**Storage Architecture Choice, SAN or NAS**

LaNet Merrill
SGI SAN & HA Product Marketing Manager

The intensity of focus on storage grows as IT spending shifts from servers to storage. The explosion of the Internet, the accompanying electronic data revolution and the expectation that everything should be online all the time are driving an unprecedented demand for storage. One only needs to look at the growth in the market capitalization of companies, or ask storage managers how much new storage capacity they will bring online this year to understand the role of storage as part of the infrastructure of the electronic data.

As demand rises technological innovation surely follows. As technological innovation accelerates, user confusion over the number of products and technologies available increases as well. The storage market is no exception to this rule.

One of the primary areas of innovation in storage today is in the application of networking technology to storage connectivity. This has given rise to two new, and commonly confused, storage topologies: Network Attached Storage (NAS) and Storage Area Networks (SAN). Many people don't understand the difference between NAS and SAN. Some think of them as competing technologies; journalists occasionally depict a battle between NAS and SAN for storage technology domination, while investment analysts mull over whether NAS vendors or SAN providers are a better investment choice. Users often ask which they should choose as their storage architecture.

NAS and SANs are both the products of the merge between storage and networking technologies. But, far from being competitive, they are in fact complementary technologies that productively coexist in many data centers and are even starting to converge. Together, they represent the future of storage and storage networking.

SAN Described

A SAN is quite simply a data transportation network dedicated to storage. More precisely, the Technical Dictionary published by the Storage Networking Industry Association (SNIA) defines a Storage Area Network as: "A network whose primary purpose is the transfer of data between computer systems and storage elements and among storage elements. A SAN consists of a communication infrastructure, which provides physical connections, and a management layer, which organizes the connections, storage elements and computer systems so that data transfer is secure and robust."

Unlike the traditional direct-attach storage model, a

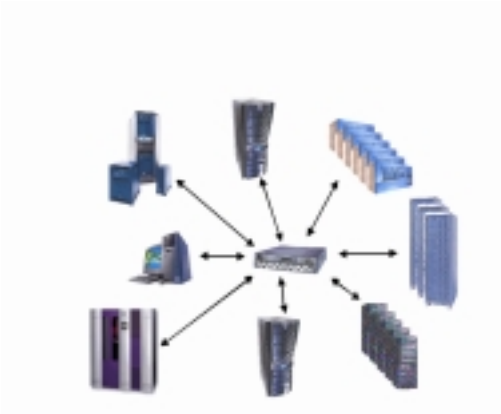

SAN attaches storage devices to servers in a networked fashion, using hubs, switches, routers and bridges to build the topology. Both the systems and the storage devices can, in theory, be heterogeneous in nature, though interoperability concerns limit some users to building homogeneous SANs. All commercially viable SANs today are built with Fibre Channel technology using SCSI Fibre Channel protocol layer (FC4).

SANs provide a number of advantages over direct-attached storage. They provide connectivity between servers and storage devices, making possible the sharing of storage resources between multiple servers, and thus enabling system managers to consolidate storage on a few large storage platforms. They also provide connectivity between the storage devices themselves, opening the way for direct movement of data between storage devices, vastly improving efficiency of data movement and processes, such as data backup or replication. The use of Fibre Channel, or most any other networking technology proposed for SANs, enables longer connectivity distances and higher performance than currently possible with SCSI technology.

Over time, SAN technology will ease the task of centralized storage management, and increase the implementation of remote management and data protection strategies, storage consolidation, system clustering and cross-platform data sharing.

True data sharing is the key to a SAN, it is the killer application. Shared file systems allows multiple servers to access the same data simultaneously and eliminates the need for data replication and the maintenance overhead of managing multiple copies of data. True data sharing allows any server to have the capabilities of direct read, write, and update access to source files.

NAS Described

NAS describes NFS file storage attached to a LAN network connection. The SNIA's Technical Dictionary defines Network Attached Storage as: "A term used to refer to storage elements that connect to a network and provide file access services to computer systems. A NAS Storage Element consists of an engine, which implements the file services, and one or more devices, on which data is stored. NAS elements may be attached to any type of network".

Note that the SNIA's definition says that a NAS system may be connected to any type of network. This is an important future consideration. Today, however, NAS systems are connected to a local area network (LAN). Because they can serve multiple heterogeneous clients, NAS devices provide a form of heterogeneous data sharing.

NAS systems can be either single purpose or a general purpose server running a NAS software package. The clients request access to files using standard Network File System (NFS) or Common Internet File System (CIFS) commands.

