Windows Servers in a Storage Area Network Environment

White Paper

Published: April 2004

Abstract

Storage area network (SAN) technology, once restricted to large organizations and mainframe environments, is becoming more common in mid-sized organizations that need to store and share access to large amounts of data. With the development of Internet SCSI (iSCSI) technologies, the benefits of SAN technologies are likely to extend even to small organizations.

This white paper describes how SANs are used with Windows products such as Microsoft® Windows Server™ 2003 and Microsoft® Windows® Storage Server 2003, as well as with Microsoft® Exchange Server 2003 and Microsoft® SQL Server™ 2000. While the Windows® operating system has been used in SAN environments since the introduction of Microsoft® Windows Server NT® 4.0, the release of Windows Server 2003 has dramatically enhanced SAN interoperability. That commitment to SAN interoperability will continue with future Microsoft releases.

Contents

Introduction

Storage area networks (SANs) are high performance networks dedicated to delivering block (unformatted) data between servers and storage. SANs originated in the 1980s as a mainframe solution designed to accommodate the then specialized need to store huge amounts of data in a manner that was scalable, flexible and highly available. Originally restricted to high end single vendor solutions, by 2003, with the rapid growth of the Internet and increased reliance on e-commerce, adoption of SANs (often in multi-vendor configurations) has become more common. According to IDC[1],the global market research firm for IT industries, SAN adoption has grown from a market share of a little over 20% in 2000 to nearly 45% in 2002. This growth has been at the expense of direct attached storage (DAS), which has declined in market share from 66% to 45% over the same time period. While no longer solely a high end solution (although the most advanced and fully developed SAN solutions are still to be found in mainframe environments), SAN adoption is predominantly found in organizations that need to store a terabyte or more of data[2].

As the need for more flexible and scalable storage solutions extended beyond the mainframe market to the open server market, it became clear that operating systems designed with the expectation that storage is local to the server would need to be extensively re-architected to allow customers to realize the SAN benefits of shared storage.

In the Windows environment, a number of changes were made to enable the operating system to function more effectively in a SAN environment. The first step in this process was to enable SAN connectivity. This advance was put into effect with the release of Windows NT Server 4.0. Since that release, Windows architecture has undergone continued development to enable more effective SAN use and performance. The improvements are the subject of this paper.

As a direct result of these advances, Windows use in the SAN market has dramatically increased. In Fibre Channel SAN environments, Windows adoption has grown from 26% of the market in 2000 to more than 40% by the end of 2002[3]. Most of that growth has happened in the lower price bands[4], and many analysts expect Microsoft support for iSCSI (which enables SANs using existing IP infrastructure) to continue to help drive SAN adoption over the next five years[5], particularly at the less expensive price ranges.

---

[1] IDC Market Analysis (2003): Worldwide Disk Storage Systems Forecast and Analysis, 2003-2007, Table 5.

[2] In organizations storing more than a terabyte of data, adoption is around 40%. It is less than 10% in those organizations that store less than a TB (PriMetrica survey, cited in InfoStor January 2003).

[3] IDC 2003, Table 39.

[4] IDC 2003, Table 41.

[5] IDC 2003, Table 77.

Storage Area Networks

Both direct attached storage and storage area networks use the SCSI protocol to move data in blocks, rather than files[6]. In fact, from the vantage point of most operating systems, DAS and SAN storage are indistinguishable, despite the differences in their network topologies (see Figures 1 and 2). So what do storage area networks offer that traditional direct attached storage does not?

SANs are designed to enable centralization of storage resources, while at the same time overcoming the distance and connectivity limitations posed by directly attached storage. Parallel SCSI interconnects limit direct attached storage devices to a distance of 25 meters, and can connect a maximum of only 16 devices. Fibre Channel SANs extend the distance limitation to 10 kilometers or more, and enable an essentially unlimited number of devices to attach to the network. These factors allow SANs to effectively uncouple storage from the server and to pool on a network where it can be shared and easily provisioned, without the problems of scaling associated with DAS.

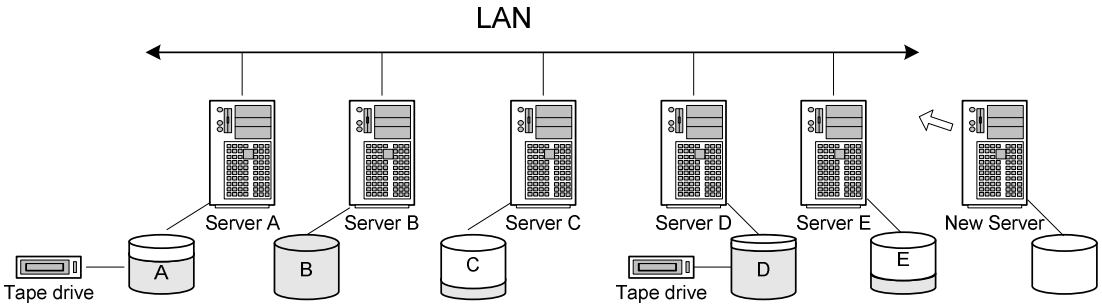

Figure 1. Direct Attached Storage

Figure 1 shows storage directly attached to servers on the LAN. (Shading indicates in-use storage resources). Despite that fact that servers C and E have excess storage, there is no easy way for these resources to be redeployed to either server B or D; the applications on server B will fail since they can no longer write to the full disk attached to server B. Once a server has no more room for additional storage, the most common way to add storage resources is to add a new server. The disadvantages of this approach are increased capital expenditures and greater management complexity (especially for backups).

