

# Using Network Attached Storage as a Gateway to Storage Area Networks

Dave Rosenberg  
Hewlett-Packard Company  
825 14<sup>th</sup> Street SW  
Loveland, CO 80537  
[dave\\_rosenberg@hp.com](mailto:dave_rosenberg@hp.com)

# Using Network Attached Storage as a Gateway to Storage Area Networks

## Contents

Overview .....	
NAS Architecture.....	
SAN Architecture.....	
NAS and SAN Fusion .....	
How to Incorporate NAS into an Existing SAN .....	
The Future of NAS/SAN Fusion .....	

# Using Network Attached Storage as a Gateway to Storage Area Networks

Network Attached Storage (NAS) is defined as a purpose-built device (appliance) that presents one or more file systems to client systems via a network. A Storage Area Network (SAN) is defined as a dedicated network of interconnected storage and servers that offers any-to-any communication between these devices and allows multiple servers to access the same storage device independently. A SAN can be a very large pool of storage for running enterprise level applications. Can you connect a NAS device into a SAN so network clients get efficient file system access to all that SAN storage? Absolutely!

## Overview

As the need for storage continues to grow at an unprecedented rate, the question of how to meet that need and keep IT budgets under control remains at the forefront for most enterprise system managers. NAS and SAN are two recent technologies helping to address this problem.

### NAS

NAS devices have been around for a few years and are rapidly growing in popularity. Often called filers, NAS represents a simple and quick method for adding storage to a network. NAS devices are essentially dedicated file servers consisting of a processor, network interface, disk storage and NAS software. Most NAS boxes use operating system and file system software specifically tuned to provide efficient network file access. NAS boxes do not run databases; they do not run third party applications; and they do not provide print serving. NAS boxes only serve files, so they are made to serve them very well. Most NAS devices are "heterogeneous," meaning they can provide file access to both Windows (SMB/CIFS) and UNIX (NFS) based clients. Some also can be used as Netware filers.

NAS boxes are known for their ease of installation on the network. Most require only a Local Area Network (LAN) connection and power. A web browser running on a network client is used to configure the NAS. NAS boxes can be installed and file storage access made available to clients without ever taking down the enterprise. That makes them extremely flexible to use.

NAS is often described in terms of a "head" and "storage." The NAS head consists of the dedicated processor, network interface and NAS software. The storage part is the disk controllers and disk drives.

Benefits of NAS technology include:

- Multiple client file system access - This capability is not available with SANs. While SANs do offer multiple server access to devices, most applications expect multiple client access to files. This is one of the unique features of NAS.
- Ease of understanding - While Fiber Channel has matured in its expanded use, some still regard it as less understood than the IP technologies that users have experienced for years. As such, users understand the technical foundation of NAS; this understanding provides high confidence in its use.
- Ease of maintenance - NAS systems are comprised of elements from one manufacturer and are well integrated. Many SANs are integrated with multiple components from multiple vendors, thus creating greater complexity for customers who attempt to maintain them.

Component failures on the NAS can be fixed without application server downtime.

- Inexpensive implementation - NAS systems allow users to leverage their existing network infrastructures, and there is no need to build a back-end SAN to implement NAS.

- Protection against failures - Many NAS systems can automatically replicate its files to other disk storage in “clustered” configurations, thus providing an alternative storage to draw from if the primary storage goes down or becomes is corrupted.
- Virus protection – Some NAS systems include anti-virus software to prevent hostile attacks against files.

## SAN

SANs provide an efficient method for linking storage systems with servers and represent a significant improvement over traditional Direct Access Storage (DAS). With DAS, fixed storage systems are “hard wired” to a single server. If a server needs additional storage capacity, the only option is to upgrade the system. IT budgets are, of course limited, now more than ever, so not all servers in an enterprise can be upgraded to meet user need. Furthermore, once upgraded, a server is seldom downgraded. As a result, the additional storage remains on the system forever, whether or not the server’s users still need it. This creates a mismatch between available system resources and actual user requirements. The enterprise may have the necessary storage, but it’s not on the server the where users need it. SANs can reduce this resource mismatch problem, providing significant cost savings. SANs create a pool from one or more disks or storage systems and allow IT managers to allocate necessary storage from the pool to individual servers as needed. Unlike a NAS device, servers in a SAN are often general-purpose. A SAN server can be running a database, email or any other application. Of course, a server on a SAN can also be a file server. SANs allow scaling of storage by decoupling the direct link of storage and servers. SANs allow system management to focus from device management to management of the entire storage environment. SANs also allow multi-vendor storage to be managed as one coherent pool.

