Emerging Network Storage Strategies

BY IRA GOODMAN

Because of the explosive growth of data on today's networks, administrators are faced with unprecedented storage and backup challenges. Two new storage strategies, SAN (Storage Area Network) technology and NAS (Network Attached Storage), may help them keep pace.

IS departments are hard at work. Teams of programmers are creating intranets and extranets that spawn scores of web pages and gigabytes of web statistics. Others are building complex web sites that support e-commerce applications. Still others are working on huge data warehouses that not only replicate much of the data already on an organization's databases (albeit cleaned up and neatly groomed), but also give rise to multitudes of new analyses and statistical reports based on "mined" data. IS workers are also developing hundreds of thousands of new ERP, multimedia, and spiffy easy-to-use graphics-based applications. All of these increase productivity and efficiency, but they also demand gigabytes and even terabytes of new disk space. And, of course, all of this additional data has to be backed up.

With industry studies now showing data on today's networks growing by 50 to 100 percent per year, storage vendors are working overtime to keep pace. The result has been several interesting innovations, and two of the most talked about appear at first to be mirror images of each other: NAS and SAN. However, in practice, they are quite different. NAS signifies a Network Attached Storage device that provides high availability disk storage. It is attached to the network directly and not through a server.

SAN is not a device, and today it is more of a concept than a concrete reality, although a few SAN products are already on the market. The idea of a Storage Area Network, a dedicated storage network connected to a LAN, is a powerful one and presents some very exciting possibilities. As Michael Petersen, a storage analyst for Strategic Research Corp. likes to say, "It's not a matter of if you will implement a SAN. It's a matter of when."

KEY STORAGE CHALLENGES

To understand what the newest innovations in storage are, we first need to understand the needs they are designed to meet. Today system administrators face four major challenges:

- **Scalability:** As I've already noted, storage requirements are accelerating at a frightening rate. An organization that uses only a few hundred gigabytes of storage today may find itself needing terabytes in a few years, especially if management decides to launch a data warehousing project or begin selling its products on the Internet.

- **Accessibility:** Expending the time and effort to build a data warehouse or set up an e-commerce site is useless if the applications for these projects are not accessible very quickly. Both NAS and SAN have been developed for speedy access. For example, even though most SAN equipment is currently using SCSI commands for communication, the hardware uses fibre channel as a communication medium, which is far faster than using a SCSI bus.

- **Availability:** The necessity of having computer systems available 24 hours a day, seven days a week is quickly becoming the norm. Commerce is now global, and the Internet never sleeps. System administrators must find ways to keep a system up in spite of the need for routine maintenance and critical backup procedures and the unscheduled downtime caused by component or system failure. Both NAS and SAN technologies address this availability problem.

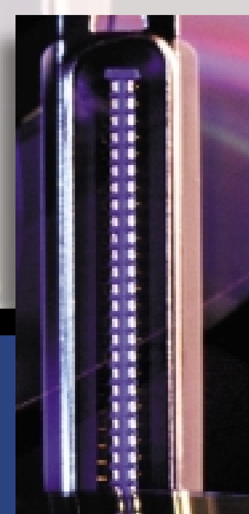

technical SUPPORT
Supporting Enterprise Networks and Operating Environments

AS SEEN IN TECHNICAL SUPPORT *MAGAZINE*
SEPTEMBER 1999 VOLUME 7, NUMBER 9

- **Recoverability:** Strategic Research studies set the financial impact of system failure for brokerage operations at $6.5 million per hour and credit card sales authorization at $2.6 million per hour. It is easy to see from these figures that the need to ensure that critical data is protected from a wide range of disasters, both natural and manmade, is a high priority for system administrators. Up-to-date copies of data must be stored in multiple locations so data can be recovered rapidly. SAN technology addresses this issue with particular effectiveness.

All of these challenges, combined with the heterogeneous nature of most network systems, make it easy to see why new storage technologies had to be developed.

