Storage Area Networks (SANs)

Vladimir V. Riabov, *Rivier College*

INTRODUCTION

The volume and value of enterprise data have been growing faster than the speed at which traditional backup company utilities have been increasing. Enterprises have become more dependent on their online systems and cannot ensure the availability of these data by relying only on traditional, network-bottlenecked, server-attached storage systems. Solutions to this problem have been offered by SAN technology, which provides enterprises with serverless zero-time-windows backup. In this new backup approach, backed-up data are removed to a secondary remote storage device, and the enterprise server becomes off-loaded, permitting high-performance continuous access to both applications and data. A SAN fabric is usually based on fibre channel technology that allows up to 10-km long-distance connections. This feature has a significant advantage in a campus environment where reliable backup resources can be shared among several divisions. For example, fibre channel SANs for the Healthcare Enterprise allow 24 × 7 continuous operation, patient record backup, and medical image archiving (Farley, 2001).

The SAN becomes a key element of the enterprise environment where data availability, serviceability, and reliability are critical for a company's business. Many enterprise solutions (e.g., ATTO FibreBridge products, rack mount solutions, ATTO FibreCenter3400R/D, host bus adapters, the ATTO Diamond array, Compaq StorageWorks products, EMC Connectrix solutions, LSI Logic E4600 Storage System) are available today. They can be effectively used in multiple platform storage-infrastructure solutions for data-intensive applications such as e-commerce, online transaction processing, electronic vaulting, data warehousing, data mining, Internet/intranet browsing, multimedia audio/video editing, HDTV streaming, and enterprise database managing applications.

This chapter describes fundamentals of storage area networks (SANs), their architectural elements (interfaces, interconnects, and fabrics), technologies (fibre channel-arbitrated loop transport protocol, Brocade's configurations, InfiniBand switched-fabric architecture, crossroad systems with a storage router, virtual interface architecture, direct access file system, IP storage technologies, SANs over IP, fibre channel over IP, Internet fibre channel protocol, Internet SCSI, storage over IP, fabric shortest path first protocol, storage resource management, and adaptive network storage architecture), solutions, standards, associations, initiatives, forums, coalitions, vendors, and service providers.

SAN FUNDAMENTALS

What Is a SAN?

A SAN (*storage area network*) is a networked high-speed infrastructure (subnetwork) that establishes direct access by servers to an interconnected group of heterogeneous storage devices such as optical disks, RAID arrays, and tape backups, which are effective for storing large amounts of information and backing up data online in e-commerce, online transaction processing, electronic

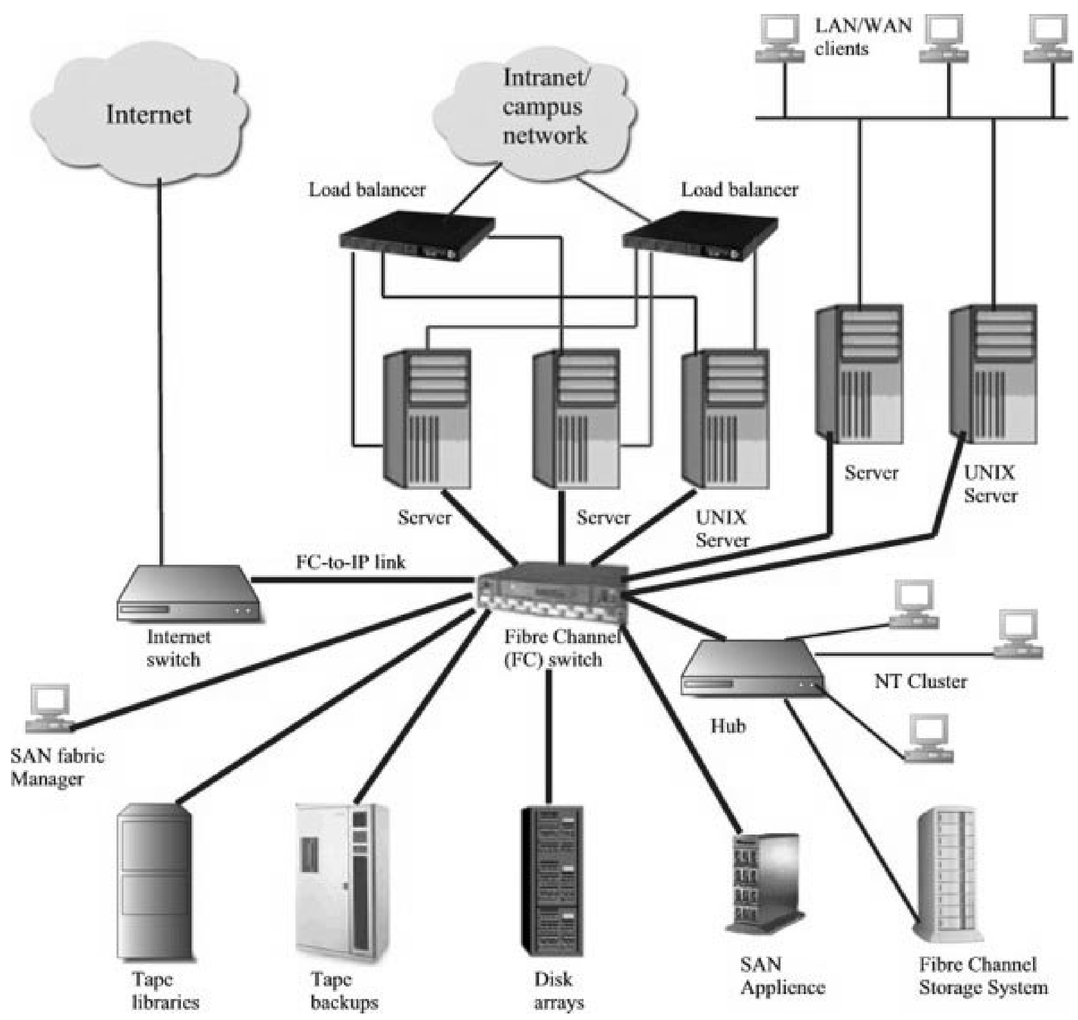

Figure 1: Storage area network as a networked high-speed enterprise infrastructure.

vaulting, data warehousing, data mining, multimedia Internet/intranet browsing, and enterprise database managing applications. SANs provide additional capabilities (fault tolerance, remote management, clustering, and topological flexibility) to mission-critical, data-intensive applications. A SAN is typically a part of an enterprise network of computing resources (Sachdev & Arunkundram, 2002). A simple model of the storage area network as a networked high-speed infrastructure is shown in Figure 1.