SAN benefits include:

- Serving raw devices – SAN provides raw block level storage to its servers, which can run any type of application. NAS cannot serve raw blocks. If your application requires access to a raw device, NAS is not an option.
- Flexibility and scalability - SAN storage can grow by simply attaching additional storage capacity to the SAN.
- Easier data backup - The cross sectional bandwidth of SAN makes large-scale backup and recovery much easier and faster. With NAS, performing a full backup of disk space used for the file system can task the I/O system much more than in any other application.

## NAS and SAN - Together

Because of increasing volume requirements, some applications for NAS are pushing the bandwidth and performance limits of traditional NAS systems. This trend has created a requirement for dramatically greater NAS storage to meet these increased data needs. The net effect of these additional requirements is a convergence of NAS and SAN technologies, referred to as NAS/SAN fusion. Combining NAS and SAN technologies allows IT managers to capitalize on the advantages of both networked storage approaches. It also enables corporations to better cope with the unrelenting pressures of data storage growth and limited IT resource investment funds.

## NAS Architecture

The purpose of a NAS device is to serve files to client workstations. A NAS system physically resembles a general-purpose server, consisting of a processor, operating system, network interface and storage. However, because NAS is dedicated to file serving, the software that operates on a NAS platform is limited to provide file serving.

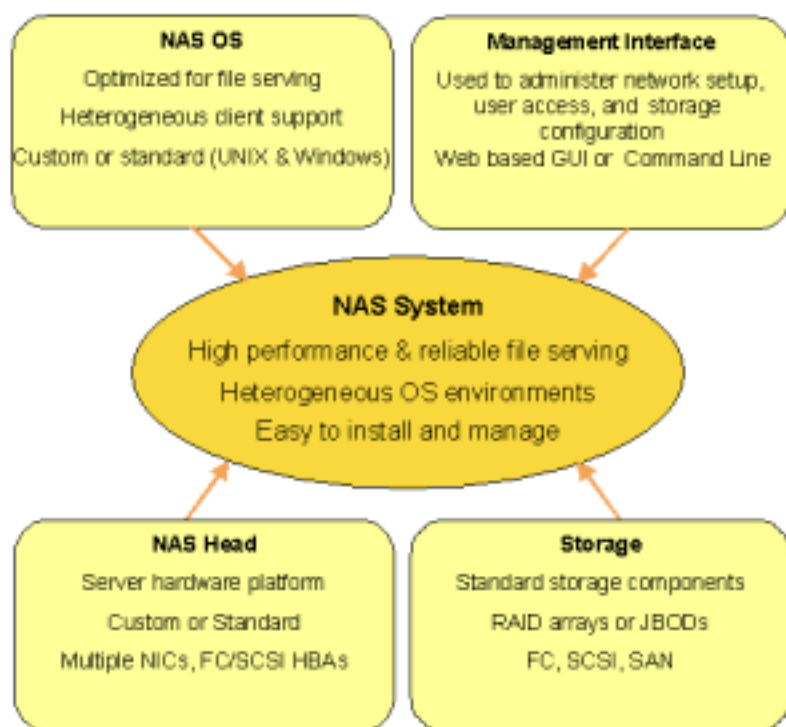


Figure 1 – NAS Architecture

## NAS Operating System

The NAS operating system can be a standard operating system such as Windows Server or UNIX. A NAS operating system can also be proprietary. Proprietary NAS operating systems can be made specifically to speed file serving. This specialization gives them a performance advantage over industry standard operating systems that are designed to support general-purpose applications. Most NAS operating systems support both Windows and UNIX file structures, so that they can serve files to both Windows and UNIX clients. Disk caching and other techniques often are used to speed file serving.

## NAS Head

The NAS Head primarily consists of the server hardware platform, network interface cards (NIC) and Fibre or SCSI Host Bus Adapters (HBA). The server platform can be custom for the NAS but typically is based on a manufacturer's standard server. A standard Ethernet NIC provides connection to the network and client workstations. The HBA(s) provides connection to the storage components. Often multiple HBAs will be used to improve performance by increasing processor to storage system bandwidth.

