clients or expose LUNs to multiple application clients. This solution makes it possible to share a library or any other mapped device by preventing a misbehaving host from disrupting another host's use of the library. It also provides a proxy-based solution for implementing RESERVE and RELEASE support.

Storage routing has seen limited success as parallel-to-parallel SCSI bridges, or extenders. This niche storage industry allows extending the interface's bus length—usually less than 20 meters—while possibly offering a way to share parallel SCSI devices. These simple extenders, which limit a host's access to one SCSI target or ID-based object, cannot gracefully accommodate access to an individual SCSI target's LUNs. Intelligent storage routers resolve this problem for any simple parallel-to-par-

allel SCSI extension, or any protocol-to-protocol conversion, by restricting granular access to the LUNs of attached storage.

Figure 3 shows a basic storage router implementation for libraries on a SAN fabric. In this configuration, any one of the Fibre Channel application clients can access the SCSI devices attached to the storage router, resulting in a device that resembles a native Fibre Channel storage device.

In addition to protocol conversion, many intelligent storage routers support server-free backup using the SCSI EXTENDED COPY command and advanced SAN-based management of parallel SCSI devices. More importantly, storage routing also allows sharing of equipment and continued use of existing resources—while they still have a useful life—on newly developed SANs. In addition, intelligent storage routers let non-Fibre Channel-capable application clients access Fibre Channel storage and networks by terminating the SCSI protocol from the host and reissuing Fibre Channel commands to the storage. This solution reverses the extended life equation by allowing existing application clients that support SCSI to access new Fibre Channel storage topologies.

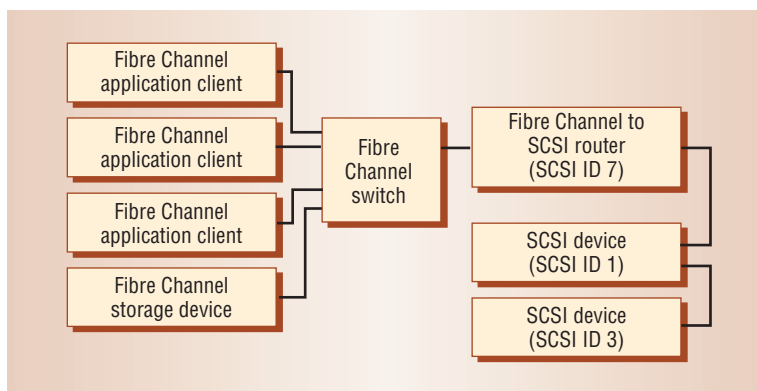
### FIBRE CHANNEL AND SAN TOPOLOGIES

Overcoming Fibre Channel's preeminence in the storage industry, now surpassing the installation rate of parallel SCSI implementations, will present a challenge for new topologies. However, growth in IP storage technologies will begin to force SAN technologies away from the Fibre Channel infrastructure at a faster pace than the SCSI-to-Fibre Channel conversion.

NAS infrastructures are inherently IP based, as file-sharing technologies favor client systems over user-based LAN or IP composition. However, this will not lead to automatic adoption of block-based storage over IP, even though IT organizations understand in detail both technologies. Potential overlap of NAS and SAN, or the convergence of NAS and SAN into common repositories of resources and infrastructure, will drive common topologies. Why repeat wires, switches, and hubs of divergent technologies when you can use the same wires for both block- and file-based storage? This question is not easy to answer for IT managers who have invested heavily in Fibre Channel, so early adopters of complete IP storage solutions will be those departments that have little to no Fibre Channel infrastructure.

### Fibre Channel

To date, implementers of most SAN installations



**Figure 3. Basic storage router and protocol bridging topology. Any one of the Fibre Channel application clients can access the SCSI devices attached to the storage router, resulting in a device that resembles a native Fibre Channel storage device.**

have built on the years of work invested in making Fibre Channel hosts, devices, and wires interoperate. The basic serialization of block storage over Fibre Channel has been understood much longer than the solutions have been available.

Encapsulating the basic SCSI CDB in header data, command and sequence tracking, and packet error correction guarantees in-time and in-order delivery of commands and data, promoting the topology's stability and application client usability. This involves establishing an I\_T\_L\_Q nexus at command delivery time via the host and device's Fibre Channel worldwide name (WWN) association. Parallel SCSI targets discover devices through bus arbitration, whereas devices in a Fibre Channel infrastructure discover through loop initialization or name services.

After hosts have discovered the targets, however, using the target from the view of the application client's upper-layer protocols is the same except for Fibre Channel's data packet encapsulation. By restricting the application client and device server, or target, to SCSI command processing, encapsulation does not disrupt applications using the storage—put simply, the wires change, not the host system applications.