A SAN can be considered as an extended and shared storage bus within a data center, consisting of various storage devices and specific interfaces (e.g., fibre channel, ESCON, HIPPI, SCSI, or SSA) rather than the Ethernet (Peterson, 1998). In order to be connected to the enterprise network, the SAN utilizes technologies similar to those of LANs and WANs: switches, routers, gateways, and hubs (see Figure 1). Wide area network carrier technologies, such as asynchronous transfer mode (ATM) or synchronous optical networks, can be used for remote archival data storing and backup. As an important element of modern distributed networking architectures of storage-centric enterprise information processing, SAN technology represents a significant step toward a fully networked secure data storage infrastructure that is radically different from traditional server-attached storage (Clark, 1999). In addition, SANs provide improved options for network storages, such as a creation of remote or local, dedicated or shared data storage networks with access capability faster than *network attached storage (NAS)*.

SANs are based on the storage-centric information-processing paradigm, which enables any-to-any connectivity of computers (servers) and storage devices over a high-speed enterprise network of interconnected fibre channel switches that form the SAN fabric. The incidence of unconnected clusters of information is eliminated or significantly reduced by SANs. According to this concept, a SAN resides behind the server and provides any users or devices on the enterprise network ("clients") with fast access to an array of data storage devices. It can be viewed as multihost connected-and-shared enterprise storage. Adding new storage devices and server elements resolves traditional network-bottlenecks and small-scale problems of interfaces, such as the small computer systems interface (SCSI) and network attached storage (NAS), and can easily expand the scale of the SAN (Thornburgh & Schoenborn, 2001). Another advantage of SAN technology is that backups can be made over the SAN fibre channel subnet, and, in this case, backup traffic is totally removed from the enterprise network.

The SAN represents a new segment of the information services industry called storage solution providers (SSP). However, isolated SANs cannot realize SSPs' services, such as real-time data replication, failover, storage hosting, and remote vaulting.

Benefits of SANs

A SAN makes physical storage capacity a single, scalable resource and allows the flexible allocation of virtualized storage volumes (e.g., RAIDs, JBODs, and EMC, SUN, and DELL storage devices). The SAN can manage backup tasks that were a huge administrative and computer-resource burden under old storage architectures. The storage management cost savings can be higher than 80%. A cost-effective, scalable SAN enhances overall system performance. It can integrate legacy SCSI devices, which allows increasing their system-wide effective usable capacity by up to 30% (InfraStor, 2002).

SANs are an integral part of a large financial services enterprise, ISP, government organization, research laboratory, electronic publisher, digital video production group, TV-broadcasting station moving to digital services, or educational institution, or any organization with increasing data storage needs.

There are several key reasons for implementing a SAN (InfraStor, 2002). The first three concern business issues of return on the investment in data storage, as well as the protection of existing investments:

SANs are cost-effective (reduced cost of storage management, including backup and recovery; increased user productivity; cost-effective implementations of high availability disaster protection, using remote clusters and remote mirrored arrays);

SANs reduce business risk (faster disaster recovery; reduced revenue loss from down-time; reduced lost-opportunity costs);

Legacy investments are protected (SANs can be implemented without abandoning existing storage infrastructures such as devices using SCSI connections).

The next four address critical technical issues that face data-center managers at a time when the volume of data to be managed and made available in many organizations is increasing at a 60% annual rate (InfraStor, 2002):

SANs provide scalability (add servers and storage independently);

SANs allow flexibility (reconfigure storage and servers dynamically without interrupting their services; load sharing and redistribution);

SANs enhance overall system performance (more effective use of existing server compute cycles; real-time backup without impacting LAN/WAN; multiple server-to-storage paths; networked storage arrays that can outperform bus-attached storage; compatibility with parallelized database applications);

SANs are an integral part of any high-availability plan (facilitation of shared on-line spares and remote backup or mirroring; reduced down-time requirements; storage independent of the application and accessible through alternate data paths such as found in clustered systems).

The implementation of a SAN can realize significant overall cost savings in data-center operations and can increase user productivity. The opportunity to avoid escalating costs depends on decentralization of data and applications. A key element in the consolidation of data storage must include the implementation of a basic SAN infrastructure in order to provide (InfraStor, 2002):

Bandwidth to service clients;

Maintenance of data availability without impacting LAN bandwidth;

Scalability for long term, rapid growth with protection of legacy investments;

Flexibility to provide optimum balance of server and storage capacity;

Manageability for ease of installation and maintainability;

Shared access to data resources for real-time backup and recovery.

Distributed environments require high-cost maintenance in terms of staff resources. The consolidation of distributed NT-based storages to a virtualized SAN-based resource can save 80% or more of the costs of management (InfraStor, 2002).

SAN Applications

SAN applications cover the following areas of data transfer (Peterson, 1998): (1) the externalization of data storage out of the server-SAN-attached-storage (SAS) and NAS-with-SAN-interconnects network architectures; (2) clustering, a redundant process that provides failover, high availability, performance, and scalability through the use of multiple servers as a data pipe and allows data storage resources to be shared; (3) data protection solutions for backup, remote clustering, file mirroring, and replicating and journaling file systems by creating data storage redundancy on a dynamic basis; (4) data vaulting, which is the process of transferring archived data to less expensive media; (5) data interchange from one storage system to another or between different environments; and (6) disaster recovery, which is similar to data interchange, moving copies of data offsite, and is built on remote vaulting (backup) processes or on remote array mirroring or clustering. Several new applications benefit from 2 Gb/s fiber channel SANs (Hammond-Doel, 2001): multimedia audio/video servers that provide the ability to stream higher resolution files, medical imaging, prepress that speeds up design and file preparation, and video editing of uncompressed HDTV data.