EARLY ATTEMPTS TO IMPROVE SCALABILITY

By definition (at least in pre-SAN days), networks linked together computers, each of which had its own disk for storage. As the client/server model took hold, server machines tended to be far larger than client machines and would often take on the lion's share of the processing burden. But with the advent of VLDBs (Very Large Databases) and huge data warehouses, storage demands quickly outgrew what was available on a single server.

At first, system administrators merely attached more disk storage to each server as demand grew. This fast and simple approach resulted in what has lately come to be called JBODs (Just a Bunch of Disks). Hundreds of gigabytes of storage might be scattered throughout a site, and such a fragmented storage "strategy" is very difficult to manage and maintain. Both accessibility and availability suffers.

One storage solution that has been around for many years is RAID (redundant array of independent disks). RAID devices were an early attempt to achieve greater reliability and performance. Although RAID provides many advantages, it just hasn't been powerful enough to address the exploding scalability issue, and RAID devices are now being welcomed into the larger SAN solution.

Another interim solution to the storage problem is clustering. In this strategy, two servers work together and share their storage. If one server fails, the other can pick up its workload and, most important of all, provide access to the failed server's stored data. Each server can be attached to a RAID device, but again this is often simply not enough. In addition, the level of management complexity is much higher in such set ups.

THE NEXT STEP: NAS

Network Attached Storage (NAS) "appliances" (sometimes called "filers") are an attempt to deal with both the complexity of clustering servers and the need for fast access and high storage availability across a network.

Comparable to large disk arrays or disk "farms," NAS appliances attach to the network directly and not through a server, so they are both platform and operating system independent. Each NAS appliance contains an embedded processor running a simplified operating system or microkernel that understands network protocols and is optimized for I/O services. Since it is dedicated to one function — storage availability — a NAS operating system does not have the conflicting demands upon it that a general-purpose operating system does. Thus, a heterogeneous collection of servers can access more than a terabyte of storage quickly from a single NAS appliance.

To network applications, a NAS appliance looks like an ordinary server. To the client, it looks like a large storage device. The NAS advantage is that these appliances directly attach to a standard messaging network and are addressable via standard file system protocols, such as NFS (Network File System) and CIFS (Common Internet File System). But for all its performance and connectivity, NAS appliances have several drawbacks.

THE LIMITATIONS OF NAS

Some current NAS implementations have limitations, and a few of these are described in the following sections. However, it is important to remember that NAS appliances are a developing technology and these limitations may disappear in the next few years.

Bandwidth

Since NAS appliances are attached to the network and are used as a common storage device for a number of servers, network bandwidth limitations can be a problem. Large amounts of data must travel back and forth between the NAS appliance and the servers, creating quite a lot of network traffic. In addition, because the general speed of today's networks is usually between 10 and 100 Megabits per second, NAS appliances may be unable to deliver the performance that a critical file server needs.

Overhead

Even if the bandwidth restrictions were eliminated, network overhead can be a problem for sites using NAS appliances. Data must be transmitted via a network stack, and stack instructions take up generous amounts of CPU cycles and memory. This activity can be considerable when a NAS appliance is attached to a network.

Backup and NDMP

Because NAS appliances are essentially DASD, they are not appropriate for backup, which is usually done to a less volatile medium, such as tape. However, the introduction of NAS has made the complications of network storage more obvious than ever and has inspired the development of NDMP (Network Data Management Protocol), an open network protocol that defines common functional interfaces between an NDMP server and backup software.[1]

The aim of the NDMP protocol is to facilitate centralized control of data management while minimizing network traffic during backup. If NDMP is accepted as a standard, customers will be able to choose the best mix of NDMP-compliant hardware solutions for their environment much more easily. In addition, backup software vendors can concentrate more of their resources on enhancing their products instead of on incorporating yet another set of device drivers for the newest hardware.

Is It a Server or a Storage Device?

One final complication of using NAS appliances is the definition of exactly what they are. NAS appliances are dedicated devices directly attached to the network. As such, they have a bare-bones proprietary operating system, but are not full servers. On the other hand, they are not normal peripheral devices either. This creates difficulties for new storage strategies based on SAN technology. Should NAS appliances be treated as servers or storage devices? How should they be addressed and how should they be driven? This quandary opens up the question of whether or not NAS

1. For complete information about NDMP, visit the web site at www.ndmp.org.

will become an integrated part of SAN technology or remain a specific solution for particular environments.