[6] NAS, or network attached storage, differs from SANs by serving data in file format rather than as blocks.

Through shifting from a server-centric storage model to a network centric one (Figure 2), SANs facilitate on-demand resource provisioning. Since all servers have access to the same storage pool, accommodating peak storage needs is a matter of shifting resources to servers on an as-needed basis, rather than systematically overbuying storage resources for each server. SANs also facilitate the sharing of maximally up-to-date data, equipment consolidation (including shifting from discrete tape drives to shared tape libraries), effective clustering solutions, high performance I/O, and a reduction in network traffic. The net results of deploying a storage area network are more efficient storage resource management, better data protection, and improved performance.

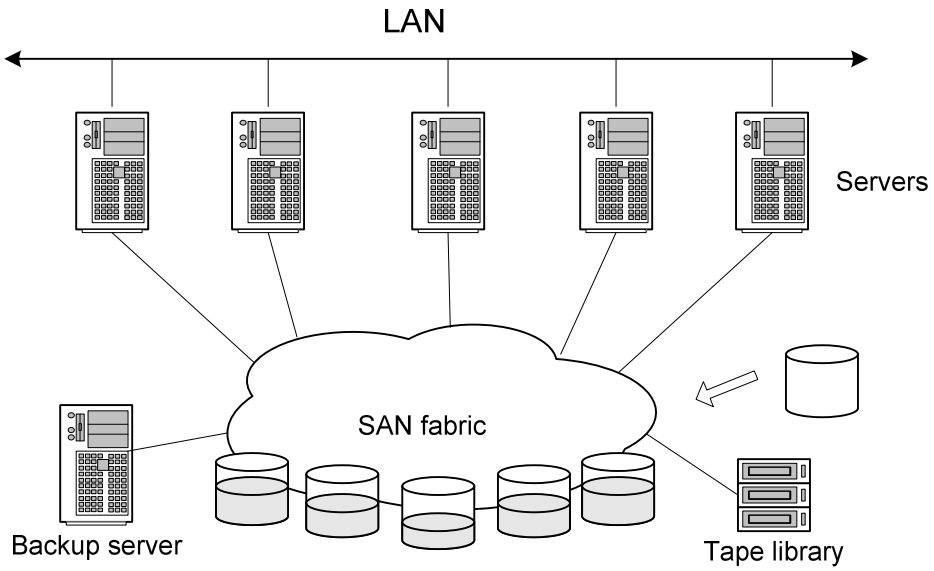

Figure 2. Storage Area Network

The benefits of SANs, however, are not free. Much of the cost lies in the network components of the SAN solution. Given the distance and connectivity limitations of the parallel SCSI bus—and given the fact that IP networks, although ubiquitous, were not originally designed to carry SCSI command data blocks (the raw I/O without file formatting)—a new interconnect solution was required before the goals of centralized storage could be realized.

Fibre Channel SANs

Fibre Channel interconnects delivers high performance block I/O to storage devices. The most widely adopted SAN interconnect technology, Fibre Channel, is based on serial SCSI technologies and overcomes the parallel SCSI limitations to enable essentially unlimited device connectivity over long distances. And, unlike parallel SCSI devices which must arbitrate (or contend) for the bus, FC channel devices, using switch technology, can transmit information between multiple servers and multiple storage devices at the same time.

However, the capital investment and the expertise required to install Fibre Channel networks components (Host Bus Adapters (HBAs), cabling and switches) can be substantial. (HBAs, like network adapter cards, install in the server and function to transmit block data.) And redundant configurations, while ensuring continuously available data, add considerably more cost and complexity.

iSCSI SANs

Internet SCSI (iSCSI) SANs make use an existing IP knowledge base, leveraging much of the existing in-place network hardware, software, and expertise. They also make it simpler and less expensive to set up an iSCSI network to transport block data between servers and storage than it would be to install a Fibre Channel network. Since the iSCSI standard did not receive final ratification until 2003, its technologies are just now beginning to be implemented. iSCSI technologies offer a potentially lower-cost alternative to Fibre Channel SANs, and they may help bring SAN benefits to more organizations than was previously possible. For both security and performance, a dedicated iSCSI network separating network traffic from storage traffic is recommended.

The current down side to iSCSI delivery of block storage is that I/O (input/output) transfer is slower than that which Fibre Channel can deliver. iSCSI performance can be improved by using Gigabit Ethernet cabling and switches, and iSCSI chips or host bus adapters on which to offload the server TCP/IP processing overhead. Finally, for both security and performance, a dedicated iSCSI network separating network traffic from storage traffic is recommended.

Making a SAN Perform

Separating storage onto a dedicated network—whether Fibre Channel or iSCSI—offers greater disk management efficiencies, more effective system and resource management, improved network performance, higher availability of stored resources, and more cost-effective data protection than direct attached storage can offer.