Fibre Channel technology has become the standard for SAN deployments, moving beyond the original arbitrated loop (AL) definitions to switched fabrics that let multiple devices access a large number of port connections. Designed to replace direct attached storage (DAS), most Fibre Channel implementations that began to ship in the mid-1990s broke parallel SCSI's 15-device limitation and permitted much longer distances between host and devices.

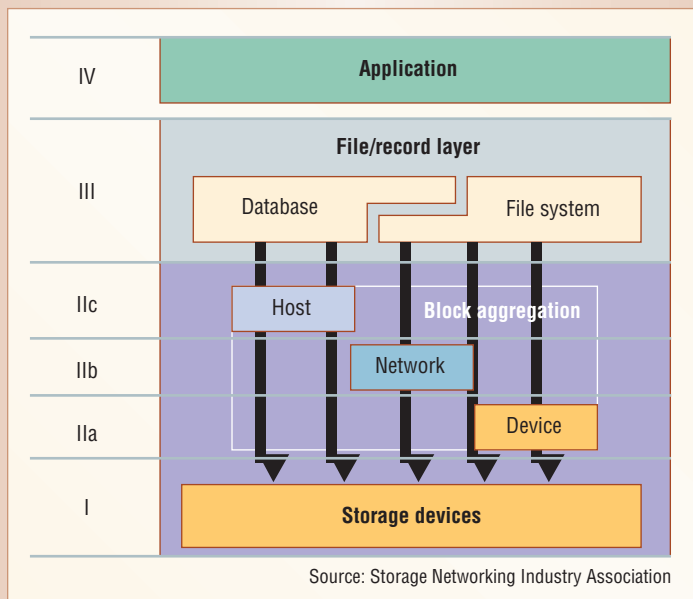
Yet even with these advancements, Fibre Channel-AL-based solutions did not solve the problems associated with multiple hosts on one bus or loop or let multiple hosts use or share a single device easily. Sharing storage, even in today's SAN deployments, is neither easy nor common. Fibre Channel fabric innovators spent many person-hours defining interconnections between like ports, like devices, and their switches, but not between each other's products. Initial shipments of Fibre

## Shared Storage Model

The Storage Networking Industry Association's Shared Storage Model, shown in Figure B, describes how independent storage functions are layered and the methods those layers use while interacting. Rather than serve as a specification, design guide, certification method, or advertising tagline, SSM is designed to promote mutual understanding of how networked storage devices and features interoperate.

According to the SNIA ([http://www.snia.org/tech\\_activities/shared\\_storage\\_model/Shared\\_Storage\\_Model.pdf](http://www.snia.org/tech_activities/shared_storage_model/Shared_Storage_Model.pdf)), SSM "may be used to describe common storage architectures graphically, while exposing what services are provided, where interoperability is required, and the pros and cons of each potential architecture.... The model does not represent any value judgments ... [but] makes it possible to compare architectures, and to communicate about them in a common vocabulary." This lets vendors differentiate their products and customers make buying decisions more easily.

SSM offers a common method for describing SAN components and their relationship to the network as a whole. Despite its name, the model has not helped achieve what most IT managers would deem true "shared storage"—the ability to have any machine access any storage, regardless of type or format, at any time. As a lofty and unlikely goal, the SSM strives to define what *could* achieve this level of sharing, even though most customers will never experience it.



**Figure B. SNIA's Shared Storage Model. Roman numerals are used for the layers to avoid confusion with the International Organization of Standardization and Internet Engineering Task Force networking stack numbers.**

Channel switches and infrastructures, sold as homogeneous platforms only, guaranteed interoperability by the manufacturer and distributor; however, these guarantees did not cover expansion with a different manufacturer's product.

The push for interoperability in the SAN industry is more disjointed than that for DAS-based storage, in large part due to the number of associations that support it. These include the Fibre Channel Industry Association (FCIA), Desktop Management Task Force (DMTF), American National Standards Insti-

tute (ANSI) InterNational Committee for Information Technology (INCITS) T10 and T11, Network Data Management Protocol (NDMP), and Storage Networking Industry Association (SNIA). Many of these bodies have evolved from marketing initiatives to full standards bodies driving adoption of SAN, NAS, and storage networking as a whole.

Since its inception in the late 1990s, the SNIA has become the leader in storage networking standard definition, focusing on interoperability. The SNIA's mission is "to ensure that storage networks become efficient, complete, and trusted solutions across the IT community" (<http://www.snia.org/about/>). Companies of any size and market share can participate for the benefit of the industry as well as their own advantage.

The SNIA has described many interoperability efforts, including expanded use of DMTF's Common Interface Method, an object-based technique for managing multiple storage device types. As described in the "Shared Storage Model" sidebar, SNIA's Technical Committee has also defined "a framework that captures the functional layers and properties of a storage system, regardless of the underlying design, product, or installation" ([http://www.snia.org/tech\\_activities/shared\\_storage\\_model/Shared\\_Storage\\_Model.pdf](http://www.snia.org/tech_activities/shared_storage_model/Shared_Storage_Model.pdf)). While not intended as a standard or even an architecture, SSM will drive adoption of common theories by storage vendors who believe in its premise. In addition, the SNIA has formed an Interoperability Committee to create standardized tests, test suite definition and creation, and continued lab work.

