

Module – 7

Network-Attached Storage (NAS)



Module 7: Network-Attached Storage (NAS)

Upon completion of this module, you should be able to:

- Describe NAS, its benefits, and components
- Discuss NAS file-sharing protocols
- Describe different NAS implementations
- Describe file-level virtualization

This module focuses on benefits and components of Network-Attached Storage (NAS). It also focuses on NAS file-sharing protocols, different NAS implementations, and file-level virtualization.

Module 7: Network-Attached Storage (NAS)

Lesson 1: NAS Components and Benefits

During this lesson the following topics are covered:

- File sharing technology evolution
- Benefits of NAS
- NAS components
- NAS file sharing protocols
- NAS I/O operations

This lesson covers a comparison of general purpose file server and NAS. It also describes key components of NAS, file sharing protocols (NFS and CIFS), and NAS I/O operations.

File Sharing Environment

- File sharing enables users to share files with other users
- Creator or owner of a file determines the type of access to be given to other users
- File sharing environment ensures data integrity when multiple users access a shared file at the same time
- Examples of file sharing methods:
 - ▶ File Transfer Protocol (FTP)
 - ▶ Distributed File System (DFS)
 - ▶ Network File System (NFS) and Common Internet File System (CIFS)
 - ▶ Peer-to-Peer (P2P)

File sharing, as the name implies, enables users to share files with other users. In a file-sharing environment, a user who creates the file (the creator or owner of a file) determines the type of access (such as read, write, execute, append, delete) to be given to other users and controls changes to the file. When multiple users try to access a shared file at the same time, a locking scheme is required to maintain data integrity and, at the same time, make this sharing possible.

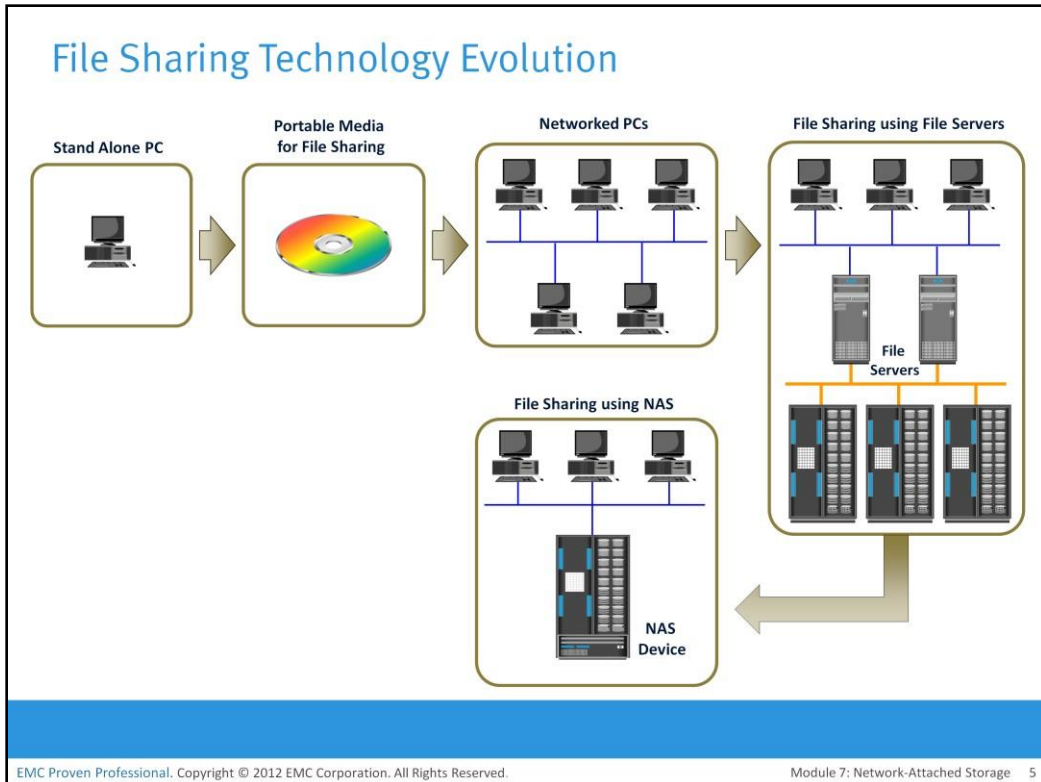
Some examples of file-sharing methods are; File Transfer Protocol (FTP), Distributed File System (DFS), client-server models that use file-sharing protocols such as NFS and CIFS, and the peer-to-peer (P2P) model.