Effective fabric management is the framework upon which effective SAN solutions are built. All the networked storage resources—the hosts, interconnects, and storage subsystems—must come together as a complete system to deliver a high performance storage solution. The software controlling SAN functionality (of both the fabric and storage array) will reside on the server, the switch and/or the storage array. Such software must:

- Be able to detect when storage resources are added or removed from the storage network, even in multivendor storage configurations
- Ensure that servers effectively share storage resources on a SAN, without interfering with ownership of storage resources, even in multiplatform environments
- Enable the detection and management of SAN components (such as links to the storage array), ensure they are functional, and if not, be able to fail over to an alternate component with no loss of data

- Help to ensure that SANs are secure
- Maximize throughput using high performance interconnects to and from the storage array
- Help realize solutions that support high availability and disaster recovery goals
- Enable administrators to configure storage resources on the SAN in a way that realizes the benefits of a centralized and consolidated storage model

There are two approaches to obtaining an integrated SAN solution. One is to purchase a single vendor end-to-end solution. Such solutions have been available for Fibre Channel SANs for several years[7]. Single vendor solutions have the advantage of being complete; the disadvantages are a lack of flexibility in adding to the solution and potentially high costs.

Alternatively, many organizations are installing multivendor Fibre Channel SAN solutions, which together can deliver all the same SAN benefits for lower cost. While these solutions can create significant interoperability challenges, the challenges are diminishing as standards for ensuring interoperability emerge. In contrast, one of the significant advantages of iSCSI SANs is that interoperability issues are much less problematic.

The next two sections outline the basics of SAN topology for those unfamiliar with SAN technologies.

Basic SAN Configuration

The simplest SAN configuration, whether Fibre Channel or iSCSI, is to deploy a minimum of two production servers attaching to a storage array.

In a Fibre Channel SAN, shown in Figure 3, each server contains an HBA which connects by means of a Fibre Channel switch to a disk controller on the storage array. HBAs, although they reside on the server, are also part of the storage network. They serve first to provide the interface between the server and the attached Fibre Channel network, and second to provide I/O processing, offloading most of the server processing required for transferring data. The resulting performance is very high and very scaleable.

[7] End to end iSCSI solutions, because the technology is so new, are not yet commonplace.

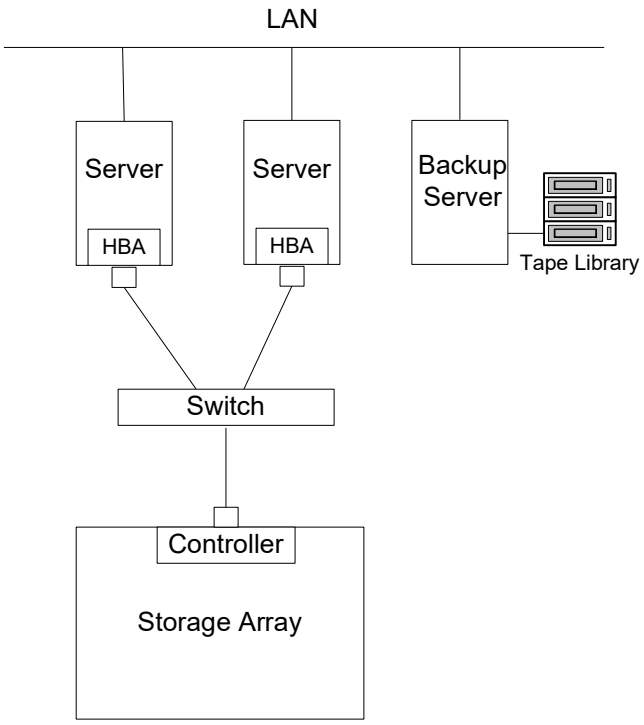

Figure 3. Basic SAN Configuration, No Redundancy

The Fibre Channel switch is an intelligent device that collects information about the SAN network topology and attached devices. Switches

- Enable simultaneous communication between multiple endpoints (such as a server and storage disk)
- Provide capabilities for managing those connections between devices so that access is controlled (through zoning)
- Provide physical connectivity for failover and load balancing

Although the simple configuration shown in Figure 3 provides basic connectivity to shared storage, it does not provide any protection from single point of failure problems, such as loss of a switch, controller or interconnect cable. For continual access to storage from each server, redundant components must be added: dual HBAs, switches and disk controllers. Once dual components are added in, effective fabric management becomes critical to SAN functioning.

Note that an iSCSI SAN has the same basic setup as that shown in Figure 3, except that a Gigabit Ethernet adapter replaces the HBA, and the connection to the disk controller on the iSCSI storage target is through an Ethernet switch.

Advanced SAN Configuration

More advanced SAN configurations may support dozens or hundreds of servers. These configurations must be designed with redundant components—controllers, switches, ports, HBAs and cabling—to ensure that there is no single point of failure. An example of a redundant configuration is shown in Figure 4. Note that there is no cable connection between the two switches. This ensures that the two fabrics remain independent.