The first effective application of SANs has been serverless backup, which provides enterprises with full-time information availability. All backup related tasks have been relegated to the SAN. Large enterprises can store and manage huge amounts of information (several terabytes or more) in the SAN high-performance environment. Enterprise servers are connected to storage devices (e.g., RAID arrays) via a high-speed interconnection, such as fibre channel. The SAN any-to-any communication principle provides the ability to share storage resources and alternative paths from server to data storage device. A SAN is also able to share the resources among several consolidated servers.

A cluster of interconnected servers may be connected to common storage devices in the SAN environment and be accessible to all clients. Modern enterprises employ this clustering technology to resolve several challenging application problems (Barker & Massiglia, 2001, p. 244), i.e., providing customers, partners, and employees with *continuous* application service, even if the enterprise systems fail, and supporting application performance growth as demand grows, without service disruption to customers. Clusters provide load balancing, high availability, and fault tolerance and support application scaling. In some implementations, the clustered servers can be managed from a single console. Clustering methodology is effectively used in e-commerce, online transaction processing, and other Web applications, which handle a high volume of requests.

SAN methodology has its roots in two low-cost technologies: SCSI-based storage and the NAS-based concept. They both successfully implement storage-network links, but are limited to a low volume of data flows and rates. SCSI still remains the most popular "bus-attached" server–storage connection in SAN-attached storage (SAS) systems, especially at the stage of transition from SCSI bus devices to fibre–channel switches using the SCSI-fibre protocol converter in a new enterprise storage ("data center") environment. In the network attached storage (NAS) system, storage elements (i.e., a disk array) are attached directly to any type of network via a LAN interface (e.g., Ethernet) and provide file access services to computer systems. If the NAS elements are connected to SANs, they can be considered as members of the SAN-attached storage (SAS) system. The stored data may be accessed by a host computer system using file access protocols such as *NFS* or *CIFS*.

SANs provide high-bandwidth block storage access over long distance via extended fiber channel links. However, such links are generally restricted to connections between data centers. NAS access is less restricted by physical distance because communications are via TCP/IP (InfraStor, 2001). NAS controls simple access to files via a standard TCP/IP link. A SAN provides storage access to client devices, but does not impose any inherent restrictions on the operating system or file system that may be used. For this reason, SANs are well suited to high-bandwidth storage access by transaction-processing and DBMS applications that manage storage access by themselves. NAS, which has the inherent ability to provide shared file-level access to multiple OS environments, is well suited for such requirements as Web file services, CAD file access by combined WinNT/2000, UNIX, and LINUX devices, and wide-area streaming video distribution (InfraStor, 2001). A balanced combination of these approaches will dominate in the future. See the chapter *NAS* (*Network Attached Storage*).

SAN Architecture

The SANs architectures have been changed evolutionarily, adapting to new application demands and expanding capacities. The original fibre-channel-based SANs were simple loop configurations based on the fibre channel arbitrated loop (FC-AP) standard. Requirements of scalability and new functionality had transformed SANs into fabric-based switching systems. Numerous vendors offered different solutions of problems based on fabric switching. As a result, immature standards created various interoperability problems. Homogeneous high-cost SANs were developed. Ottem (Ottem, 2001) refers to this phase as the legacy proprietary fabric switch phase. The latest architectural approach is associated with a standards-based "Open" 2Gb fabric switch that provides all the benefits of fabric switching, but based on new industry standards (FC-SW-2) and interoperability architecture that runs at twice the speed of legacy fabric. The standards-based switches provide heterogeneous capability. The latest feature reduces prices of the SAN's components and management costs of running a SAN. Characteristics of three generations of SANs are summarized in Table 1.

The Open 2Gb fiber channel allows doubled SAN speeds, enables greater flexibility in configuring SANs for a wide range of applications, and is especially useful for managing 1.5-Gb high-definition video data. In the HDTV applications, a single fiber can carry a full high-definition video stream without having to cache, buffer, or compress the data. Other examples (Ottem, 2001) include storage service providers that must deliver block data from to users at the highest possible speeds and e-commerce companies that have to minimize transaction times. The 2-Gb fiber channel provides the high-speed backbone capability for fibre channel networks, which can be used to interconnect two SAN switches. This configuration increase overall data throughput across the SAN even if servers and disk subsystems continue to operate via 1-Gb channels.

A SAN system consists of software and hardware components that establish logical and physical paths between stored data and applications that request them (Sheldon, 2001). The data transforms, which are located on the paths from storage device to application, are the four main abstract components (Barker & Massiglia, 2001, p. 128): the *disks* (viewed through ESCON, FCP, HIPPI, SCSI, and SSA interfaces as abstract entities), *volumes* (logical/virtual disk-like storage entities that provide their clients with identified storage blocks of persistent/retrieved data), *file systems*, and application-independent *database management systems*. In a system with a storage area network, five different combinations (Barker & Massiglia, 2001) of these data transforms and corresponding transformation paths serve different applications and system architectures by various physical system elements. The disk abstraction is actually the physical disk drive. The abstract volume entity is realized as an external or embedded RAID controller, as an out-of-band or in-band SAN appliance, or as a volume manager

Table 1 Three Generations of SANs

	Fabric	Main Characteristics	Applications
First-Generation SANs	1Gb Loop	FC-AL protocol; 1Gb speed; enabled first SANs	SCSI replacement
Second-Generation SANs	1Gb Proprietary; Legacy Fabric	FC-SW protocol; 1Gb speed; proprietary switch-to-switch connections; expensive	LAN-free backup; HA clustering
Third-Generation SANs	2Gb Open Fabric	Open FC-SW-2 protocol; 2Gb speed; standards-based switch-to-switch connections; competition-driven price reductions price reductions	Serverless backup; heterogeneous storage consolidation; high-definition video; data; virtualization

serving a database or an application. Storage servers (such as NAS devices), database servers, and application servers may contain the abstract file systems. These devices and servers can be clustered to increase scaling and application availability. In that case, their volume file and management systems should be cluster-aware (Barker & Massiglia, 2001).