FTP is a client-server protocol that enables data transfer over a network. An FTP server and an FTP client communicate with each other using TCP as the transport protocol.

A distributed file system (DFS) is a file system that is distributed across several hosts. A DFS can provide hosts with direct access to the entire file system, while ensuring efficient management and data security.

The standard client-server file-sharing protocols, such as NFS and CIFS enable the owner of a file to set the required type of access, such as read-only or read-write, for a particular user or group of users. Using this protocol, the clients mount remote file systems that are available on dedicated file servers.

A peer-to-peer (P2P) file sharing model uses peer-to-peer network. P2P enables client machines to directly share files with each other over a network. Clients use a file sharing software that searches for other peer clients. This differs from client-server model that uses file servers to store files for sharing.



Traditional methods of file sharing involves copying of files to a portable media, such as floppy diskette, CD, DVD, or USB drives and delivering them to other users with whom it is being shared. However, this approach is not suitable in an enterprise environment in which a large number of users at different locations need access to common files.


Network-based file sharing provides the flexibility to share files over long distances among a large number of users. File servers use client-server technology to enable file sharing over a network. To address the tremendous growth of file data in enterprise environments, organizations have been deploying large numbers of file servers. These servers are either connected to direct-attached storage (DAS) or storage area network (SAN)-attached storage. This has resulted in the proliferation of islands of over-utilized and under-utilized file servers and storage. In addition, such environments have poor scalability, higher management cost, and greater complexity. Network-attached storage (NAS) emerged as a solution to these challenges.

What is NAS?

NAS

It is an IP-based, dedicated, high-performance file sharing and storage device.

- Enables NAS clients to share files over IP network
- Uses specialized operating system that is optimized for file I/O
- Enables both UNIX and Windows users to share data



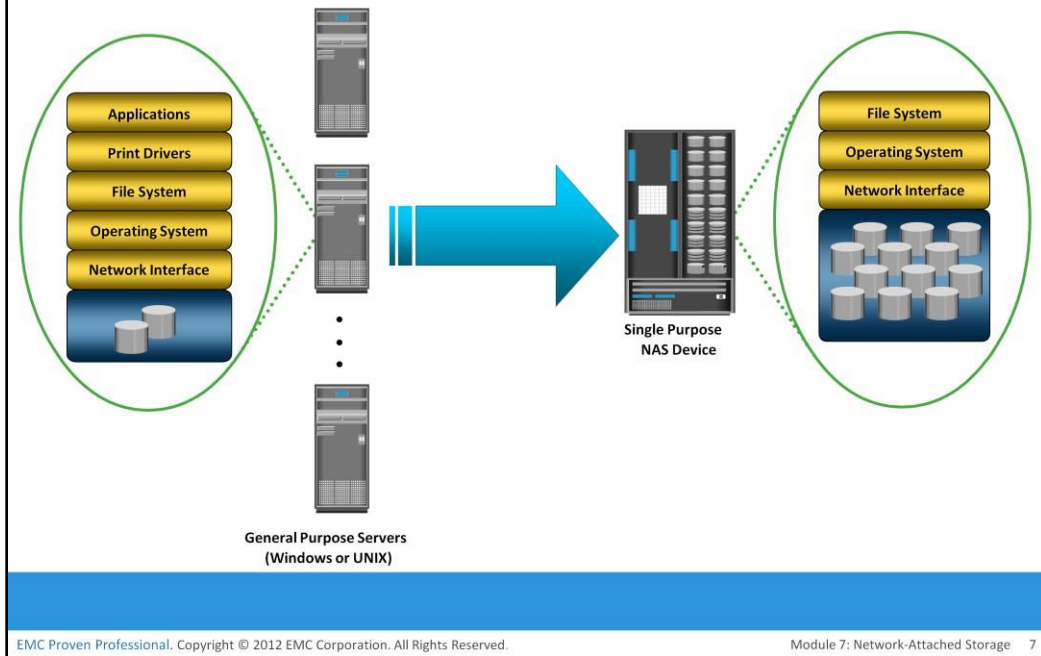
EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Module 7: Network-Attached Storage 6

NAS is a dedicated, high-performance file sharing and storage device. NAS enables its clients to share files over an IP network. NAS provides the advantages of server consolidation by eliminating the need for multiple file servers. It also consolidates the storage used by the clients onto a single system, making it easier to manage the storage. NAS uses network and file-sharing protocols to provide access to the file data. These protocols include TCP/IP for data transfer, and Common Internet File System (CIFS) and Network File System (NFS) for network file service. NAS enables both UNIX and Microsoft Windows users to share the same data seamlessly.