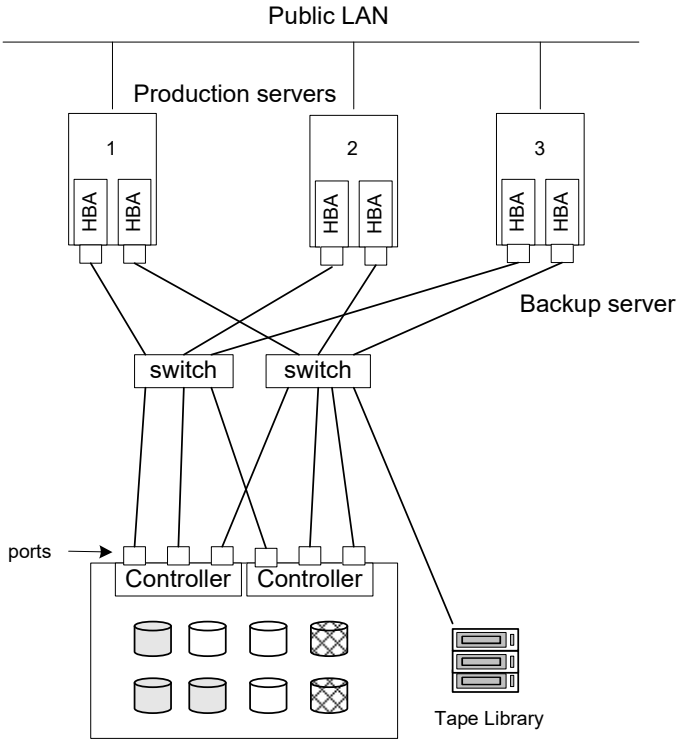

Figure 4. Advanced SAN Configuration, Showing Redundancy

Controlling Device Access

To effectively manage storage resources among servers, access to those resources must be restricted in such a way that servers do not see one another's storage. Managing storage access can be controlled by the switch (through zoning), and/or by the HBA or storage device controller[8] (through logical unit number (LUN) masking).

[8] Although it is possible to do LUN masking using the HBA, Microsoft recommends controller-based LUN masking for use with the Windows platform.

In switch based zoning, the switch keeps a list of either the port addresses or the World Wide Names that are allowed to communicate with one another; these ports are members of the same zone. If ports outside the zone attempt to access resources, they will be denied. As fabrics grow, more switches and zones can be added to increase both connectivity and aggregated bandwidth[9].

Controller-based LUN masking restricts access to the storage devices themselves. In Figure 4, shading indicates related logical units of storage in the array: thus for example, the cross-hatched disks together form a single LUN. By masking specific LUNs to all but a single server, other servers are prevented from accessing LUNs that do not belong to it.

A Note on Fibre Channel Arbitrated Loop (FC-AL) Configurations

Both of the configurations mentioned here employ switches. Using a hub rather than a switch is a simple and inexpensive alternative; in fact, hub-based FC-AL topology was the initial interconnect used when Fibre Channel SANs were first deployed. However, for a number of reasons, the industry has moved away from using hubs, and they are now rarely seen in configurations where the goals are high performance, reliability, and maximum realization of SAN benefits.

[9] Because dynamic zoning can be disruptive, Microsoft only recommends static zoning.

Windows Servers in a SAN Environment

Before the release of Windows NT Server 4.0, organizations running applications on the Windows platform could only accommodate their rapid data growth by adding more servers with direct attached storage (whether internal or external to the server); any excess storage capacity could not be shared among servers. While for many organizations, this scale-out model was, and still is, an effective solution, for others, such upgrades are a temporary, time consuming and costly fix that limit the usefulness of deploying the Windows platform. This limitation was eliminated after Windows servers could reliably connect to a SAN.

Platform Evolution

As SANs were more widely adopted, the Windows platform developed increasingly sophisticated support for integrating with storage area networks (see Figure 5):

- **Windows NT 4: Basic SAN connectivity**. For organizations running I/O intensive applications such as SQL or Exchange, the basic SAN connectivity of Windows NT, enabled through third party drivers, meant that organizations could take advantage of the SAN consolidated and shared storage model in the Windows environment.
- **Windows 2000 Server: Scaling Capacity**. With the release of Windows 2000 Server, the Windows platform became competitive with alternative operating systems for organizations with large data centers or online transaction processing needs. With 32 processors and 64 GB RAM, Datacenter editions supported scaling up to more powerful solutions; 32-node network load balancing and 8 node failover clustering supported more sophisticated scale-out solutions.

  In terms of SAN functionality, in-box drivers gave limited Fibre Channel capabilities to the Windows platform. While this was a step forward, the platform continued to require skilled and careful configuration. Vendor drivers continued to develop capabilities, such as fabric support and more advanced manageability functions.