Some NAS vendors attempt to adhere to the appliance model of computing. That is, NAS devices are designed to do one thing, file serving. These NAS devices can contain embedded processors hosting a specialized operating system designed to

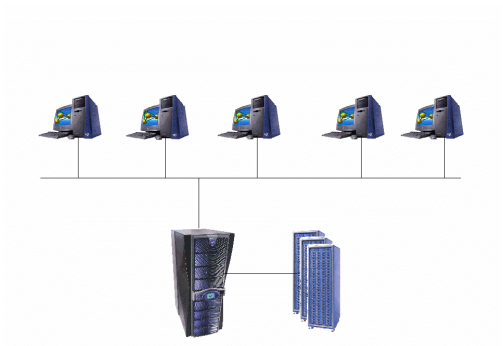

enable the NAS device to serve up files to clients with very high LAN performance.

General purpose system vendors often supply NAS software with their systems. Both NFS and CIFS server software packages are available. In choosing a server for NAS applications, the critical factors are IO bandwidth, disk and filesystem capacity and performance and network performance.

Appliance style NAS devices are typically designed to be very simple to install and configure. The storage they provide is often housed within the device's enclosure, though some NAS devices allow for the attachment of external storage.

General purpose system NAS packages generally include administration tools that make them convenient to install and configure.

The drawback of NAS for data sharing storage is the throughput and overhead of a LAN. LANs have fundamental bandwidth limitations, and may steal more than 50% of the CPU cycles on both ends for processing fast transfers. Shared SAN file systems provide a compelling alternative. They offer the benefits of NAS from the file system perspective by providing direct access to that storage rather than a LAN network-mounted file system.

Security

System managers must also consider security when designing and implementing a SAN. For example, zoning is a fundamental security feature of a SAN. Zoning enables users to logically segment a SAN to control access and visibility to servers and storage subsystems. With zoning, storage administrators can arrange fabric connected devices, servers, or workstations into virtual private SANs within the physical configuration of the SAN fabric. Zone members see only members in their zones and, therefore, access only one another. A device not included in any zone is not available to the devices in the zones.

There are two types of zoning: hardware and software. Software zoning uses worldwide names to define membership in a zone and hardware zoning is specified by physical port. For optimal security, both types of zoning should be supported, because zone overlap enables a storage device or server to reside in more than one zone and to be shared among different servers, which may be in separate zones. Another important zoning factor is its scalability.

Switch venders are now including authentication features like standard Public Key Infrastructure (PKI) and Access Control Lists (ACL) in their operating systems. PKI includes the use of digital certificates and digital signatures available in the market today. ACLs include restricting connections to devices, restricting connections of switches to the existing fabric and secure management communications.

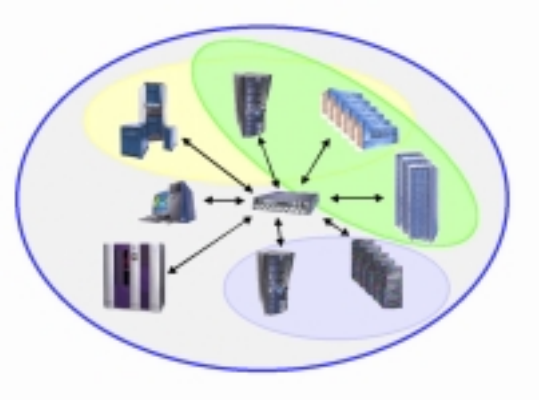

NT and UNIX have fundamentally different models of file system security. Security features reside in the platform operating system as well as the file system. UNIX uses user IDs (UIDs) to identify users, and group IDs (GIDs) to identify groups. The permission bits control read access (r), write access (w), and execute access (x). The full set of UNIX permission information stored with each file consists of:

- UID of the owner
- GID of the owner
- User perm bits (controls rwx for owner)
- Group perm bits (controls rwx for the group)
- Other perm bits (controls rwx for anyone else)

When UNIX users log in to a UNIX machine and filesystem their password is authenticated by the password and group maps on an NIS (network information service) master, users defined in these maps have existence on all machines participating in the NIS domain. NIS is equivalent to NT's SAM (security account manager). Different encryption and hatch is created for the password each time the user logs in.

When performing validation, UNIX determines whether the request is from the file's owner, someone in the file's group, or anyone else, and then uses the appropriate permission bits. SGI's IRIX security allows for United States Department of Defense (DoD) C2-level security certification. IRIX also includes access control lists (ACLs) and least-privilege capabilities (POSIX P1003.1e/2c), which allows allocation of system privileges individually.