Any SAN-based client–server system consists of three architectural components: *interfaces*, *interconnects* or *network infrastructure components* (switches, hubs, routers bridges, gateways, multiplexors, extenders, and directors), and *fabrics*. The SAN interfaces are fibre channel, ESCON, HIPPI, SCSI, and SSA. The SAN interconnects link these storage interfaces together, making various network configurations. Routers and bridges perform protocol transformation in SANs. Switches increase the overall SAN bandwidth by connecting system elements for data transmission and allow advantages of the centralized storage repositories with the shared applications and central management. The most common SAN fabrics are switched fibre channel, switched SCSI, and switched SSA, all of which physically link the interconnects and determine the SAN's performance and scalability. Some fabrics embed operating systems that provide for SAN security, monitoring, and management. Hosts are connected to the fibre channel SAN through host bus adapters (HBAs), which consist of hardware and interface drivers. Fibre channel HBAs support negotiation with network-attached devices and switches and allow the host to minimize its CPU overhead.

SAN Operating System Software Components

The SAN software plays an important role in providing an environment for various business and management applications, called *system applications* (Barker & Massiglia, 2001, p. 13), such as clustering, data replication, and data copying. The *management applications* (zoning, device discovery, allocation, RAID subsystems, and others) manage the complex environment of distributed systems. These applications can significantly reduce the cost and improve the quality of enterprise information services.

SAN TECHNOLOGIES AND SOLUTIONS

The SAN infrastructures support multiple protocols, such as SCSI, SNMP, VI, ESCON/FICON, TCP/IP, and SSAIP, over a single physical connection. This unique capability provides the SAN system with a coupled functionality of an interface to storage devices and a server interconnect.

In the early 1990s, fibre channel was developed by the Fibre Channel Systems Initiative (FCSI) and adopted later by the ANSI X3T11 Committee as a high-speed interface for connecting storage devices to servers and other network configurations. These interconnect standards provide SANs with the vital properties of connectivity, bandwidth, interconnectivity, protocol efficiency, distance range, recoverability, failure tolerance, and cost options. The fibre channel standards specify electrical and optical transmission media, as well as conventions for signaling and transmission/functional protocols. Optical medium (with SC, LC, and MT-RJ connectors) supports reliable signaling over long distances. Fibre channel provides data rates in the range from 133 Mbit/s to 4 Gbit/s over low-cost copper cabling (shielded twisted-pair wire or coaxial cable with DB-9 and HSSDC connectors) or higher-cost multimode fiber-optic cable. Fibre channel fabrics have transceivers, called gigabit interface converters (GBICs), which convert optical to electrical signals to cable connectors. The fibre channel technology supports distances up to 10 km.

Fiber-Channel-Arbitrated Loop Transport Protocol (FC-AL)

The fibre channel methodology has means to implement three topologies: point-to-point links, arbitrated loops (shared bandwidth loop circuits), and bandwidth-switched fabrics that provide SANs with the ability to do bandwidth multiplexing by supporting simultaneous data transmission between various pairs of devices. Any storage device on the loop can be accessed through a fibre channel switch (FCSW) or hub. The fibre channel switch can support entry-level (8–16 ports) to enterprise-level (64–128 ports) systems. Under the ANSI X3T11 standards regulation, up to 126 storage devices (nodes) can be linked in the fiber channel arbitrated loop (FC-AL) configuration, with the storage interface bandwidth about 100 Mbits/s for transferring large files. More than 70 companies, including industry-leading vendors of disk arrays and computer and networking systems, support the FC-AL voluntary standards. The FC-AL topology is used primarily to connect disk arrays and FC devices. Originally developed as the high-speed serial technology of choice for server–storage connectivity, the FC-AL methodology is extended to the FC-SL standard that supports isochronous and time-deterministic services, including methods of managing loop operational parameters and QoS definitions, as well as control. The FC-VI regulation establishes a fibre channel-virtual interface architecture (FC-VIA) mapping standard. See the chapter on fibre channels for more information about various implementations of the technology in various network configurations, including SANs.

Because of the high cost of the FC interconnect components and separation of storage and servers at the wide area network scale (resulting in slow capabilities of WAN–SANs with fibre channel), alternatives to FC technologies have been developed. The *ipStorage* technology (Barker & Massiglia, 2001, p. 187) employs TCP/IP as a storage interconnect. The Internet Engineering Task Force (IETF) has proposed the iSCSI (*Internet SCSI*) standards that address the issues of long distances (WAN-scale), reducing the interconnect cost, high security, and complex storage network topologies. The iSCSI is layered on top of the TCP/IP protocol hierarchy and can instantly access all modern transmission media and topologies. TCP/IP and related protocols have been implemented in the server-based systems that allow the most general storage networks to be constructed with the iSCSI methodology. The main challenge is a reduction of the iSCSI processor overhead of operating iSCSI packets below the Fibre Channel overhead level.

InfiniBand Solutions

InfiniBand is a new emerging interconnect technology, developed by the standards of the InfiniBand Trade Association (founded by Compaq, Dell, Hewlett-Packard, IBM, Intel, Microsoft, and Sun Microsystems) that offers the most general low-cost server system topologies (InfiniBand Trade Association, 2003). It is expected that InfiniBand interfaces will be embedded into all Intel-based servers (Barker & Massiglia, 2001, pp. 188–192; Intel InfiniBand Architecture, 2003) and will allow Windows and Linux servers to be available for resolving complex problems of data centers by adopting clusters and multihost enterprise RAID subsystems. The InfiniBand technology implements a switched-fabric architecture with the packet-switching communication protocol (PSCP) that relates to the virtual interface (VI) architecture methodology. SANs, parallel processing systems, and systems area networks can effectively use InfiniBand as a high-performance/low-latency interconnect. See the chapter *InfiniBand*.

Crossroads Systems With a Storage Router

InfiniBand technology has been successfully implemented by Crossroads Systems, Inc., which promotes storage solutions based on protocol-independent connectivity at Gigabit/s speeds and unparalleled manageability for various storage devices. The Crossroads' storage routers (e.g., Crossroads™ 10000) support peer operations between storage devices and multiprotocol servers on fibre channel storage networks.

Brocade's Configurations

Brocade Communication Systems, Inc., has developed an intelligent fabric services architecture that creates a scalable and secure environment for enterprise mission-critical storage applications such as data backup and business continuity. The Brocade SANs (SilkWorm™ family of fabric switches and software) provide enterprises with any-server-to-any-storage-device connectivity and consolidate storage resources and servers, as well as sharing backup resources (Beauchamp, Judd, & Kuo, 2002).