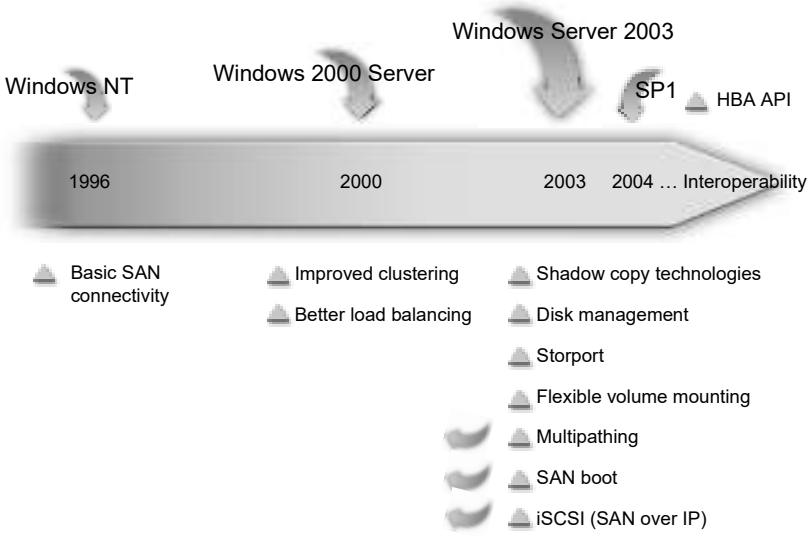

Figure 5. Designing in Support for SANs

- **Windows Server 2003: Designed for SANs**. The Windows Server 2003 platform has been enhanced with a number of new services and drivers designed specifically to support higher performance and fabric management of Fibre Channel SANs. While some of these SAN management capabilities have been back-ported to Server 2000 (see arrows in Figure 5), the majority are unique to Windows Server 2003. These include shadow copy technologies, enhanced disk management capabilities (such as the ability to grow volumes), flexible volume mounting, among others.

  The newest addition to Microsoft's support for SANs are the Microsoft Windows Management Instrumentation (WMI)-based HBA API and support tools, which will be released publicly in early 2004. Not only is the API critical to Fibre Channel SAN fabric management, but it provides the most complete implementation of this industry standard yet released to the public[10].

  In addition to support for Fibre Channel SANs, through its release of the iSCSI Initiator service and drivers, Microsoft is helping to promote SAN adoption over IP.

Storage Architecture

The underlying storage architecture of the Windows operating system has been redesigned and enhanced to provide more robust SAN functionality. Changes to the storage stack components were designed to provide the storage administrator with the tools needed to accomplish two broad goals: 1) effective management of the SAN infrastructure itself, including both the SAN fabric and storage array; and 2) support for SAN enabled solutions. The next sections describe how the various Windows storage stack components address each of these goals.

[10] See the INCITS T11 Web (http://www.t11.org/index.htm) site for more information about the capabilities of the HBA API. Note that the Windows version of this interface does not require multiple libraries from the vendors. All support is built in to the HBA drivers and surfaced to the API library by using the Windows Management Instrumentation infrastructure. This removes the confusion caused by conflicting DLLs.

Windows-Based SAN Management

Management of pooled storage resources that are shared among servers is fundamentally more complex than managing direct attached storage. Shared resources necessitate a discovery mechanism to detect and report to servers when storage resources are added or removed from the pool, as well as a mechanism to allocate those resources to the servers. Resource sharing is effective if a server accesses only those storage resources assigned to it; accessing the resources of another server can result in data corruption. Also, the storage network itself, consisting of HBAs, switches and storage arrays, is inherently more complex than the simple SCSI bus that connects direct attached storage.

Discovery of Storage Resources

The number of devices that attach to a storage area network can run into the hundreds or thousands; in DAS environments, the number of devices that can attach to each SCSI bus is limited to fifteen. Determining whether or not a device is present at a particular address on the direct attached SCSI bus is a simple matter of the server polling each address and waiting to see if a device responds. This method, for fifteen devices, can be practically instantaneous.

In SAN environments, on the other hand, where complex device topologies of switches and storage disks translate into hundreds of thousands of addresses, a polling-based discovery process would cause the network to come to a halt. Time is not the only constraint. Because SANs are designed to be continuously operational, devices must be able to attach and detach from the network without disrupting services. This poses another problem for servers: at any given point-in-time, how does a server determine which resources are connected to the SAN?

As has been discussed earlier, intelligent network devices, such as HBAs and switches, can maintain information about the status of the fabric. In a Fibre Channel network, these devices communicate by means of the Fibre Channel protocol, so passing this information back to servers (which speak SCSI) is problematic. Moreover, because hardware vendors have each interpreted the Fibre Channel specification slightly differently, interoperability among hardware components on the SAN is not guaranteed. Recognizing this as a considerable problem, Microsoft is helping to provide customer solutions, as explained in the following sections.

Fibre Channel Information Tool

To promote interoperability, the Storage Networking Industry Association (SNIA) developed, and, in 2002, released to the INCITS T11[11] group, the industry standard application programming interface (API) for managing Fibre Channel HBAs and discovering SAN resources[12]. Microsoft adopted the T11 standard, and as of May 2004, delivered the Fibre Channel Information Tool (FCINFO.EXE), the first complete implementation of the HBA API standard.

Used with vendor HBA drivers that support the HBA API standard and Microsoft WMI interfaces, the Fibre Channel Information Tool, a command line interface that supports scripting, can be used to gather information from both host bus adapters and switches.

There is one additional significant difference between the solution that SNIA has proposed and the implementation put forward by Microsoft. SNIA uses a polling-based method for discovering fabric events. Notification of fabric changes does not happen in real time. Under some circumstances, this delayed-event notification can interfere with I/O processing. In contrast, the Microsoft implementation, based on WMI (Windows Management Instrumentation), uses a true "eventing" model. Notification of changes to the fabric are instantaneous, and therefore not problematic for I/O.

Microsoft support for the HBA API industry standard means that information about any component of the SAN that the HBA connects to will now be available within Windows. This information includes:

- World wide names
- Driver and firmware versions (of HBAs and switches)
- Fibre Channel port attributes, such as port speed
- Link conditions
- Zoning information
- Device information (even for devices masked to the server)

Access to this information enables Windows to monitor the heath of the network connections, to determine whether the topology of the network is operating as designed, and to discover when resources are added or deleted. In short, Windows now has the ability to effectively manage a Fibre Channel storage area network fabric.