## Storage

The storage component of a NAS can vary greatly. Some NAS systems, especially at the low end, are packaged such that the storage resides with the head in a single chassis. In the middle to upper performance and capacity range, the storage component is separated from the NAS head. This provides for greater storage options such as increased capacity, RAID, redundant storage networking, etc.

## Management Interface

The management Interface is used to administer the NAS, including setting an IP address, defining users, creating storage shares and exports, etc. The user interface for a NAS can be either command line driven or

graphical. Most GUI management interfaces are accessible over the web. Command line interfaces are managed via telnet.

### NAS in the Network

The figure below shows a NAS system configured into a network. Files are shared as directories over the network. Windows clients see the shared storage as directories that can be mapped as local disk drives. UNIX systems see the storage as exported file systems that can be “mounted.”




Figure 2 – Traditional NAS system on a LAN

### SAN Architecture

As defined earlier, a SAN is a dedicated network of interconnected storage and servers that offers any-to-any communication. A network of switches and high-speed fiber provides the “any-to-any” connection. Data is transferred between servers and disks (or disk arrays) at the block level. The switches determine data block routing between a particular server and a particular disk device. Together, the high-speed fiber and the switches are referred to as a “fabric.” Servers connect to the fabric using a Host Bus Adapter (HBA). Although the fabric can allow any server in the SAN to communicate with any disk system on the SAN, in reality, block access must be very well controlled. If two servers were to access the same data, corruption could occur. Specific software must be used when multiple servers are to share the same SAN storage. This special software is usually used by systems designed for fail-over or data mirroring.

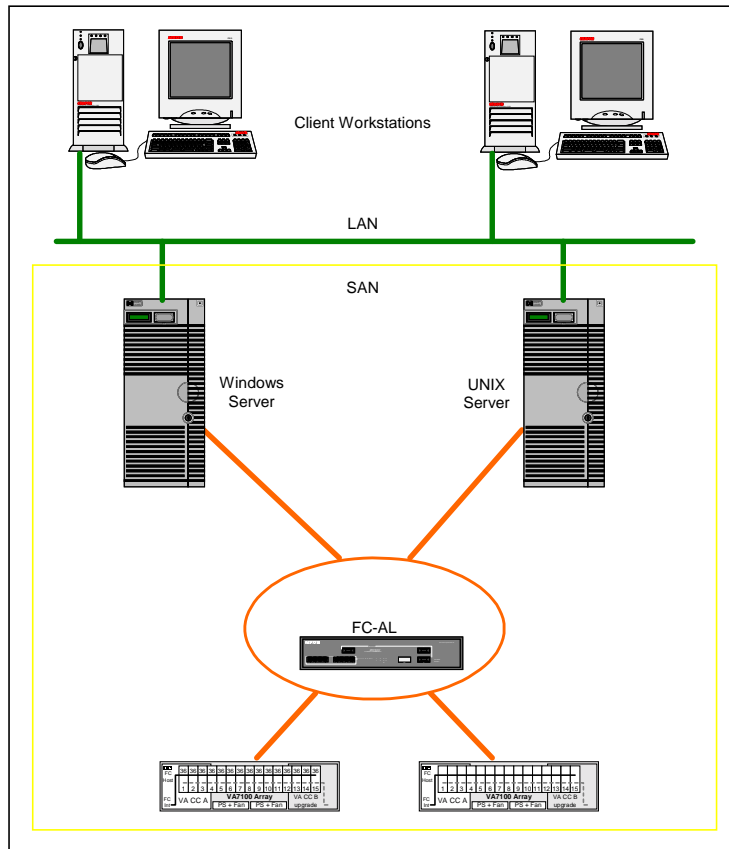


Figure 3 – Storage Area Network with Windows and UNIX Servers

To prevent one server on a SAN from corrupting another server's data on the same SAN, storage on a SAN is partitioned and allocated to individual servers. Partitioning of disk system storage is based on Logical Unit Numbers (LUN). A LUN is "the portion of a storage array capacity that appears as a physical disk." A RAID manager creates LUNs and can present multiple LUNs via a single port. A JBOD (just a bunch of disks) is typically a single LUN. Only storage capacity that has been assigned a LUN is available to operating systems.

There are three primary methods for insuring allocated storage remains accessible to only a single server.