Two types of ACLs exist, access and default. An access ACL governs the discretionary access to a file or directory. A directory can have an associated default ACL that governs the initial access for files find subdirectories created within that directory. A user who wants to gain access to the files in a directory must be listed on both ACLs and must be allowed, by IRIX file permissions, to successfully gain access. However, DAC, (Discretionary Access Controls) uses the default ACL only to initialize ACLs during file or directory creation and not to determine access permissions. Therefore, adding default ACLs to a directory does not protect existing files in that directory.

It should be noted that ACLs do not apply to an exported NFS file system. NFS exports data to a client *system*. NFS relies on the client to authenticate its own users. Given the complete local control a user might have over a client system, the NFS security model can be subverted relatively easily. NFS exporting and remote log ins for a UNIX file system must use security features available like exporting read only, secure RPC and secure shell

CIFS shares data with *users* (not machines, as with NFS). Authentication of a CIFS user's identity is done by the server, not the client (as with NFS).

NAS and NFS security resides with NT authentication domain controllers and IP access control. NT uses security IDs (SIDs) to identify both users and groups. The NT permissions for each file consist of:

- SID of the owner
- SID of the owner's group
- ACL (Access Control List) for the file

The ACL contains one or more access control entries (ACEs). Each ACE contains a SID, indicating the user or group to which the ACE applies, and a set of permission bits. NT permission bits include the three UNIX bits read, write, and execute as well as "change permissions" (P), "take ownership" (O), "delete" (D), and others. An ACE can either grant or deny the specified permissions. One ACE might grant read and write permission to the engineering group, but another ACE might specifically deny write permission to John Smith. Even if John is in the engineering group, he will be denied write access.

A problem with NT authentication happens when an NT domain administrator uses Server Manager for Domains to add a client machine into a domain, he or she really creates a client machine account in the PDC (primary domain controller) SAM (security account manager) database. The name of this account is the NetBIOS name of the client machine being added, in uppercase with a dollar sign ($) appended. Unfortunately, rather than asking the administrator what password to use for this new account, NT uses a default password: the NetBIOS name of the client machine, uppercase, in little-endian 2-byte Unicode format (which is the standard Microsoft Unicode format). It is of course hashed using the standard NT hashing method, but it's still a known value. This protocol has a serious flaw, in that the initial password is a known value, determined only by the NetBIOS name of the machine joining the domain.

This means that if someone can capture the network traffic of a machine joining a domain, he or she can follow the above password change sequence and learn the new machine password.

NFS and NT also differ in how they authenticate users. NFS is a connectionless protocol, and each NFS request includes the UID and GIDs of the user making the request. The UNIX client determines the UIDs and GIDs when the user first logs in, by looking at the files /etc/passwd and /etc/groups. NT networking is session based, so the identity of the user can be determined just once, when the session is first set up. At session connect time, the client sends the user's login name and encrypted password (actually the challenge and the client's response) to the file server, and the server determines the session's user SID and group SIDs. Servers commonly forward the name and password to an NT domain controller (DC) and let the DC perform authentication. NFS was designed to run on trusted networks behind a firewall and uses IP based security, IP based security can be spoofed on non-trusted public networks. The NFS server trusts that the users are who they claim to be. With Secure NFS, users must be able to decrypt a special key stored on the NIS or NIS+ server before the NFS filesystem will allow the user to access his or her files.

Coexistence

SAN describes a shared file networked storage topology, and NAS describes an NFS file server attached to the LAN. The questions asked by the system managers, then, typically come down to some variation of the following: Can SAN and NAS be used together, or must I choose to base my infrastructure on one or the other? When do I choose which technology?

The first question arises because, just as NAS/NFS provides shared access to file system data, one of the promises of SAN is also to provide high performance storage and data sharing. The good news is that the choice between SAN and NAS is not an either or decision. SAN topologies and NAS devices do, in fact, peacefully coexist in many data centers. For example, a SAN in the data center may network together servers with a number of large storage devices on which their data resides, while one or more NAS devices are attached to the LAN providing NFS file access to clients.

The choice of which technology to use is driven mainly by the requirement being addressed, and partly by timing. If the requirement is to provide shared file access to a large number of heterogeneous clients, NAS is generally the answer. NAS devices meet this need today with great efficiency. Because NAS systems are built on existing LAN and file system protocols, NAS technology is relatively mature in comparison with SANs. Most SAN file sharing solutions that exist, are generally aimed at specialized markets, such as high performance computing or video editing.