---

[11] T11 is the standards body responsible for FC device level interfaces.

[12] The initial SNIA design included two layers of code libraries: a high level "wrapper" library that all vendors could use to link to top level management layer applications, and a lower level library. Each was specific to a particular vendor device. Although it was obvious that hardware vendors would write their own lower layer libraries, that left it up to each application vendor to implement a wrapper library, meaning that each application vendor would have to ensure that it correctly interfaced with the private library of each hardware vendor. Microsoft solved this problem by providing the wrapper library for applications that run on Windows and requesting that hardware vendors implement WMI (Windows Management Instrumentation) in their drivers. The wrapper layer interfaces directly with the vendor device drivers, without the need for vendor specific libraries. Vendor drivers are tested to ensure that they are compatible with Windows.

iSCSI SANs: Plug and Play

The iSCSI initiator service enables the host computer system to discover target storage devices on the IP SAN and to determine whether or not it has access to those devices. By using the device discovery protocol iSNS (Internet Storage Name Service), an iSNS server (a free download from Microsoft) tracks information about initiators and targets and reports any changes to the initiator service. Because the iSCSI initiator service interfaces with Plug and Play, iSCSI storage targets can be dynamically discovered, and immediately made available to Windows servers.

Framework for Simplified Storage Resource Management

New storage devices on the network must be configured, formatted with a file system if required by the application, and allocated to the servers before use. In Windows Server SAN environments containing multi-vendor storage arrays, it is common for the storage administrator to have to move between storage consoles and use multiple management interfaces to accomplish this task.

The Virtual Disk Service (VDS), new to Windows Server 2003, provides a means by which to employ a single interface to configure all direct attached storage or storage on a Fibre Channel SAN[13]. VDS offers hardware vendors a general purpose interface, already coded for common disk management functions. Hardware vendors can build on this interface, writing the only code necessary to translate these common functions into their hardware-specific functionality. (This vendor-supplied code is sometimes referred to as a "hardware provider".) Storage administrators can use applications based on this service to manage storage arrays for all vendors that have supplied a VDS hardware provider.

The Virtual Disk Service interface supports common disk management functionality (such as drive letter assignment, NTFS formatting, and dynamic disk creation) through the scriptable DISKPART command line interface (CLI). More advanced disk management functionality (such as LUN creation, extension and masking) can be accomplished using the DISKRAID CLI, which is available through the Windows Server 2003 Deployment Kit *(*http://go.microsoft.com/fwlink/?LinkID=4941*)*. Full functionality is explained in the paper, "Storage Management using Windows Server 2003 and Windows Storage Server 2003 Virtual Disk Service and Volume Shadow Copy Service."

[13] VDS does not currently work with iSCSI SANs. Configuration of an iSCSI target must be done by using the storage array management tool.

Effective Resource Sharing

Once disk resources have been carved up into volumes spanning either single or multiple disks, these resources are formatted with the file system (NTFS) and assigned drive letters prior to use. This is done through a process known as volume mounting, under control of the Mount Manager, a component of Windows.

In SAN configurations, unlike DAS configurations, multiple servers have access to the same storage resources; however, to prevent overwriting and corruption of resources, only the server that actually owns the resources must be allowed to mount the file system.

Flexible Volume Mounting

In earlier versions of the Windows operating system, any storage volume that became visible to a server was automatically mounted, even if the storage capacity belonged to another server. This is especially problematic in a reboot scenario, in which servers coming back online must be able to correctly identify resources that belong to it and remount those resources.

Enabling Automount (part of the DISKPART command set functionality) in Windows Server 2003 has the following consequences, depending on the prior state of the volumes:

- **Volume previously mounted**: After a reboot, if a volume is determined to have been previously mounted by the server, the volume will be correctly re-mounted.
- **Volume previously unmounted**: If, after a reboot, it is determined that a volume has *not* previously been encountered by the server, then it will not be mounted unless the administrator explicitly mounts the volume.

Disk Signatures

A disk signature is a unique ID assigned to the disk for identification purposes. Note that disk signatures should not be used to determine ownership of resources across operating system platforms, since operating systems cannot read the disk signature of a different platform.

Disk signatures are used to ensure that assigned drive letters are persistent after rebooting. They are also used in clustering configurations, in which the same disk can be accessed by different servers. If the operating system fails to interpret the two views as a *single* disk signature accessed through two paths, the “duplicate” disk signature may be overwritten. In instances where multipathing drivers fail to work correctly, disk signatures could also be incorrectly changed.

In Windows Server 2003, the algorithm has been improved so that disk signatures are no longer altered unless there are multiple unique devices with the same signature.

SAN Security

Storage area networks were initially considered to be secure, in part because the Fibre Channel protocol was not well known, and in part because SANs originated in secure data centers. However, as SANs become more commonly deployed, and as many of them are now well outside the security of the datacenter, keeping mission critical data secure from both internal and external threats is becoming a more critical need.