OTHER STORAGE NETWORKING TECHNOLOGIES

The following emerging technologies introduce new system architectural approaches in storage networking. SAN developers and users are trying to adapt them to a new enterprise environment that is characterized by host-level heterogeneous complexity, management flexibility, new TCP/IP network communication services, file-access-protocol developments, and the repartitioning of the functionality of the file management systems.

VI (Virtual Interface) Architecture

The virtual interface (VI) architecture is a midlayer protocol specification that regulates virtual intercommunication between applications running on different remote servers (i.e., in a cluster). This methodology significantly reduces the latency and the volume of the system I/O operations by using message and data buffer pools that are insensitive to heterogeneous operating environment or other applications. The reduction of the I/O-related interrupts increases the CPUs' time for processing various other system tasks. Developers of the VIA technology (Compaq, Intel, Microsoft, etc.) utilize this architecture as an efficient way of message communication between the SAN nodes at the application level, creating only a small overhead of intercommunication between the remote applications. This methodology has been successfully implemented in database managers and NAS devices (Barker & Massiglia, 2001, pp. 189–190). Several efforts (i.e., the Direct Access File System and Network File System initiatives) have been made to improve file system performance by utilizing VIA-type advanced transport protocol features. The Emulex Corporation promotes the GN9000/VI™ 1 Gb/s VI/IP PCI host bus adapter, which is based on the virtual interface (version 1.0) architecture, supports standard TCP/IP-reliable data delivery, IP routing, and the direct access file system (DAFS) standard, and speeds data access over standard Gigabit Ethernet networks.

Direct Access File System (DAFS)

The direct access file system is a new file access/transfer protocol that is based on CIFS/NFS characteristics and VIA-type transport protocol features. DAFS/VIA technology supports direct file transferring between the storage system and clients. In the SAN environment, data can be directly transferred among a number of servers.

IP Storage Technologies

Another *block-mode data* mechanism has been used by the IETF IP Storage Working Group in developing standards for a new IP-based transport-through-network technology that encapsulates fibre channel and SCSI high-speed interfaces and provides direct access to data on disks, tapes, and optical storage devices. IP storage technology allows embedding low-cost SANs into IP-based enterprise infrastructures over existing Gigabit Ethernet networks. See the chapter *IPStorage*.

SANs Over IP

To avoid the distance limitation of the fibre channel interconnects, enterprises build remote SANs that can be interconnected by means of the SAN-over-IP technology originally developed by the Computer Network Technology Corp. The distant SANs appear as local storage entities. This technology improves enterprise management and data access, disaster recovery, business continuity, disk mirroring, electronic tape vaulting, and wide area clustering. The Storage Networking Industry Association (SNIA) offers three technologies for integrating fibre channel SANs into the IP backbone. These methodologies include fibre channel over IP (FCIP), iFCP, and Internet SCSI (iSCSI). The FCIP, iFCP, and iSCSI transport protocol descriptions are presented in (Clark, 2002) and on the CNT Web site (CNT, 2002).

Fibre Channel Over IP (FCIP)

FCIP is the simplest point-to-point IP tunneling solution for intercommunicating remote SANs with fibre channel fabrics. The FCIP gateways establish TCP/IP connections over a WAN path to transport the fibre channel encapsulated frames. A typical discrepancy in data communication rates between an FCIP-attached WAN link and fibre channel fabric generates various flow control issues that can be resolved by TCP sliding-window algorithms. Several FC-FCIP management issues cannot be properly determined for the FCIP pipes because the TCP/IP transport component ends at the external nodes of the fibre channel network. These problems have been addressed and successfully resolved in the iFCP and iSCSI approaches.

Internet Fibre Channel Protocol (iFCP)

The gateway-to-gateway Internet fibre channel protocol (iFCP) supports a means of integrating fibre channel end devices into a single IP SAN. By using iFCP, the fibre channel fabric services can be provided to the remote FC devices over a TCP/IP network. The iFCP IP storage switches can directly connect fibre channel storage arrays, HBAs, hubs, switches, and routes. The iFCP is a protocol stack, which can be implemented in an IP storage controller interface or integrated into Gigabit Ethernet IP storage NIC (known as ANSI X3T10 and X3T11 Standards) (Clark, 2002, pp. 126–139). It supports any-to-any IP routing of storage data. A mismatch in data communication rates between an iFCP-attached WAN link and fibre channel fabric generates various flow control issues that can be resolved by TCP sliding-window algorithms. The IPSec, public or private keys, and zoning methods can provide security across the Internet. One of the important applications of the iFCP technology is the support of multiple TCP/IP connections for concurrent storage transactions.

Internet SCSI (iSCSI)

In contrast to the FCIP concept, the iSCSI methodology, which follows the SCSI client/server model, is based on the implementation of a light switch technology in IP storage networking (Clark, 2002, pp. 139–149) and excludes fibre channel elements. The iSCSI servers (*targets*) are present in disk arrays and client nodes (*initiators*) that occupy host platforms. The iSCSI protocol over the TCP/IP layer is used for block data transport between these entities over the IP network. Data can be directly written into application memory through a steering and data synchronization layer located below the iSCSI sublayer. IPSec, Keyberos, public key, and other methods can provide security across the Internet. SANs use the iSCSI adapters with TOEs to minimize processing overhead and realize high-performance features of the iSCSI technology. The enterprise solutions with IP SANs can also support the Gigabit and faster Ethernet on iSCSI-switch infrastructures.

Storage Over IP (SoIP)

Based on the SoIP remote storage technology, the Nishan Systems Corporation developed IP Storage switches of the IPS 4000 Series™and a suite of storage management software that allow configuration and monitoring of large-scale storage networks. See the chapter *IP Storage*.

Fabric Shortest Path First (FSPF)

Fabric shortest path first (FSPF) is the OSPF-based standard routing protocol for fibre channel that determines the next shortest route for data traffic, updates the routing table, and detects the failed routes (Vacca, 2002, p. 152). The optical, link, or switch failures can be effectively handled by FSPF with minimal impact on the interconnected devices in the FC/SAN environment.

