

Storage Area Network Overview

Bob Walder
The NSS Group

Today's organizations are demanding high-speed data protection for complex and resource-intensive data such as image, audio, and video files—as well as for large databases. In addition, the amount of data stored on distributed servers is increasing constantly, while backup windows continue to shrink. This trend has resulted in complex multiple server environments with low scalability, high administrative costs, and insufficient protection. To overcome these issues, organizations need a storage management solution that provides:

- Fast backup speeds that match the allotted backup window without adversely affecting network band-width.
- Flexibility to connect and share remote devices and servers.
- Improved scalability to expand the storage infrastructure without rebuilding.
- Interoperability between disparate systems.
- Centralized management to lower the overall total cost of ownership.

In addition to data protection, organizations wishing to maintain their competitive edge require the ability to quickly restore or recover critical information ranging from customer data to internal operations. This task has become extremely difficult due to the sheer amount of data spread across WANs and distributed heterogeneous systems, while contending with bandwidth saturation and usability issues. And while in the past it may have been acceptable to have servers "down" for a period of time, this is no longer the case. Today, system downtime often results in loss of business, decreased market share, and possible disaster for an organization—regardless of its size or industry.

Many enterprises have taken a distributed approach to storage management. Backups are either performed over the LAN, where several systems are backed up to a central storage device, or locally, where a system has a backup device directly connected to it. The SCSI bus, while a mainstay of storage connectivity for over two decades, still has some limitations. Ultra Wide SCSI today only delivers up to 40Mbps and sustains up to 15 devices on the chain.

What Is the Storage Area Network?

The LAN world is about to go through another revolution in terms of storage. This revolution—known as a Storage Area Network (SAN)—involves moving network storage from its traditional location—inside, or directly connected to, file servers—to a separate network of its own. Disk, tape, and optical storage can then be attached directly to this network, which is based on a "fabric" of fibre, switches and hubs that connects storage devices to a heterogeneous set of servers on a many-to-many basis.

A SAN is thus a dedicated storage network that carries I/O traffic only between servers and storage devices—it does not carry any application traffic, which eliminates the bottlenecks associated with using a single network fabric for all applications. A SAN can also enable direct storage-to-storage interconnectivity, and lends itself to the exploitation of new breeds of clustering technology and to getting the best out of Network Attached Storage devices that can intelligently provide disk and tape capabilities to one or more servers.

Fibre Channel—for so long a technology with no applications—is the critical enabler for the SAN. SANs utilize high-speed fibre optic or copper cabling to interconnect between server and storage

devices, resulting in data transfer speeds of up to 200 Mbps in a dual loop configuration or 100 Mbps in redundant mode.

Fibre channel also supports multiple servers and enables device sharing between servers on the loop. Fibre optic bus lengths can reach 10 kilometers (or 6.25 miles) without the use of extenders or switched fabric technology (switched fibre channel SANs connected to each other). Furthermore, SANs are capable of supporting and mapping SCSI, HIPPI, IP, ATM, and other network and channel protocols.

The following table illustrates the key benefits of Fibre Channel over the more traditional storage model using SCSI:

	Ultra Wide SCSI	Fibre Channel SAN
Data Transfer Rate	40 Mbps	100 Mbs
Scalability	15 Devices	126 (FC-AL), Virtually Unlimited (Switched)
Max. Length	10 feet, inflexible cable	6.25 miles, easy to interconnect
Hot Swap Support	No	Yes
Manageability	Server-dedicated device	Load balancing multiple servers across multiple devices
Connectivity	Costly reconfiguration required	Hot swap new devices into hub/switch
Availability	None	Easily redirect job to another server on the loop

Fibre Channel has relieved the connectivity and bandwidth limitations associated with SCSI and allowed SANs to be implemented today for large-scale storage sharing, since it provides the ability to transmit data at very high speeds over long distances. Fibre-enabled servers, disk arrays and other intelligent storage devices are connected to the "fabric" by fibre through sophisticated switches and hubs.

The FC-AL (Fibre Channel Arbitrated Loop) configuration uses a hub to connect the servers to the storage devices, and the hubs arbitrate the signals from any one server to a storage device, thereby disallowing simultaneous conversations across its ports. The Switched Fabric SAN, on the other hand, utilizes high speed, low latency micro-switches, allowing simultaneous conversations across all ports.

Switched Fabric thus enables better throughput and forms the basic building block for fibre channel fabrics, thereby allowing virtually unlimited scalability. However, the price per port is typically much higher than FC-AL. As a result, an organization looking to adopt a fibre channel solution must weigh cost against scalability requirements.

While FC-ALs, in theory, can support up to 126 devices while sustaining their transfer rate, Switched Fabric can support almost unlimited devices. Both of these fibre channel interfaces also support hot swapping, allowing administrators to plug in additional servers and/or storage devices without bringing the loop or the servers down. Clearly, businesses that require connectivity over

great distances, high speeds, and a large number of devices on the bus should strongly consider a fibre channel interface.

Fibre channel SANs offer the benefit of centralized backup, device management from multiple servers, and management of these multiple storage devices. In addition, centralized management helps in isolating, identifying, diagnosing, and recovering from load management problems all from either a centrally managed console or any server on the loop. This powerful solution also offers improved fault tolerance.