Access Control

Limiting access to resources on the storage network is the most common security method employed. These methods include:

- **Physical Lockdown**. Limiting access to the datacenter minimizes the risk of physical tampering with the SAN components (such as array controllers) critical to device functioning, data access, and data availability. This method is appropriate with either Fibre Channel SANs or iSCSI SANs.
- **Zoning**. In Fibre Channel SANs, both switches and storage devices can be zoned so that only authorized servers or users can access specific ports on HBAs, switches and storage arrays. The new HBA API makes zoning configurations available to Windows, although it does not allow altering zones or zone sets.
- **LUN masking**. Also used with Fibre Channel SANs, this method controls access to storage resources at a more specific level: that of a particular storage logical unit within a target. The Virtual Disk Service DISKRAID commands can control the LUN masking functionality that is included with storage devices. Although LUN masking is also sometimes available as an HBA function, it is risky to use if the default state is to allow access to all LUNs; moreover, it adds an extra layer of configuration and hence additional opportunity to make mistakes. A new HBA can compromise LUNs belonging to other servers.

Authentication

It is not sufficient to restrict access to storage resources. It is also necessary to verify that the servers or users that are requesting access are, in fact, who they say they are. Most methods of authenticating user access occur at the level of the application requesting the data, rather than at the storage level. Authentication methods are appropriate for both Fibre Channel and iSCSI SANs, although currently only iSCSI SANs implement this functionality.

All points where the SAN connects to the LAN must be password protected to prevent unauthorized access. Switches, routers and other network components should also be password protected, enabling only authorized use. (The common practice of sharing a single administrator password is strongly discouraged, since it decreases the effectiveness of such protection.)

Encryption

Encryption is one of the most robust methods of ensuring that if the security of a storage area network is compromised, the data itself will not be usable. Data can be encrypted both while in transmission across LAN or WAN networks (using the IP security protocol, IPsec), as well as while at rest on the storage array. The drawback to encryption is that the encryption/decryption process is time intensive and costly in terms of processor cycles.

Encryption-based security solutions are under development. However, current host-based encryption degrades processing too much, and hardware solutions, available for iSCSI SANs, are still very costly.

Future Work

Current methods for authenticating, authorizing and auditing users and devices on the storage are neither end-to-end nor centralized. While many storage components such as switches and RAID arrays do provide some mechanisms for authorizing users before they can perform administrative tasks, few tie into the corporate security infrastructure. RADIUS (Remote Authentication Dial-In User Service), a widely used protocol that assists in authenticating users, is starting to be employed on SANs as well.

The Microsoft implementation of RADIUS, through Internet Authentication Service (IAS), integrates with Microsoft® Active Directory® directory service in Windows Server 2003. This implementation helps administrators centrally manage network access to LANs; Microsoft is currently working with storage partners to extend the use of the RADIUS server to deliver centralized authentication and authorization for SANs.

Windows Support for SAN solutions

The previous section discussed ways in which the Windows platform can be used to manage the SAN itself. This section focuses on how Microsoft enables the benefits of SANs to be realized on the Windows platform.

Centralized Storage Management

Consolidating storage resources onto the network allows for a centralized approach to storage management. The central benefit of SANs, provisioning storage resources from a single pool, eliminates the problem of some server having excess storage while others have no more free disk space. For application servers like Exchange where the message stores tend to increase dramatically over time, this flexibility in scaling is an enormous advantage.

With distributed storage resources, most maintenance functions, such as backups, are time intensive. Each server must be maintained and protected separately. Pooling storage enables centralization of such processes, not only reducing the time and resources (both equipment and personnel) necessary for maintenance, but also improving overall network performance. For example, removing the tape drive from an individual server and attaching it to the SAN can remove high volume data transfers from the LAN, thereby significantly improving network performance. The tape library can be shared for non-concurrent access with other servers on the SAN.

Many organizations are not simply consolidating their servers. They are also moving toward diskless servers (which maximize server density per square foot) to improve data center manageability. In these configurations, all server storage is on the SAN. To boot, the server must be able to access an operating system image from the SAN.

Microsoft support for boot from SAN enables organizations to realize many of these consolidation benefits. For full details on boot from SAN in a Windows environment, see the white paper “Boot from SAN in Windows Server 2003.”

Data Protection

In addition to the archival protection that tape backups provide, data can be protected against disk failure through redundancy, by using techniques such as mirroring and striping. While RAID (redundant array of independent disks) technologies can readily be deployed in direct attached configurations, the method is only useful when devices are not located more than a few meters away from the server. In contrast, RAID arrays on a SAN can be separated by several kilometers, facilitating disaster recovery scenarios that are not possible with direct attached storage.

Sharing Data on the SAN

Windows Server 2003 now contains the built-in technology used to make high fidelity point-in-time copies of data, known as shadow copies or snapshots. The Volume Shadow Copy Service provides the infrastructure to make point-in-time copies of single or multiple disk volumes. These shadow copies can be offloaded from the production server to a backup server where they can be backed up to tape. Such offloaded backups help keep production server performance high.

In a SAN environment, shadow copy technology facilitates sharing of data among servers, thereby enabling a very rapid restore process. Shadow copies of source data made on one system can be unmasked to a second server—a form of "virtual" transport that enables maximally up-to-date data sharing, without risking data corruption or loss of the original data.

Data Consistency and Application Performance

As a method of protecting frequently used data, there are two drawbacks to the backup process. One is that, for backups to be complete and fully restorable, the file or database should not be in use. Attempting to back up data while applications are open can result in data inconsistencies. As an alternative, open files are often not backed up, or full backups are restricted to once a week at a time when applications are not in use. Such backups, however, quickly become out of date.