Adaptive Network Storage Architecture (ANSA)

The Procom Technology Corporation has developed the adaptive network storage architecture (ASNA) approach, which delivers both block level and file access to data.

Procom's NetFORCE3000™ Series provide filer functionality (together with advanced features of security, high stability, backup, and recovery) to an enterprise storage that can result in high-performance information management systems. The ASNA technology has been successfully applied to database management, data warehousing, e-mail delivering, and 24 × 7 rich media applications.

Storage Resource Management (SRM)

The SRM technology provides applications for managing logical and physical storage-system resources (virtual devices, disk volumes, file resources, storage devices and elements, and appliances). SRM tools allow storage-system administrators to configure and monitor SANs and other storage resources. During the administrative monitoring, the transport or storage data remain unchanged. Vendors of the SRM tools, products, and services include SUN Microsystems (*Sun StorEdge*™), HighGround Systems, Inc. (*Storage Resource Manager*, *Media Mirror*), and Storage Computer Corp. (*Storage Administrator*) (Toigo, 2001).

STANDARDS

American National Standards Institute (ANSI)

The American National Standards Institute (ANSI) coordinates SAN voluntary standards (ANSI, 2003). The ANSI X3T10 and X3T11 working committees are associated with storage networking issues including SCSI I/O interface standards (X3T10) and fibre channel interface standards (X3T11). The first set of fibre channel standards (ANSI X.3230–1994) (Vacca, 2002, pp. 75–78) describes standards for a switch fabric (FC-SW2), the interconnect that supports high volumes of throughput and bandwidth for disk output and input, as well as a management information base (MIB) management standard that permits fibre channel devices (switches) to be managed by any vendor's software, which includes an implementation of the simple network management protocol (SNMP). SAN users can find a brief description of other X3T11 Fibre Channel standards in (Barker & Massiglia, 2001, pp. 384–386).

Several FC SAN equipment vendors, including Brocade Communications Systems, have refined SAN standards in the areas of management, discovery, data transport, and WAN connectivity. This allows the fibre channel SAN to become an integral part of the enterprise framework (Vacca, 2002, pp. 78–79).

Distributed Management Task Force (DMTF)

The Distributed Management Task Force (DMTF) has introduced management standards for computer systems and enterprise environment (DMTF, 2003). The SAN-related management standards (Barker & Massiglia, 2001, pp. 386–387) cover a set of the Web-Based Enterprise Management (WBEM) XML-based technologies. They support an object-oriented approach in developing a business's management environment, using the Common Information Model; architectures and frameworks for desktop, laptop, and server management (desktop management interface); standard data models for a network, its elements, policies, and rules (directory enable networks); and functional calls for transaction monitoring (known as the application response measurement (ARM) standard).

Storage Systems Standards Working Group (SSSWG)

As a division of the Institute of Electrical and Electronics Engineers (IEEE), the Storage Systems Standards Working Group (SSSWG) develops models and architectures of storage systems, including SANs (SSSWG, 2003). The SSSWG project authorization requests (Barker & Massiglia, 2001, pp. 387–388) include the Guide to Storage System Design; Media Management System (MMS) Architecture; Session Security, Authentication, Initialization Protocol (SSAIP) of the MMS; Media Management Protocol (MMP) for both client and administrative applications; Drive Management Protocol (DMP) of the MMS; Library Management Protocol (LMP) of the MMS; the Media Manager Interchange Protocol (MMIP) for information exchange between autonomous Media Managers; the Media Manager Control Interface Protocol (MMCIP); the C Language Procedural Interface for implementation of the MMS's components; MMS User Mount Commands for establishing "command line interfaces" (CLI); MMS standard administrative and operational commands for administering and operating an MMS; and "MOVER" specifications of a storage system data mover architecture and its interfaces.

Internet Engineering Task Force (IETF)

The Internet Engineering Task Force (IETF) organization defines a variety of the Transmission Control Protocol/Internet Protocol (TCP/IP) standards that are widely used in the enterprise environment with SANs (IEFT, 2003). The IETF standards related to storage networking (Barker & Massiglia, 2001, pp. 388) include the Simple Network Management Protocol (SNMP) for managing and monitoring devices and systems in a network; the Internet Protocol over fibre channel (IPoFC); and policy for quality of services (QoS).

STORAGE NETWORKING ASSOCIATIONS, INITIATIVES, FORUMS, AND COALITIONS

The following organizations promote storage networking technologies and products, develop standards, undertake marketing activities in information technology industry, educate, train, and create the knowledge base for implementing SAN technology:

SNIA (Storage Networking Industry Association)

The SNIA, an international association of developers of storage and networking products, is focusing on the creation of a forum of IT companies, system integrators, and application vendors for delivering architectures, education, and services in storage networking, as well as defining the specifications and infrastructures, and proposing standards for storage networking systems, including SANs, SAN attached storage (SAS), and network attached storage (NAS) (SNIA, 2003).

Fibre Channel Industry Association (FCIA)

The FCIA is an international organization of manufacturers, systems integrators, developers, systems vendors, industry professionals, and end users. In June 2002, this organization included more than 190 members and affiliates in the United States, Europe, and Japan. The FCIA is committed to delivering a broad base of fibre channel infrastructure to support a wide array of industry applications within the mass storage and IT-based arenas. FCIA working groups focus on specific aspects of the technology that target markets, which include data storage networking and SAN management. The overview of fibre channel's SAN and networking applications and examples of fibre channel solutions for high-performance networks of heterogeneous storage, server, and workstation resources can be found on the technology section of the FCIA Web site (FCIA, 2003).

Fibre Alliance (FA)

The Fibre Alliance is the networking industry consortium originally founded by a group of storage networking companies, including EMC Corporation, to develop and implement standards for managing heterogeneous fibre-channel-based SANs. In collaboration with the Internet Engineering Task Force (IETF), this group develops the definition of Simple Network Management Protocol management information bases (SNMP MIB) for storage network and device management (Fibre Alliance, 2003).

Jiro

Jiro is a Sun Microsystems technology that delivers intelligent management services for networked devices. Using the principles of Java and Jini platform-independent application development interfaces (Jini, 2003), Jiro technology provides the architecture for connecting and managing complex distributed environments such as storage area networks. The Jiro technology brings higher levels of interoperability, adaptability, and manageability to enterprise networks with storage resources (Jiro, 2002).