On the other hand, many system managers are struggling with the need to consolidate storage and servers to improve centralized management. Or, they want to take advantage of device-to-device data movement for applications such as backup or to reduce time needed to exchange data between severs for application workflow. In this case, SAN topologies can provide unique capabilities to address these requirements.

Application Requirement

| | |
|---|---|
| Storage or server consolidation | Use SAN |
| Data sharing between systems | Use NAS (SAN) |
| Access by UNIX, Linux, NT … | Use NAS (SAN) |
| Large I/Os or bandwidth | Use SAN |
| Small I/Os or transaction | Use NAS |
| LAN-free backup | Use SAN |
| Large number of clients | Use NAS |

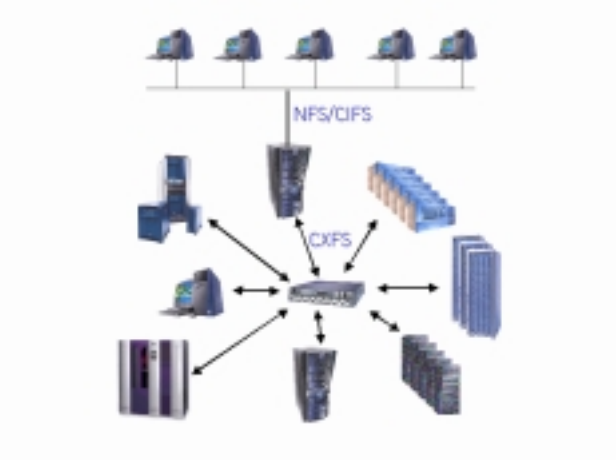


In the table above, (SAN) indicates that SAN storage sits behind the NAS server. This includes the host systems directly accessing the SAN's fabric infrastructure. While NAS provides data and file sharing using file based access methods (NFS and CIFS) over standard network interfaces (Ethernet, Gigabit, etc.), SAN technology including Fibre Channel Switches, and Fibre Channel RAID arrays provide storage for the NAS server and the direct attached host systems. Think of NAS as being an extra layer of file connectivity that sits in front of the storage being served by the SAN.

Summary

Reasons and benefits of using SAN include:

Reduce data replication
High performance applications
LAN-free backup
Overcome distance issues
Server and storage consolidation
Direct access shared files

Reasons and Benefits of using NAS include:

Simple to Configure and Install
Agnostic Platform and OS Connectivity
File sharing using NFS and CIFS
Small block data transfer performance
Mature Technology
Large number of clients

The Future

While SAN and NAS today are similar, but distinct, technologies, the lines between them are already beginning to blur. This technology convergence will likely take two forms. The first, which is already underway, is the use by NAS systems of SAN infrastructures for their back-end storage. While the storage capacity of many NAS systems is contained within the NAS device's enclosure, some SAN/NAS devices allow for highly scalable storage and performance outside the device

In many respects, this gives the system administrator the best of both worlds. Clients requiring file access get the benefit of an optimized NAS file server or/and the benefit of a high performing SAN architecture. The system manager can take advantage of the efficiencies of storage consolidation by exporting the NAS file system on a shared file system SAN storage device. And the system management staff benefits from the plug-and-play features of NAS setup and administration.

The second avenue of potential technology convergence is more speculative. Many vendors are exploring new network based storage options and protocols, including iSCSI FCIP, SoIP, InfiniBand and VI, although most of these are several years away from becoming a reality. The strongest push is in the IP enabled storage arena, which interconnects SANs with IP networks. Storage-over-IP transfers large blocks of data or files over local Gigabit Ethernet or wide area networks. In a storage-over-IP network, SCSI or next generation InfiniBand-based servers, tape libraries, and disk arrays can connect to and operate with devices on a Gigabit Ethernet network.

According to Metagroup users should monitor developments in iSCSI, but must realize that the trend to carry SCSI storage traffic over an IP network is immature. Multiple proposals are vying to become the standard, and we do not expect a resolution before 2002. The ability to carry SCSI storage traffic over an Ethernet network for communications within data centers will not mature until 2004-05.

Conclusion

We have seen that SAN and NAS technologies can offer the system manager with distinct and complementary capabilities. SAN topologies offer the ability to consolidate storage and improve data protection and storage management processes with a dedicated, high-performance storage network. The true value of SANs is realized once multiple servers are allowed to access and share the same file systems. Deploying SAN file systems will result in reduced total cost of ownership, improved overall performance, and more flexibility. NAS systems offer mature technology, low-administration file serving and file sharing for heterogeneous systems. We can also see that care must be taken to improve the security of these architectures. Used together, they provide a potent one-two punch for addressing data center requirements.