A NAS device uses its own operating system and integrated hardware and software components to meet specific file-service needs. Its operating system is optimized for file I/O and, therefore, performs file I/O better than a general-purpose server. As a result, a NAS device can serve more clients than general-purpose servers and provide the benefit of server consolidation.

General Purpose Servers Vs. NAS Devices



A NAS device is optimized for file-serving functions such as storing, retrieving, and accessing files for applications and clients. As shown in the slide, a general-purpose server can be used to host any application because it runs a general-purpose operating system. Unlike a general-purpose server, a NAS device is dedicated to file-serving. It has a specialized operating system dedicated to file serving by using industry standard protocols. Some NAS vendors support features, such as native clustering for high availability.

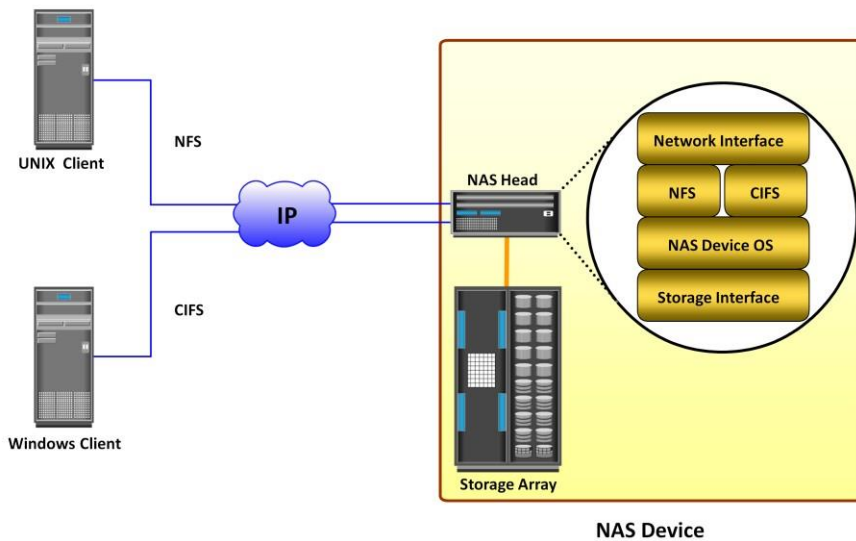
Benefits of NAS

- Improved efficiency
- Improved flexibility
- Centralized storage
- Simplified management
- Scalability
- High availability – through native clustering and replication
- Security – authentication, authorization, and file locking in conjunction with industry-standard security
- Low cost
- Ease of deployment

NAS offers the following benefits:

- **Improved efficiency:** NAS delivers better performance compared to a general-purpose file server because NAS uses an operating system specialized for file serving.
- **Improved flexibility:** Compatible with clients on both UNIX and Windows platforms using industry-standard protocols. NAS is flexible and can serve requests from different types of clients from the same source.
- **Centralized storage:** Centralizes data storage to minimize data duplication on client workstations, and ensure greater data protection.
- **Simplified management:** Provides a centralized console that makes it possible to manage file systems efficiently.
- **Scalability:** Scales well with different utilization profiles and types of business applications because of the high-performance and low-latency design.
- **High availability:** Offers efficient replication and recovery options, enabling high data availability. NAS uses redundant components that provide maximum connectivity options. A NAS device supports clustering technology for failover.
- **Security:** Ensures security, user authentication, and file locking with industry-standard security schemas.
- **Low cost:** NAS uses commonly available, and inexpensive Ethernet components
- **Ease of deployment:** Configuration at the client is minimal, because the clients have required NAS connection software built in.

Components of NAS



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Module 7: Network-Attached Storage 9

A NAS device has two key components: NAS head and storage. In some NAS implementations, the storage could be external to the NAS device and shared with other hosts. The NAS head includes the following components:

- CPU and memory
- One or more network interface cards (NICs), which provide connectivity to the client network. Examples of network protocols supported by NIC include Gigabit Ethernet, Fast Ethernet, ATM, and Fiber Distributed Data Interface (FDDI)
- An optimized operating system for managing the NAS functionality. It translates file-level requests into block-storage requests and further converts the data supplied at the block level to file data
- NFS, CIFS, and other protocols for file sharing
- Industry-standard storage protocols and ports to connect and manage physical disk resources

The NAS environment includes clients accessing a NAS device over an IP network using file-sharing protocols.