- Server level control
- Control access within the network
- Control access within the storage array

**Server level control** requires special software on each server on the SAN, essentially delegating access control to each server, making the server itself responsible for accessing storage assigned to it. Server level control allows storage sharing at a block level, controlling multiple OS access to LUNs. Server level control allows assignment and reassignment of storage from one server to another without the need to reboot servers. Host control can offer a simplified management through automatic server and storage discovery.

**Control access within the network** uses a system called zoning. Zoning can be either "hard" or "soft". Hard zoning specifies which ports within a switch can communicate with the other ports in the switch. Hard zoning is very secure and often used to demonstrate a physical separation between storage systems shared between multiple users. Zoning can be somewhat inflexible, however. Because it uses physical port connections to

control access, all the LUNs routed to a particular port are accessible to all other ports permitted to communicate with it. Soft zoning is similarly limited however, it uses a mapping of “world wide names” to specify which server can access which port on a storage device. A worldwide name is a unique hardware address within Fibre Channel that identifies each HBA in servers. It is similar to the MAC address of a network card used in Local Area Network topologies.

Example:

- A Windows NT server and a UNIX based server are connected via fabric in a SAN environment. The Windows NT host desires exclusive access to LUN 0 and LUN1 on Array 1 and shares its associated file systems to client workstations.
- A UNIX server is on the same SAN,. It desires access to LUN 2 and 3 on Array 2.




Figure 4 – Storage Area Network with Zoned LUN Allocation

**Control access within the storage array** places the responsibility of access control with the storage device using a method called “LUN masking.” Because access control resides with the storage device, this method provides very secure access control. In LUN masking, the RAID controller within the storage system makes LUNs available only to server HBAs with specific worldwide names. LUNs that are not to be accessed by a particular server are hidden from view or “masked.” This makes it impossible for any server to access a LUN unless the storage device has permitted it. Information about which server HBA can access the storage device is kept in a LUN masking table in the raid controller of each array. Since connection tables must be maintained within the storage devices, this method does have the drawback that control must be managed device by device rather than centrally.



Example:

- A Windows NT server and a UNIX based server are connected via fabric in a SAN environment. The Windows NT host desires exclusive access to LUN 0 on Array 1 and shares its associated file systems to client workstations.
- A UNIX server is on the same SAN, it desires access to LUN 2 and 3 on Array 2 and shares its associated file systems to client workstations.
- LUN 1 is not allocated to any server.

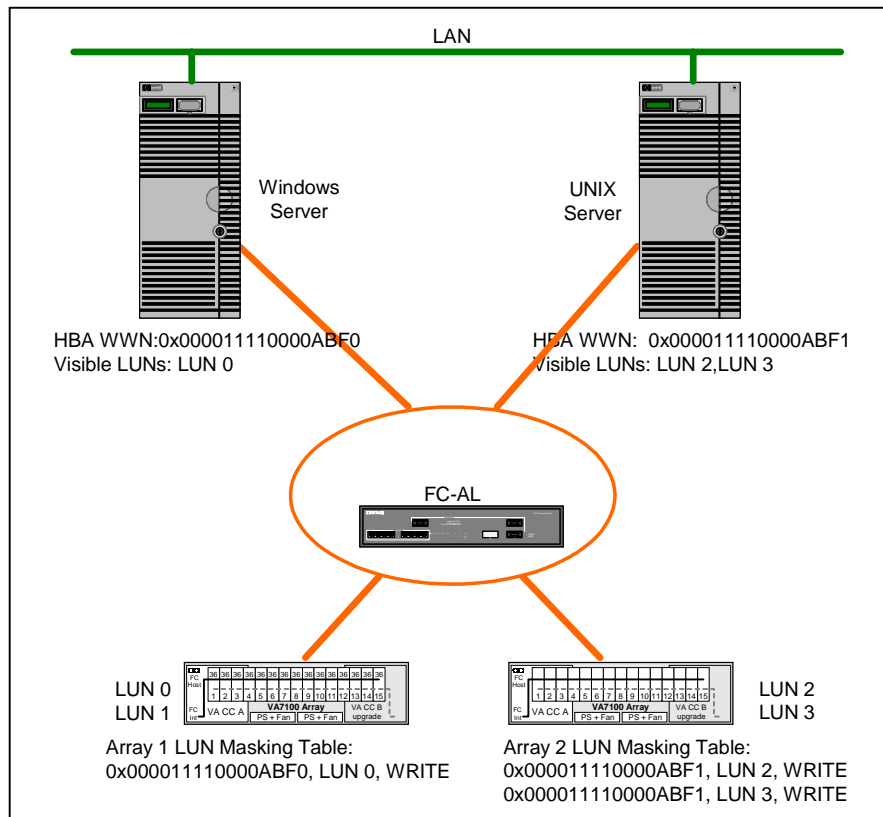


Figure 5 – Storage Area Network with Server LUN Masking

## NAS and SAN Fusion

Combining NAS and SAN is a union of two technologies, resulting in significant benefits.