The SAN community continues to focus on Fibre Channel-based topologies. Although newer IP-storage-based efforts will build on Fibre Channel's development, they probably face even more obstacles to adoption. For IP storage to become an effective SAN platform, node naming, device discovery, and error correction require clear definitions, the same hurdles Fibre Channel deployments had to overcome before sales achieved critical mass.

## SANs

The technologies underlying simple Fibre Channel installations also support basic SAN installations. Switched fabrics let many hosts detect multiple storage devices and, with appropriate gate keeping, ensure effective storage allocation. Most SAN deployments do not provide shared storage—a given host can use any device, irrespective of type or format—but rather segment specific storage devices for their appropriate hosts, furthering the DAS concept. Put simply, hosts must have complete and unfettered


access to the storage allocated to them. Further, most hosts deployed in a switched-fabric SAN will not tolerate the dynamic removal of existing storage devices or the addition of new ones.

In closed Fibre Channel-based SAN environments, it is common to establish switch zones that limit specific hosts to specific storage targets, along with storage routers for LUN mapping, to provide a virtual DAS topology that can expand to include more hosts and devices. Figure 4 depicts a simple Fibre Channel-based SAN with multiple inter-switch communication paths and limited redundancy between host and storage device. As storage traffic flows through a switch, it interrogates the Fibre Channel packet header information to guarantee that only packets with appropriate port worldwide names go to the identified Fibre Channel device. The storage router further restricts access by presenting SCSI devices used in particular zones as LUNs of its WWN or as distinct WWNs.

Ethernet segments or virtual private networks make it possible to apply SAN topologies to IP storage fabrics. Because IP storage traffic can easily use installed base devices such as IP switches and routers, the ability to separate IP-based hosts from their associated storage extends the DAS model to IP SANs. The advent of Internet SCSI (iSCSI) as a replacement for Fibre Channel-based SANs is months if not years away, but this technology will clearly become the next-generation method for configuring SAN fabrics.

As with Fibre Channel storage, iSCSI isolates hosts' upper-layer protocols from the fabric's underlying hardware and topology mechanics; replacing an existing system's Fibre Channel infrastructure with IP-based devices allows the hosts and storage to continue running as expected. Because iSCSI also encapsulates the SCSI CDB and associated data into packets, then transports those packets to the right location based on IP routes established in the LAN, only the conversion to iSCSI delivery mechanisms is required. The main technical obstacles to iSCSI adoption are interoperability and speed, but forthcoming solutions that address discovery, naming, booting on the fabric, and other issues will facilitate implementation soon.

Most current corporate IT installations center on Fast Ethernet and would therefore not sustain the throughput needed to support block-oriented storage traffic. Although these environments support user-based data traveling to NAS servers, IT departments need a compelling reason to upgrade to a full Gigabit Ethernet infrastructure for SAN deployments, including the trained staff required to under-



**Figure 4. Simple SAN with limited Fibre Channel path redundancy. Each switch guarantees that only packets with appropriate port worldwide names go to the identified Fibre Channel device, while the storage router presents SCSI devices used in particular zones as logical units of its WWN or as distinct WWNs.**

stand storage—IP or Ethernet experts are not storage experts simply because the storage is IP based.

Fibre Channel and iSCSI aside, SAN deployments will continue increasing as IT departments learn that system and storage consolidation topologies solve many of the physical problems presented by DAS topologies, such as having to cable storage devices directly to the host. SANs also make it possible to repurpose existing equipment with storage routers, while their flexible fabric configurations ease deployment of new fabric-aware application clients and storage.

**S**torage topologies, in both cabling and protocols, exist to satisfy specific storage requirements. The storage's relative location to the application client usually determines which topology represents the best or simplest solution for a given installation. Because most topologies underutilize storage devices, it is important for customers to evaluate different models' usage characteristics before deployment. For example, although SAN and NAS topologies target large enterprises, small and medium-sized businesses can benefit from their use. With information repositories growing at exponential rates, understanding storage topologies will become increasingly important for IT managers at all levels. ■

*Robert Griswold is chief technologist and director of marketing at Crossroads Systems, a developer and manufacturer of storage networking solutions based in Austin, Texas. He received a BS in electrical and electronic engineering from California State University, Sacramento. His research interests include SCSI and storage history, operating systems, and computing futures. Griswold is a member of the IEEE, the IEEE Computer Society, and the Tau Beta Pi Engineering Honor Society. Contact him at bobntrudy@aol.com.*