National Storage Industry Consortium (NSIC)

Since April 1991, the National Storage Industry Consortium has consolidated the efforts of over 50 corporations, universities, and national labs in the field of digital information storage. The corporate members are major information storage manufacturers and companies from the storage industry infrastructure, including SANs. As a nonprofit organization, NSIC supports precompetitive joint research projects, involving collaboration among users and integrators of storage systems, storage system and device manufacturers, storage component and media manufacturers, suppliers, universities and national laboratories (National Storage Industry Consortium, 2003).

SANS' MARKET, VENDORS, AND SERVICE PROVIDERS

Evolution of SANs' Market

SAN technologies allow existing enterprises to effectively manage more transactions, customers, suppliers, and services. Company operations are significantly improved by providing continuous high availability through uninterruptible access to data, increasing scalability through multiple-channel data transmission, and reducing the network and server's CPU overhead. Additional opportunities for the IT enterprises are also associated with the Internet, which allows them to increase the volume of data and rates of their transmission. According to the International Data Corporation (IDC), since the mid-1990s, the number of users of e-commerce services has increased exponentially up to several hundred millions. SANs have revolutionized the IT enterprise's infrastructure and improved its e-business applications including e-commerce, e-mail, online transaction processing, data replication, and enterprise database management. Global continuous delivery of multimedia secured information has become the main service of modern e-business enterprises.

By adding networking and intelligence features to data storage, fibre channel SAN switches enable the solution of several challenging e-business storage problems, such as linking high-performance workstation clusters, connecting high-performance tape storage on disk farms, giving server farms a high-speed data-transmission pipe, clustering disk farms, and linking Ethernet, FDDI, ATM, and Token Ring LANs to the backbone. Intelligent SAN systems allow improving enterprise performance significantly, decreasing latency, supporting direct access to the storage shared by multiple servers, reducing network traffic on the front-end network, and removing storage management tasks from servers.

In December 2000, the IDC Corp. estimated that worldwide disk storage systems sales were about $31.7 billion in 2000. Networked-based systems with SAN installations represented 20% of the 2000 revenues. Trends indicate that both SAN and NAS implementations are accelerating, as SANs have grown 70% year over year. This trend is expected to continue during the next 5 years. Networked storage on the whole experienced a 43% capital asset growth rate during the same period, with SANs growing at 33%.

SAN Vendors and Service Providers

Table 2 represents a list of SAN vendors, storage-networking service providers, and their products. The complete list of SAN deployment companies can be fount in (Vacca, 2002, pp. 495–497) and on the Network Buyers Guide Web site (Network Buyers Guide, 2003).

The main providers of SAN solutions are EMC Corporation (about 40% of the market), Compaq Computer Corporation (13%), Sun Systems, Inc. (11%), IBM Corporation (10%), Hewlett-Packard Company (7%), Dell Computer (3%), Hitachi Data Systems (2%), Brocade Communication Systems, Inc., SANgate, TrueSAN Networks, Inc., and XIOTech Corp.

Table 2 SAN Vendors and Service

Company	Product Type	Source
ADVA Optical Storage Networking Technologies	DWDM system	(ADVA, 2003)
ATTO Technology, Inc.	Fibre Channel hub	(ATTO, 2003)
Brocade Communication Systems, Inc.	Fabric switch	(Brocade, 2003)
Computer Network Technology	Storage router	(CNT, 2003)
Compaq	Disaster-tolerant SAN solutions	(Compaq, 2003)
Crossroads Systems, Inc	Modular storage router	(Crossroads, 2003)
Cutting Edge	Clustered failover	(Cutting Edge, 2003)
Dell	Enterprise storage solutions	(Dell, 2003)
EMC Corporation	Networked storage solutions	(EMC, 2003)
Emulex Corporation	VI/IP PCI host bus adapter	(Emulex, 2003)
Hewlett-Packard	SAN management tools	(Hewlett-Packard, 2003)
IBM	Enterprise storage server	(IBM, 2003)
LSI Logic Storage Systems, Inc.	Storage manager software	(LSI, 2003)
McData	Enterprise FC management tools	(McData, 2003)
Media Integration	Fibre channel SAN; open storage	(Media Integration, 2003)
Nishan Systems	IP storage switches	(Nishan Systems, 2003)
Procom	Storage systems	(Procom, 2003)
Storage Computer Corp.	Storage systems	(Storage Computer, 2003)
StorageTek	Enterprise fibre channel switch	(StorageTek, 2003)
SUN Microsystems	Storage systems	(Sun Microsystems, 2003)
U.S. Design Corporation	RAID storage hardware	(US Design, 2003)

CONCLUSION

SANs, networked high-speed infrastructures, enable e-business enterprises to improve significantly their 24×7 continuous scalable services. They have become a critical part of the enterprise network infrastructure. The above-considered technologies and effective SAN solutions allow companies to shift their focus from numerous IT infrastructure problems to the successful performance of their businesses and services.

GLOSSARY

CIFS Common Internet file system, also known as the Microsoft server message block protocol; a network file system access protocol that is primarily used by Windows clients to communicate file access requests to Windows servers.

CIM Common information model; an object-oriented description of entities and relationships in the enterprise management environment.

DWDM Dense wavelength division multiplexing; a method that allows more wavelengths to use the same fiber.

FC-AL Fibre channel-arbitrated loop transport protocol.

FCSW Fibre channel switch.

FC-VIA Fibre channel–virtual interface architecture.

FSPF Fabric shortest path first; a routing protocol used by fibre channel switches.

IPoFC Internet protocol over fibre channel.

iSCSI Internet small computer systems interface.

JBOD Just a bunch of disks; a term for a collection of disks configured as an arbitrated loop segment in a single chassis.

NAS Network attached storage.

RAID Redundant arrays of inexpensive disks; a technology for managing multiple disks.

SAN Storage area network.

SAS SAN attached storage.

SoIP Storage over IP; a storage technology developed by the Nishan Systems Corporation.

SSA Serial storage architecture.

VI Virtual interface architecture; a midlayer protocol specification.