NAS File Sharing Protocols

- Two common NAS file sharing protocols are:
 - ▶ Common Internet File System (CIFS)
 - ▶ Network File System (NFS)

Most NAS devices support multiple file-service protocols to handle file I/O requests to a remote file system. As discussed earlier, NFS and CIFS are the common protocols for file sharing. NAS devices enable users to share file data across different operating environments and provide a means for users to migrate transparently from one operating system to another.

Common Internet File System

- Client-server application protocol
 - ▶ An open variation of the Server Message Block (SMB) protocol
- Enables clients to access files that are on a server over TCP/IP
- Stateful Protocol
 - ▶ Maintains connection information regarding every connected client
 - ▶ Can automatically restore connections and reopen files that were open prior to interruption

Common Internet File System (CIFS) is a client-server application protocol that enables client programs to make requests for files and services on remote computers over TCP/IP. It is a public, or open, variation of Server Message Block (SMB) protocol.

The CIFS protocol enables remote clients to gain access to files on a server. CIFS enables file sharing with other clients by using special locks. Filenames in CIFS are encoded using unicode characters. CIFS provides the following features to ensure data integrity:

- It uses file and record locking to prevent users from overwriting the work of another user on a file or a record.
- It supports fault tolerance and can automatically restore connections and reopen files that were open prior to an interruption. The fault tolerance features of CIFS depend on whether an application is written to take advantage of these features. Moreover, CIFS is a stateful protocol because the CIFS server maintains connection information regarding every connected client. If a network failure or CIFS server failure occurs, the client receives a disconnection notification. User disruption is minimized if the application has the embedded intelligence to restore the connection. However, if the embedded intelligence is missing, the user must take steps to reestablish the CIFS connection.

Users refer to remote file systems with an easy-to-use file-naming scheme:

`\\server\share` or `\\servername.domain.suffix\share`.

Network File System

- Client-server application protocol
- Enables clients to access files that are on a server
- Uses Remote Procedure Call (RPC) mechanism to provide access to remote file system
- Currently, three versions of NFS are in use:
 - ▶ NFS v2 is stateless and uses UDP as transport layer protocol
 - ▶ NFS v3 is stateless and uses UDP or optionally TCP as transport layer protocol
 - ▶ NFS v4 is stateful and uses TCP as transport layer protocol

Network File System (NFS) is a client-server protocol for file sharing that is commonly used on UNIX systems. NFS was originally based on the connectionless User Datagram Protocol (UDP). It uses a machine-independent model to represent user data. It also uses Remote Procedure Call (RPC) as a method of inter-process communication between two computers. The NFS protocol provides a set of RPCs to access a remote file system for the following operations:

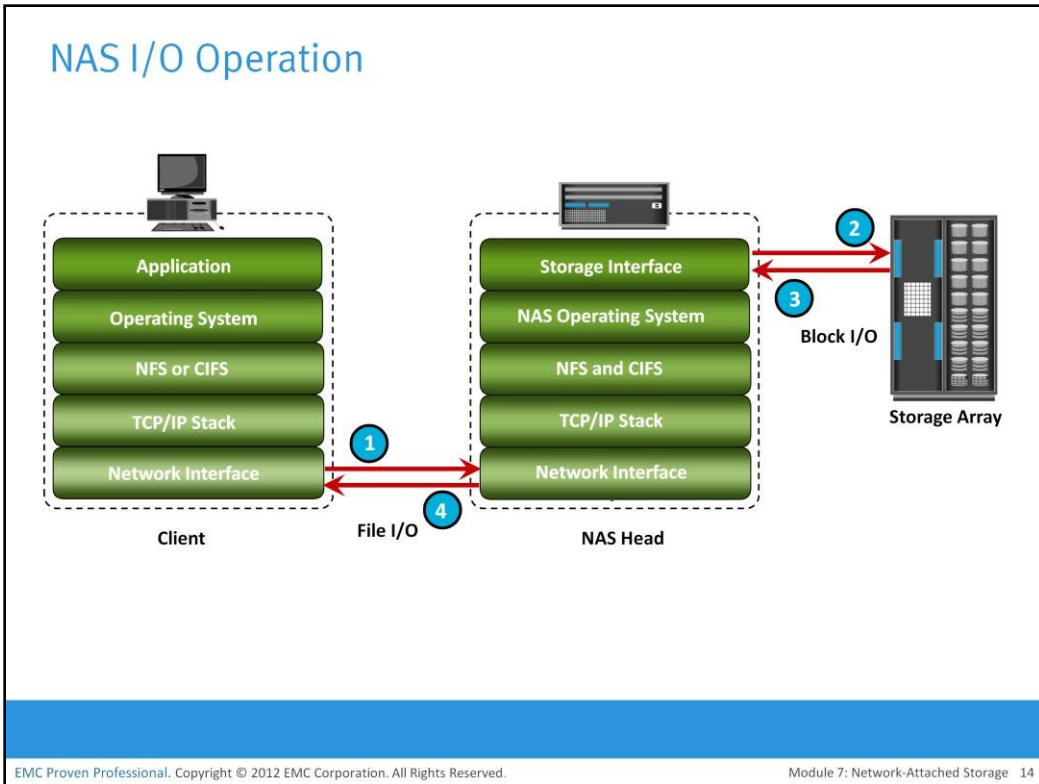
- Searching files and directories
- Opening, reading, writing to, and closing a file
- Changing file attributes
- Modifying file links and directories