Storage and Data Sharing

Just as with data migration, our external storage devices can be categorized as primary or secondary storage. Primary storage devices are usually the fastest, random access devices such as individual disks or RAID arrays, whereas secondary storage is usually a linear access device such as a tape, or a device with slow access characteristics such as an optical drive.

Because secondary storage device capacity is higher and media costs are much lower, a SAN is suitable for data archives and second/third-level data stores in data migration applications. Both primary and secondary storage devices can coexist on the same SAN, and the SAN provides both storage and data sharing capabilities in an attempt to maximize the use of primary and secondary storage devices.

Storage sharing—or storage consolidation—enables multiple computers across a corporate network to access a common set of storage devices such as disk arrays, tapes, optical drives and autoloaders. Think back to the early days of the LAN when we were sold the idea on the back of the promise of sharing expensive resources across multiple users on the network. Those resources—big disks, tapes, printers—were installed on the central file server, from where they could be accessed by everyone with the appropriate authorization.

SAN storage sharing simply introduces another level of abstraction. Now, those same resources (only much bigger) are moved out of the server and attached directly to the network, thus allowing them to be addressed directly by multiple servers. For instance, if your backup software expects to find a tape drive in the local server, you would normally have to install a drive in every server on your LAN. Now, fewer tape drives can be installed in a central array and attached to the SAN, making them accessible to every server on the network

With storage sharing, exclusive access is provided while a device is assigned. Primary storage may be assigned to a computer for a long time because the data and applications on the storage become integral parts of the computer. Secondary storage, like tape drives, may be assigned to a computer for much shorter periods of time, often only as long as is necessary to back up the computer's data files.

This exclusive access is important to preserve data integrity, especially when the same disk device is shared between two completely different operating systems. Allowing both systems to access the disk simultaneously could cause unpredictable—even disastrous—results. The SAN must therefore be capable of hiding mount points from end users, and even preventing the OS itself from recognizing the presence of locked devices on the SAN.

The SAN also improves the concept of **data sharing**. Although a typical LAN enables applications and end users to access data held in a central location, the SAN moves that data onto a much faster infrastructure. This allows multiple computers to transfer large files

concurrently at rates comparable to locally attached disks over the SAN without adversely affecting the corporate LAN.

Usually, of course, it is the host operating system that controls access to local hard drives, and can thus preside over access privileges and file locks when more than one application attempts to use the file at the same time. Once the disk storage is removed physically from the server, however, the SAN itself must take over and secure access to files through a volume lock manager or distributed file system software.

Data sharing requires that the participating computers be able to find and use the contents of a file. Hence, computers with different operating systems must use protocol translation modules and other software to establish a common communication dialect.

Ordinarily, data sharing is associated with primary storage devices. But it can be done with secondary storage. Tape devices are linearly accessible file systems managed by the backup software. Robotic tape libraries containing multiple tape devices and a large quantity of media can be accessed by more than one computer using SANs, which makes it possible to share data on secondary storage devices.

The Serverless Backup

The dominant storage interconnects today are the LAN (for remote server to backup server connections) and SCSI (for server to storage device connections).

While the LAN covers long distances, it also exhibits high latency, making it unsuitable for high-volume data transfer. In most cases the LAN, combined with copious amounts of data from various sources, creates a bottleneck that causes severe reductions in already tight backup windows. And although the majority of organizations today use a 100 Mbps or faster LAN, even these are proving to be inadequate to carry the burden of an organization's data communication and storage needs. Local versus remote backup and restore capabilities can make a tremendous difference in terms of suitability. Localized backup results in faster speeds than remote because it does not have to contend with communication traffic on the LAN.

But while instituting a local backup solution might seem the obvious answer, it is impractical from both a logistical and economic point of view to, for example, simply attach a tape drive or library to each application or database server. This type of fragmented solution presents an administrative nightmare, with limited reporting, management, access, and control for each distributed server. Furthermore, the SCSI bus—the thick, awkward cable that provides server to storage connectivity—has numerous limitations, including its length, the number of devices to which it can attach, and its bandwidth and burst rate.

SANs uncouple the front end of the IT infrastructure—applications, operating system, and processors—from the back-end storage. This enables businesses to meet their expanding storage requirements while still maintaining rapid response at the application, business process, and user level.

However, SANs don't completely free servers from their back-end tasks. Businesses not only store information they also have to back it up. Business servers still have to execute backup functions, meaning they have to read data from the storage device and write it to the backup device. When you consider the enormous, multi-terabyte databases of major enterprises and eCommerce Web sites, it is not hard to see how this can compromise server performance and hence SAN performance.

Serverless backup is a simple, elegant solution to the problem of how to safeguard massive and exponentially growing amounts of data without compromising system performance or limiting the bandwidth available for enterprise communications. It is therefore important that the backup system can support the SAN directly, otherwise it could force data to travel from SAN-based storage via a file server on the LAN just to enable the backup software to write the data to tape. In such cases, the file server would act as a bottleneck, slowing down the backup process and once again threatening the backup window.