A NEW STRATEGY: SAN

As we have seen, vendors continually offer more and more powerful storage devices, but aspects of their use remain problematic. What strategy could solve all of these problems and provide improved performance and round-the-clock availability? The answer is the Storage Area Network (SAN). Although at first glance SAN technology seems radical, it is an extension of a technique called subnetting, which network administrators have been using for years to isolate traffic and keep it off the main network.

In theory, a SAN is a separate high-speed network with a fibre channel backbone for shared storage. Each server on the LAN is also connected to the SAN, and these servers can be running any one of a wide variety of operating systems. Because they are directly connected to the SAN, the servers are all technically connected directly to the storage devices also. These storage devices can be of any type, including RAID arrays, automated tape libraries, standalone DASD, etc. If they want to add storage, network administrators would no longer need to purchase devices and connect them to individual servers. They would just add them to the SAN, and they would be available to all the servers on the SAN simultaneously.

However, there is one problem: Although many vendors are jumping on the SAN bandwagon by announcing support for SAN and limited SAN solutions are available now, a SAN that allows true interoperability between heterogeneous devices from multiple vendors is still in its formative stages. But when full-blown SAN technology does become available in the next year or two, it will provide many advantages.

Network Traffic

Because SAN topology organizes storage on a separate dedicated network with servers attached to both the SAN and the LAN, no data needs to travel across the LAN itself. Data will travel directly between the storage devices and the servers, bypassing the main network completely.

Performance

The SAN architecture takes advantage of fibre channel technology. Fibre channel

Backup and the Storage Explosion

From the very beginning of its history, network backup software has dealt with a relentless buildup of data on the networks it is designed to protect. At first, backup software used on networked machines merely moved data from individual machine to individual storage device without any centralized network-based control. However, this situation did not last long. Backup developers were quickly forced to add management capabilities that allowed network administrators to control backup processing themselves and not leave the task of safeguarding increasingly important data to individual users. Unfortunately, when control was centralized, so were storage devices. The manageability dilemma was replaced by a network traffic crisis each time data traveled to centrally clustered devices during backup processing. To remedy this situation, today's network backup software employs a distributed strategy in which storage devices are deployed across the network to reduce traffic and promote flexibility and scalability.

Sadly, network storage management as a whole did not evolve as quickly as backup storage management. Pressed for more storage on individual database or data warehouse servers, administrators usually just added standalone disks to the servers to relieve pressure. Deployment was not centrally controlled, and storage could not be shared across servers. If a disk crashed on one server, it had to be replaced physically. Since such occurrences are common, it's easy to see why the appeal of SAN technology, which promises flexible storage sharing, is almost irresistible.

Because the problems of control and storage management were addressed earlier for network backup than for the network as a whole, it's not surprising that elementary SAN techniques have been available in distributed backup software for more than a year. Both SAN technology and distributed backup software have these characteristics in common:

- both handle servers running heterogeneous operating systems (UNIX, NT, NetWare)
- both allow direct connection of storage devices to any server, although the techniques employed are quite different

If a device fails during a distributed backup, the backup image can easily be redirected to storage on another device, as SAN promises to do. And right now, distributed backup software allows individual servers to be directly connected to separate ports on an automated library with separate SCSI devices, forming a rudimentary SAN. As SAN technology matures, sites with distributed backup should find the transition especially easy.

offers a high bandwidth infrastructure designed to scale upwards with increasing I/O demands. A single, full duplex 1Gbps fibre channel link can achieve 200MBps peak bandwidth.[2] In addition, data access is much faster because servers are technically directly connected to the storage devices.

Accessibility and Availability

SANs are ideal for environments where high accessibility and availability are required. All of the techniques normally used in such environments (for example, mirroring and redundancy) would be easy to implement in a SAN. Target devices could easily be switched when routine maintenance is needed or a failure occurs, making storage systems very reliable and easily accessible (and recoverable) 24 hours a day, seven days a week.