CROSS REFERENCES

See *Conducted Communications Media; Standards and Protocols in Data Communications; TCP/IP Suite.*

REFERENCES

ADVA optical storage networking technologies (2003). Retrieved March 13, 2003 from http://www.san.com/

American National Standards Institute (ANSI) (2003). Retrieved March 13, 2003 from http://www.ansi.org/

ATTO SAN solutions (2003). Retrieved March 13, 2003 from http://www.attotech.com/sans.html

Barker, R., & Massiglia, P. (2001). *Storage networking essentials: A complete guide to understanding & implementing SANs.* New York: Wiley.

Beauchamp, C., Judd J., & Kuo, B. (2002). *Building SANs with Brocade fiber channel fabric switches.* Rockland, MA: Syngress Publishing, Inc.

Brocade Communication Systems, Inc. (see the SAN Info Center) (2003). Retrieved March 13, 2003 from http://www.brocade.com/

Clark, T. (1999). *Designing storage area networks: A practical reference for implementing fiber channel SANs*. Boston: Addison–Wesley.

Clark, T. (2002). *IP SANS: An introduction to iSCSI, iFCP, and FCIP protocols for storage area networks*. Boston: Addison–Wesley.

CNT (Computer Network Technology Corp.) storage over IP solutions (2003). Retrieved March 13, 2003 from http://www.cnt.com/

Compaq, Inc. (2003). Retrieved March 13, 2003 from http://www.compaq.com/

Crossroads Systems, Inc. (2003). Retrieved March 13, 2003 from http://www.crossroads.com/about/

Cutting Edge, Inc. (2003). Retrieved March 13, 2003 from http://www.cuttedge.com/

Dell, Inc. (2003). Retrieved March 13, 2003 from http://www.dell.com/

Distributed Management Task Force (DMTF) (2003). Retrieved March 13, 2003 from http://www.dmtf.org/

EMC Corporation (2003). Retrieved March 13, 2003 from http://www.emc.com/

Emulex Corporation (2003). Retrieved March 13, 2003 from http://www.emulex.com/

Farley, M. (2001). *Building storage networks* (2nd ed.). New York: Osborne/McGraw–Hill.

FCIA (Fiber Channel Industry Association) (2003). Retrieved March 13, 2003 from http://www.fibrechannel.org/

Fibre Alliance (2003). Retrieved March 13, 2003 from http://www.fibrealliance.org/

Hammond-Doel, T. (2001). *2 Gb/s fiber channel SANs*. Vixel Corporation. Retrieved March 13, 2003 from http://www.vixel.com/9000_docs/9000_wp.pdf

Hewlett-Packard, Inc. (2003). Retrieved March 13, 2003 from http://www.hp.com/

IBM Enterprise SAN Solutions (2003). Retrieved March 13, 2003 from http://www.storage.ibm.com/ibmsan/

IEEE SSSWG (Storage Systems Standards Work Group) (2003). Retrieved March 13, 2003 from http://www.ssswg.org/

InfiniBand Trade Association (2003). Retrieved March 13, 2003 from http://www.infinibandta.org/home/

InfraStor Technology Corp. (2001). *SAN vs NAS*. Retrieved March 13, 2003 from http://www.infrastor.com/tech/sanvsnas.htm

InfraStor Technology Corp. (2002). *Introduction to SAN*. Retrieved March 13, 2003 from http://www.infrastor.com/tech/SANTechIntro.htm

Intel InfiniBand Architecture (2003). Retrieved March 13, 2003 from http://www.intel.com/technology/infiniband/

Internet Engineering Task Force (IETF) (2003). Retrieved March 13, 2003 from http://www.ietf.org/

Jini (2003). Retrieved March 13, 2003 from http://www.jini.org/

Jiro (2002). Retrieved September 12, 2002 from http://www.sun.com/jiro/

LSI Logic Storage Systems, Inc. (2003). Retrieved March 13, 2003 from http://www.lsilogicstorage.com/

McData, Inc. (2003). Retrieved March 13, 2003 from http://www.mcdata.com/

Media Integration, Inc. (2003). Retrieved March 13, 2003 from http://www.mediainc.com/

National Storage Industry Consortium (2003). Retrieved March 13, 2003 from http://www.nsic.org/

Network Buyers Guide (2003). Retrieved March 13, 2003 from http://networkbuyersguide.com/

Nishan Systems (2003). Retrieved March 13, 2003 from http://www.nishansystems.com/

Ottem, E. (2001). *Third generation SANs: Open 2 Gb fabric* (White Paper WP-2G0801). Gadzoox Networks, Inc. Retrieved March 13, 2003 from http://www.infrastor.com/downloads/2Gb_whitepaper.pdf

Peterson, M. (1998). *Storage area networking*. Santa Barbara, CA: Strategic Research Corp. Retrieved March 13, 2003 from http://www.sresearch.com/wp_9801.htm

Procom Technology Corporation (2003). Retrieved March 13, 2003 from http://www.procom.com/

Sachdev, P., & Arunkundram, R. S. (2002). *Using storage area networks*. Special edition. Indianapolis, IN: Que.

Sheldon, T. (2001). *McGraw–Hill encyclopedia of networking & telecommunications*. New York: McGraw–Hill.

SNIA (Storage Networking Industry Association) (2003). Retrieved March 13, 2003 from http://www.snia.org/

Storage Computer Corporation (2003). Retrieved March 13, 2003 from http://www.storage.com/

StorageTek, Inc. (2003). Retrieved March 13, 2003 from http://www.storagetek.com/

Sun Microsystems (2003). *Data storage solutions*. Retrieved March 13, 2003 from http://www.sun.com/storage/

Thornburgh, R. H., & Schoenborn, B. J. (2001). *Storage area networks: Designing and implementing a mass storage system*. Upper Saddle River, NJ: Prentice Hall PTR.

Toigo, J. W. (2001). *The Holy Grail of data storage management*. Upper Saddle River, NJ: Prentice Hall PTR.

U.S. Design Corporation, Inc. (2003). Retrieved March 13, 2003 from http://www.usdesign.com/

Vacca, J. (2002). *The essential guide to storage area networks*. Upper Saddle River, NJ: Prentice Hall.