Serverless backup over SANs requires three major components:

- The backup application itself
- The Extended SCSI Copy Command standard
- A protocol-aware, intelligent SAN appliance that can recognize protocols from many heterogeneous systems and transmit data at high speeds to the tape and DLT libraries.

With serverless backup, the data flows across the SAN directly from the disk drive to the tape device, with no data moving through the server. The enterprise servers only need to host the backup application, and the backup application determines what needs to be backed up and sends the command to a "copy agent" embedded in an intelligent SAN appliance. The intelligent SAN appliance detects the source and destination parameters, retrieves the data from the storage devices, writes it to the tape or DLT libraries, and reports completion (or status) back to the backup application.

With backup now handled by the SAN devices, enterprise servers can continuously process applications and information, and not concern themselves with "housekeeping" tasks like data movement. Routine backup can be performed regularly during peak business times rather than during off-peak backup windows, or backup can be performed offline altogether with zero impact on the enterprise.

In addition to off-loading the LAN, the backup software provides storage sharing capabilities, so multiple computers can use a set of tape devices. Advanced media and file management capabilities should also be included in backup software packages, so secondary storage data sharing is available.

In a conventional storage architecture, the storage sub-system is accessible only to the server or the platform the sub-system is attached to. As storage requirements grow in the environment, the administrator is forced to reallocate storage capacity, denying accessibility of those resources to other platforms. The answer to these challenges is to share resources between multiple platforms. SAN connectivity enables resource sharing between multiple backup servers, enables administrators to consolidate backups to one storage sub-system. This simplifies management and enables efficient use of storage capacity.

The SAN and Data Availability

The ultimate goal of maintaining data availability on key servers requires a proactive solution rather than a remedial one. In other words, no matter how good the backup solution employed, there is often an unacceptable delay associated with restoring data following a catastrophic server failure.

SANs provide continuous client availability to storage devices if a server in the loop fails, and some backup solutions on the market today are capable of replicating data and application files in real time to secondary servers on the SAN. These solutions provide continuous access to data

even if the primary server suffers fatal damage or network connections are interrupted. When it detects an interruption, the backup solution can instantaneously and transparently switch users to a secondary server.

This same replication capability can also be used as a tool to remotely mirror data to an alternative site at local, metropolitan, or worldwide locations, providing further levels of data protection and redundancy.

SAN Management

Centralized management of all physical and logical storage resources via a single console is paramount as the size and complexity of today's network grows.

These resources include logical resources such as file systems and application specific storage repositories and physical resources that include RAID systems, tape libraries and fibre channel SAN components. Centralized management solutions must include the ability to automatically detect these resources, and correlate and analyze their capacity, configuration and performance information. These management solutions should also enable consistent policy administration across platforms and storage technologies.

Several storage devices, including servers, disks and tape devices, constitute the SAN environment. Storage Resource Management (SRM) defines applications that monitor and manage physical and logical resources. SRM includes capacity management, configuration management, event and alert management, volume management and policy management. Effective SAN management necessitates SRM tools be integrated with the SAN management solutions.

The ideal enterprise management tool includes the SAN as a storage network topology and as a sub-network to the communication LAN or WAN. The SAN is not isolated, and SAN solutions should not be built in isolation from other IT disciplines. Management solutions should offer SAN discovery topology mapping and end-to-end management for fibre channel devices in the loop.

While the tools for centralized management can be fairly expensive, the cost savings realized by improved reliability and ease of management more than offset the infrastructure costs.

SAN Summary

IT environments today are plagued with small backup windows, overburdened LANs, databases that increase in size daily, and high availability requirements for mission-critical applications.

To further complicate matters, the administrator is besieged with managing volumes of data and the everyday obstacles to effective storage management, such as viruses, hardware failure, faulty tapes, and more. SANs represent a huge stride toward a cost-effective solution, providing increased performance, fault tolerance, and scalability for long-term growth.

In addition, SANs provide total cost of ownership benefits such as:

- Minimized down time
- Improved LAN performance
- Ability to connect to existing LANs

- Reduced administrative costs
- Leveraging of existing hardware
- Improved fault tolerance
- Maximized storage resources through load balancing
- Total SAN management.

SAN awareness today is becoming widespread. As more and more organizations integrate SANs into their environments, they need reassurance that their current storage management solution takes full advantage of both existing storage technology and SAN technology.

Open standards are thus increasingly important to allow a single solution across multi-vendor heterogeneous networks.



About the Author

Bob Walder, a leading authority on network security, is one of the founders of The NSS Group. Since leaving behind the world of IT management in 1991, Bob has been at the cutting edge of new technology and has invested much of his time in advising on, testing and certifying a range of security products on behalf of end user organisations, vendors and certification bodies. He is also a regular contributor of technical articles to the major networking and security titles.

The NSS Group is Europe's foremost independent network and security testing facility. With labs in Cambridge in the UK and Carcasonne in the south of France, The NSS Group offers a range of specialist networking and security services to vendors and end user organisations throughout Europe and the United States. For more information, visit www.nss.co.uk or e-mail info@nss.co.uk.