Manageability

Since all storage would be connected to a single dedicated network and not to scattered servers or to subnets on a LAN, storage will be much easier to manage. Vendors are very aware of the manageability problems caused by the fact that many sites have a wide variety of heterogeneous operating systems and storage devices, and will be writing SAN management software accordingly.

Scalability and Flexibility

The use of fibre channel in SAN topology provides much greater scalability than SCSI connections. Fibre channel supports many more devices on a single segment than SCSI. In a fibre channel arbitrated loop (FC-AL), for example, 127 devices are supported as compared to 8 or 16 for SCSI. In addition, the distance covered by fibre channel is far greater than that of SCSI. A fibre channel single loop can attach devices for up to 10 km compared to under 100 meters for SCSI. And, if SAN does eventually support a large number of heterogeneous devices, scalability will be much easier and more flexible.

BUT WE'RE NOT THERE YET

SAN, as can easily be imagined, is a very complex technology. The kind of data sharing that it envisions would have to overcome currently incompatible file systems and data formats across disparate operating systems. The lack of standards leaves current SAN products with limited interoperability. In addition, fibre channel switches are first generation, and like many first generation technologies, they can be unreliable and expensive.

Oddly enough, another problem that some complain about with SAN technology is distance. SAN promises 10 km instead of 100 meters with SCSI, but even this distance (with extenders) isn't enough. Today's operating system and applications don't tolerate significant latency and will assume a device is down if it fails to respond immediately even when the problem is distance delays or network latency. Many administrators want globally distributed networks!

And since there are so many problems to solve with heterogeneous servers, file systems and devices, no single universal management console that would provide a single view of the entire system is likely to be developed quickly. Unfortunately, this console is absolutely necessary to enable the kind of flexible management that SAN technology promises, including the ability to resize and reconfigure disk volumes on the fly, along with built-in security to ensure that servers don't read or write data for which they are not authorized.

The result is that the SANs available today are limited in the number and type of servers and storage devices they support and are far from universally scalable. As a consequence, sites are installing multiple SANs, which is far from ideal but certainly a better situation than trying to manage scattered disks.

COMPLEXITY LEADS TO COMPETING ASSOCIATIONS

Just as the complexity of NAS led to the organization of a task group to develop NDMP, organizations have also developed to propose standards for SAN technology. Unfortunately, the complexity of the SAN problem has bred not one organization but two. The first is the Fibre Alliance, which has already proposed a MIB (Management Information Base) to the IETF (Internet Engineering Task Force), a standards organization. The MIB is an open standard that would allow every component in a SAN to be managed through SNMP (Simple Network Management Protocol) from a single interface and would cover such areas as configuring storage devices and reporting their status and capacity.

Another group, the Storage Network Industry Association (SNIA), has proposed a slightly different standard. Since the IETF will have the final say in choosing the standard, prospective users must sit tight until a decision is made. Meanwhile, vendors are marketing products in hopes that they will be so obviously superior that they will become the *de facto* standard. Also in the works are standards for LAN-free backup, file systems, and distributed resource management. For example, a proposed "storage container" would make all the devices in a pool look like a single device. A "third-party copy" interface would allow devices from different vendors to interact on backup and allow one vendor's device to take backup data from another.

SUMMARY

With the advent of SAN technology, the future of network (and backup) processing seems virtually unlimited — at least hypothetically. Direct high-speed access to storage resources that are available, but not restricted, to any server on a network would perhaps finally satisfy the need for limitless growth and constant access that today's hectic business climate demands. Certainly both NAS appliances and SAN technology will have an impact on how network administrators think about and manage storage, although exactly what form future implementations will take remains to be seen.

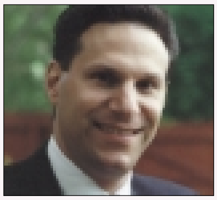

Ira Goodman is software services manager for Syncsort, Inc. He has 26 years of experience in information processing and has managed backup product support for more than a decade. He currently manages support for Backup Express, a network backup product.

2. Typically a SCSI device operates at 20MBps. Ultra-Wide SCSI operates at 80MBps.