NFS creates a connection between the client and the remote system to transfer data. NFS (NFSv3 and earlier) is a stateless protocol, which means that it does not maintain any kind of table to store information about open files and associated pointers. Therefore, each call provides a full set of arguments to access files on the server. These arguments include a file handle reference to the file, a particular position to read or write, and the versions of NFS.

Cont..

Currently, three versions of NFS are in use:

- NFS version 2 (NFSv2): Uses UDP to provide a stateless network connection between a client and a server. Features, such as locking, are handled outside the protocol.
- NFS version 3 (NFSv3): The most commonly used version, which uses UDP or TCP, and is based on the stateless protocol design. It includes some new features, such as a 64-bit file size, asynchronous writes, and additional file attributes to reduce refetching.
- NFS version 4 (NFSv4): Uses TCP and is based on a stateful protocol design. It offers enhanced security. The latest NFS version 4.1 is the enhancement of NFSv4 and includes some new features, such as session model, parallel NFS (pNFS), and data retention.



NAS provides file-level data access to its clients. File I/O is a high-level request that specifies the file to be accessed. For example, a client may request a file by specifying its name, location, or other attributes. The NAS operating system keeps track of the location of files on the disk volume and converts client file I/O into block-level I/O to retrieve data. The process of handling I/Os in a NAS environment is as follows:

1. The requestor (client) packages an I/O request into TCP/IP and forwards it through the network stack. The NAS device receives this request from the network.
2. The NAS device converts the I/O request into an appropriate physical storage request, which is a block-level I/O, and then performs the operation on the physical storage.
3. When the NAS device receives data from the storage, it processes and repackages the data into an appropriate file protocol response.
4. The NAS device packages this response into TCP/IP again and forwards it to the client through the network.

Module 7: Network-Attached Storage (NAS)

Lesson 2: NAS Implementation and File-level Virtualization

During this lesson the following topics are covered:

- NAS implementations
- NAS use cases
- File-level virtualization

This lesson describes three common NAS implementations: unified, gateway, and scale-out. It also covers server and storage consolidation use cases of NAS. Further it covers file-level virtualization and its benefits.

NAS Implementation – Unified NAS


- Consolidates NAS-based (file-level) and SAN-based (block-level) access on a single storage platform
- Supports both CIFS and NFS protocols for file access and iSCSI and FC protocols for block level access
- Provides unified management for both NAS head and storage

The unified NAS consolidates NAS-based and SAN-based data access within a unified storage platform and provides a unified management interface for managing both the environments.

Unified NAS performs file serving and storing of file data, along with providing access to block-level data. It supports both CIFS and NFS protocols for file access and iSCSI and FC protocols for block level access. Due to consolidation of NAS-based and SAN-based access on a single storage platform, unified NAS reduces an organization's infrastructure and management costs.

A unified NAS contains one or more NAS heads and storage in a single system. NAS heads are connected to the storage controllers (SCs), which provide access to the storage. These storage controllers also provide connectivity to iSCSI and FC hosts. The storage may consist of different drive types, such as SAS, ATA, FC, and flash drives, to meet different workload requirements.

Unified NAS Connectivity



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Module 7: Network-Attached Storage 17

Each NAS head in a unified NAS has front-end Ethernet ports, which connect to the IP network. The front-end ports provide connectivity to the clients and service the file I/O requests. Each NAS head has back-end ports, to provide connectivity to the storage controllers.

iSCSI and FC ports on a storage controller enable hosts to access the storage directly or through a storage network at the block level.