When using NAS in combination with SAN, the NAS device is used as a dedicated file server, providing clients with file access of SAN storage. Only the “head” of the NAS is used in a SAN configuration. Storage for the NAS head is then allocated to NAS head from that storage available within the SAN. Thus, the NAS head becomes a file server “gateway” into the SAN, and the SAN becomes a pool of storage for use by the NAS head. Storage allocated to the NAS can be increased or decreased as user demand varies.

Using NAS as a gateway to SAN storage has another advantage. General-purpose servers placed on a SAN are either Windows file system based (CIFS) or UNIX based (NFS). But most NAS devices are heterogeneous, and can provide both CIFS shares and NFS exports of the same data. Therefore using NAS as a gateway to a SAN can provide both Windows and UNIX based workstations access to the same files. Networks with platforms or protocols in transition or with planned migration will benefit from a NAS-SAN gateway. Plans to migrate from Microsoft Windows to Linux are common scenarios today. Such a migration is easily facilitated with a NAS-SAN gateway that can handle CIFS and NFS concurrently. NAS-SAN fusion provides substantive advantages of greater storage consolidation and storage management.

#### Improved NAS Backup

The complications of backing up NAS data from traditional file serving appliances are numerous. And while some clever NDMP methods have been developed to ease this complexity, the task of backing up file data from a NAS appliance remains challenging, and certainly less reliable than the comparable experience of running backups from SAN-attached disk to tape. As such, the consolidation and convenience that NAS/SAN fusion provides with regard to NAS data backup is significant.

Although separate LUNs are created and used for NAS and SAN data, the backup of all SAN-resident data is collectively controlled. This greatly simplifies the backup of NAS data when compared to traditional stand-alone NAS integrated appliances.

## How to Incorporate NAS into an Existing SAN

Incorporating NAS into an existing SAN is not very different from adding any other server into a SAN. A NAS device (head only) is basically another server on the SAN. It just happens to be dedicated to file serving.

The NAS head to be connected to the SAN must support Fiber Channel. If you purchased a NAS head specifically to integrate with SAN it should have at least one Fiber Channel HBA. If you are converting a stand-alone NAS to operate on a SAN, you may need to add a Fiber Channel HBA to your NAS head. Check with the NAS manufacturer to make sure there are the appropriate drivers, etc. available to support Fiber Channel.



NAS. A few NAS manufacturers, who happen to also be developers of SAM software, have started providing integrated NAS-SAM solutions. However, these solutions still remain limited.

## The Future of NAS/SAN Fusion

For the most part, the management tasks for NAS and SAN storage, NAS file services, NAS processor state-of-health and SAN quality-of-service are completely segregated, requiring separate human resource oversight for each of these activities. Market demand for an “open” SAN environment, where devices from multiple manufacturers can co-exist and be managed by centralized SAMs, are forcing the industry to set standards for device management.

### Storage Networking Industry Association

The Storage Networking Association (SNIA) was formed in 1997 with the mission to “ensure that storage networks become efficient, complete, and trusted solutions across the IT community. ” Working through the SNIA, developers of storage and networking products in addition to system integrators, application vendors, and service providers are working to establish standards for storage device managers. These standards will allow NAS and SAN device managers to interoperate, thus simplifying the overall management of both.



Figure 7 – Multi-vendor Open SAN with NAS

## Common Information Model Based Management

Several developers are moving towards achieving this interoperability based on the Common Information Model (CIM).” CIM is a common data model of an implementation-neutral schema for describing overall management information in a network/enterprise environment. CIM is based on an object-oriented model. Using CIM as a common base, NAS and SAN will one day be able to communicate with each other such that both can be managed from a central user interface. The IT manager will not need to use a SAM tool to allocate LUNs on the SAN, then go to the NAS management program to create volumes, then file systems and finally expose them as sharers for a client workstation to use. Rather the IT manager will tell the SAN he needs 150 GB of Windows file system space called “Payroll.” And it will be there.