A second difficulty is that backups involve transferring large amounts of data between application servers and backup servers, and the heavy load can lower application performance dramatically.

VSS solves this problem by working with "writers," application-specific software that ensures that application data remains consistent during the shadow copy creation process.

Common Microsoft applications, such as SQL, Exchange, and Active Directory all have writers that work with VSS, enabling live (open application) backups. The net result, for these applications and others, is high-fidelity point-in-time images of the data, which can be kept local for fast restores, or can be transported across the SAN for backup to tape at a later time.

Multipathing for High Availability

While centralizing storage on the SAN enables flexible storage provisioning and more effective resource management, it also carries an inherent risk: all of the storage is in one place. If access to the storage is lost, the entire organization loses access. Accordingly, SANs should be deployed with redundant components to prevent single points of failure, as well as with the switching mechanisms to redirect I/O requests from a nonfunctioning to a working redundant component.

Multipathing, the use of several redundant paths between server and storage device, can provide both high availability and better performance. Until the development of a native Windows architecture to support hardware vendor multipathing solutions, such solutions were unreliable in the Windows environment, especially in multivendor SAN configurations. With the release of the Multipath Driver Development Kit to vendors, it is now possible to ensure that each vendor's multipathing solution works correctly with Windows and with the storage devices of other vendors. See the paper "Highly Available Storage: Multipathing and the Microsoft MPIO Driver Architecture" for more details.

High Performance

SANs are designed for high performance high volume data transfers. Fibre Channel interconnects enable transfers faster than with direct attached storage. The Windows SCSI port driver, originally developed for parallel SCSI interconnects, was not designed to support the high speed performance capabilities of Fibre Channel networks.

Microsoft has designed a new port driver, Storport, for use with hardware vendor HBA miniport drivers. This driver not only delivers the high speed I/O transfers for which Fibre Channel was designed, it also provides fabric management capabilities necessary for determining fabric conditions in SAN environments. Database applications, such as Exchange Server and SQL Server™, can achieve maximum performance in such environments. Storport also provides significant improvements in availability when used with Microsoft Cluster Server (MSCS).

For further detail, see the paper “Storport in Windows Server 2003: Improving Manageability and Performance in Hardware RAID and Storage Area Networks.”

SAN Standards

One of the recurrent statements made throughout the course of this paper is that interoperability among multi-vendor SAN components has historically been unpredictable. This is especially true for components like switches and HBAs. Recognizing that, without interoperability, customers would reject storage networking solutions, industry organizations such as SNIA (Storage Networking Industry Association) were formed to help ensure that robust and effective storage solutions were developed throughout the industry.

SNIA has been instrumental in helping to develop specifications for standards reviewed for adoption by more formal standards bodies, such as INCITS, the International Committee for Information Technology Standards, and the Distributed Management Task Force (DMTF).

Microsoft Adherence to SAN Standards

To bring SAN capabilities beyond the enterprise, to midsize and smaller organizations, Microsoft has developed its SAN solutions with strict adherence to the Fibre Channel and iSCSI standards set by the leading standards bodies.

In Fibre Channel environments, adherence to the INCITS T10 and T11 committee standards has helped deliver Storport for high performance and manageability in SAN environments, the HBA API for fabric management, and MPIO for high availability and load balancing.

Microsoft support for iSCSI transfer protocol standards and work with IETF on the iSCSI security standards has helped deliver a secure alternative to high cost Fibre Channel SANs.

Windows Hardware Logo Program

To ensure that hardware solutions work with the Windows platform, Microsoft requires that hardware be tested under the Windows Hardware Logo Program for compatibility. This program is currently being strengthened, not only to ensure that third party solutions work with the Windows platform, but also to help ensure that vendor solutions work with one another.

Conclusion

Since Microsoft first delivered basic Windows connectivity to storage area networks, support for SANs has become an integral part of the company's storage initiatives. By delivering increasingly sophisticated enterprise storage solutions for both Fibre Channel and iSCSI SANs, Microsoft shows its commitment to providing businesses with robust and secure storage area network solutions.

Additional Resources

For further information on the technologies mentioned in this paper, see the following white papers:

Storage Management in Windows Server 2003 using the Virtual Disk Service and the Volume Shadow Copy Service (http://go.microsoft.com/fwlink/?LinkID=26119).

Storage Management Using Windows Server 2003 and Windows Storage Server 2003 Virtual Disk Service and Volume Shadow Copy Service (http://go.microsoft.com/fwlink/?LinkID=26119).

Microsoft Support for iSCSI (http://go.microsoft.com/fwlink/?LinkId=26635).

Highly Available Storage: Multipathing and the Microsoft MPIO Driver Architecture (http://go.microsoft.com/fwlink/?LinkId=26637).

Storport in Windows Server 2003: Improving Manageability and Performance in Hardware RAID and Storage Area Networks (http://go.microsoft.com/fwlink/?LinkId=26638).

Boot from SAN in Windows Server 2003 (http://go.microsoft.com/fwlink/?LinkId=27174).

Windows Server System is the comprehensive, integrated server software that simplifies the development, deployment, and operation of agile business solutions.
www.microsoft.com/windowsserversystem