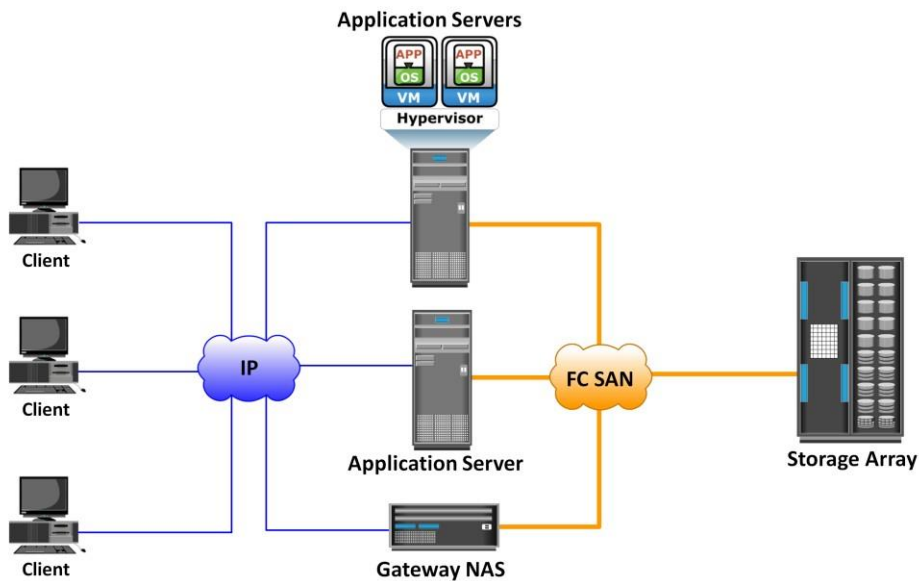
NAS Implementation – Gateway NAS

- Uses external and independently-managed storage
 - ▶ NAS heads access SAN-attached or direct-attached storage arrays
- NAS heads share storage with other application servers that perform block I/O
- Requires separate management of NAS head and storage

A gateway NAS device consists of one or more NAS heads and uses external and independently managed storage. Similar to unified NAS, the storage is shared with other applications that uses block-level I/O. Management functions in this type of solution are more complex than those in a unified NAS environment because there are separate administrative tasks for the NAS head and the storage. A gateway solution can use the FC infrastructure, such as switches and directors for accessing SAN-attached storage arrays or direct-attached storage arrays.

The gateway NAS is more scalable compared to unified NAS because NAS heads and storage arrays can be independently scaled up when required. For example, NAS heads can be added to scale up the NAS device performance. When the storage limit is reached, it can scale up, adding capacity on the SAN, independent of NAS heads. Similar to a unified NAS, a gateway NAS also enables high utilization of storage capacity by sharing it with the SAN environment.

Gateway NAS Connectivity



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Module 7: Network-Attached Storage 19

In a gateway solution, the front-end connectivity is similar to that in a unified storage solution. Communication between the NAS gateway and the storage system in a gateway solution is achieved through a traditional FC SAN. To deploy a gateway NAS solution, factors, such as multiple paths for data, redundant fabrics, and load distribution, must be considered.

NAS Implementation – Scale-out NAS


- Pools multiple nodes together in a cluster that works as a single NAS device
 - ▶ Pool is managed centrally
- Scales performance and/or capacity with addition of nodes to the pool non-disruptively
- Creates a single file system that runs on all nodes in the cluster
 - ▶ Clients, connected to any node, can access entire file system
 - ▶ File system grows dynamically as nodes are added
- Stripes data across all nodes in a pool along with mirror or parity protection

The scale-out NAS implementation pools multiple nodes together in a cluster. A node may consist of either the NAS head or storage or both. The cluster performs the NAS operation as a single entity.

A scale-out NAS provides the capability to scale its resources by simply adding nodes to a clustered NAS architecture. The cluster works as a single NAS device and is managed centrally. Nodes can be added to the cluster, when more performance or more capacity is needed, without causing any downtime. Scale-out NAS provides the flexibility to use many nodes of moderate performance and availability characteristics to produce a total system that has better aggregate performance and availability. It also provides ease of use, low cost, and theoretically unlimited scalability.

Scale-out NAS creates a single file system that runs on all nodes in the cluster. All information is shared among nodes, so the entire file system is accessible by clients connecting to any node in the cluster. Scale-out NAS stripes data across all nodes in a cluster along with mirror or parity protection. As data is sent from clients to the cluster, the data is divided and allocated to different nodes in parallel. When a client sends a request to read a file, the scale-out NAS retrieves the appropriate blocks from multiple nodes, recombines the blocks into a file, and presents the file to the client. As nodes are added, the file system grows dynamically and data is evenly distributed to every node. Each node added to the cluster increases the aggregate storage, memory, CPU, and network capacity. Hence, cluster performance also increases.

Scale-out NAS Connectivity



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

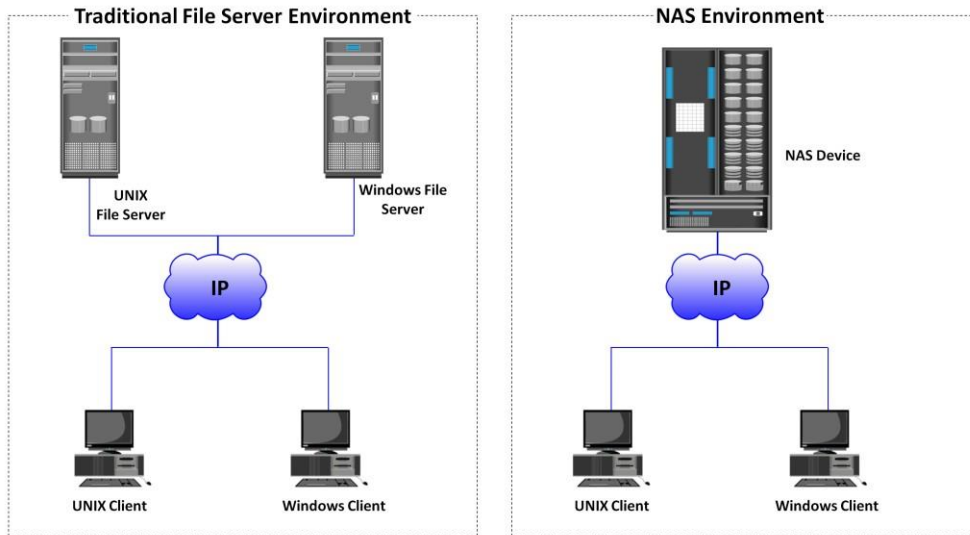
Module 7: Network-Attached Storage 21

Scale-out NAS clusters use separate internal and external networks for back-end and front-end connectivity, respectively. An internal network provides connections for intracluster communication, and an external network connection enables clients to access and share file data. Each node in the cluster connects to the internal network. The internal network offers high throughput and low latency and uses high-speed networking technology, such as InfiniBand or Gigabit Ethernet. To enable clients to access a node, the node must be connected to the external Ethernet network. Redundant internal or external networks may be used for high availability. Slide provides an example of scale-out NAS connectivity.

Note:

InfiniBand is a networking technology that provides a low-latency, high-bandwidth communication link between hosts and peripherals. It provides serial connection and is often used for inter-server communications in high-performance computing environments. InfiniBand enables remote direct memory access (RDMA) that enables a device (host or peripheral) to access data directly from the memory of a remote device. InfiniBand also enables a single physical link to carry multiple channels of data simultaneously using a multiplexing technique.

NAS Use Case 1 – Server Consolidation with NAS



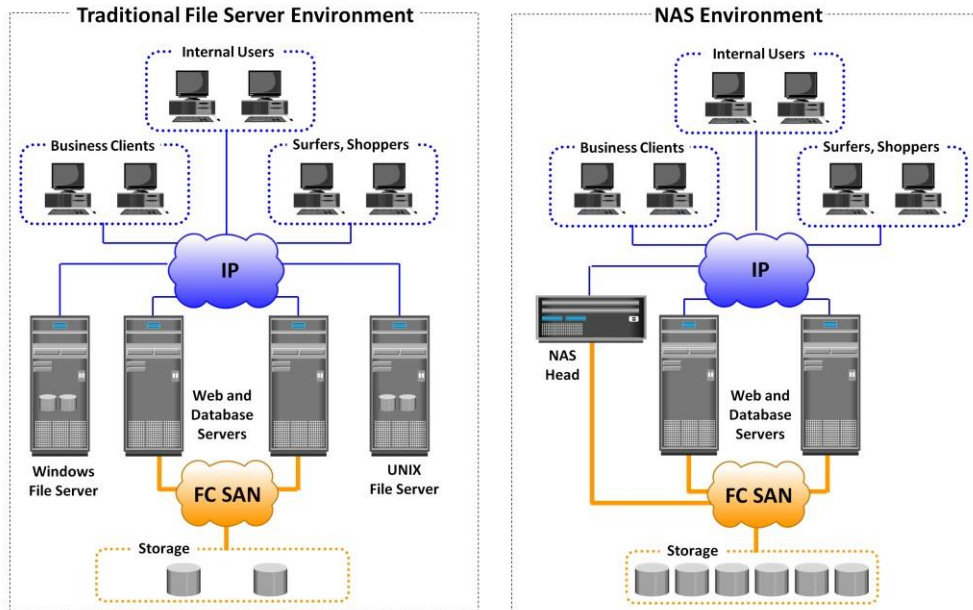
EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Module 7: Network-Attached Storage 22

The slide provides a use case that illustrates how a NAS enables consolidation of file servers. Traditionally, network file system for UNIX and Microsoft Windows are housed on separate servers. This requires maintenance of both the environments.

By implementation of NAS, both Windows and UNIX file structures can be housed together in a single system, while still maintaining their integrity. Using NAS, the same file system can be accessed via different protocols, either NFS or CIFS, and still maintain the integrity of the data and security structures, as long as the applications used for both methodologies understand the data structures presented.

NAS Use Case 2 – Storage Consolidation with NAS



EMC Proven Professional. Copyright © 2012 EMC Corporation. All Rights Reserved.

Module 7: Network-Attached Storage 23

The slide provides another use case that shows how storage resources in a traditional file server environment can be consolidated using NAS.


File-level Virtualization

- Eliminates dependency between data accessed at the file-level and the location where the files are physically stored
- Enables users to use a logical path, rather than a physical path, to access files
- Uses global namespace that maps logical path of file resources to their physical path
- Provides non-disruptive file mobility across file servers or NAS devices

A network-based file sharing environment is composed of multiple file servers or NAS devices. It might be required to move the files from one device to another due to reasons such as cost or performance. File-level virtualization, implemented in NAS or the file server environment, provides a simple, nondisruptive file-mobility solution.

File-level virtualization eliminates the dependencies between the data accessed at the file level and the location where the files are physically stored. It creates a logical pool of storage, enabling users to use a logical path, rather than a physical path, to access files. A global namespace is used to map the logical path of a file to the physical path names. File-level virtualization enables the movement of files across NAS devices, even if the files are being accessed.

Comparison: Before and After File-level Virtualization



- Dependency between client access and file location
- Underutilized storage resources
- Downtime is caused by data migrations

- Break dependencies between client access and file location
- Storage utilization is optimized
- Non-disruptive migrations

Before virtualization, each host knows exactly where its file resources are located. This environment leads to underutilized storage resources and capacity problems because files are bound to a specific NAS device or file server. It may be required to move the files from one server to another because of performance reasons or when the file server fills up. Moving files across the environment is not easy and may make files inaccessible during file movement. Moreover, hosts and applications need to be reconfigured to access the file at the new location. This makes it difficult for storage administrators to improve storage efficiency while maintaining the required service level.

File-level virtualization simplifies file mobility. It provides user or application independence from the location where the files are stored. File-level virtualization facilitates the movement of files across the online file servers or NAS devices. This means that while the files are being moved, clients can access their files nondisruptively. Clients can also read their files from the old location and write them back to the new location without realizing that the physical location has changed.

Module 7: Network-Attached Storage (NAS)

Concepts in Practice:

- EMC Isilon
- EMC VNX Gateway

The Concepts in Practice section covers EMC Isilon and VNX Gateway.

EMC Isilon

- Scale-out NAS solution
- Includes 'OneFS' operating system that creates a single file system across Isilon cluster
- Provides ability to nondisruptively add nodes to Isilon cluster
- Includes 'SmartPools' that enables different node types to be mixed in a single cluster
- Monitors component health and transparently reallocates files
- Uses 'Autobalance' that rebalances data automatically, when a new node is added to the cluster
- Uses 'FlexProtect' that protects from up to four simultaneous failures of either nodes or individual drives

EMC Isilon is the scale-out NAS solution. Isilon has a specialized operating system called OneFS that enables the scale-out NAS architecture. OneFS combines the three layers of traditional storage architectures—file system, volume manager, and RAID—into one unified software layer, creating a single file system that spans across all nodes in an Isilon cluster. It also provides the ability to seamlessly add storage and other resources without system downtime.

OneFS enables different node types to be mixed in a single cluster through the addition of the SmartPools application software. SmartPools enables deploying a single file system to span multiple nodes that have different performance characteristics and capacities. Isilon offers different types of nodes, such as the X-Series, S-Series, NL-Series, and Accelerator.

OneFS constantly monitors the health of all files and disks within a cluster, and if components are at risk, the file system automatically flags the problem components for replacement and transparently reallocates those files to healthy components.

When a new storage node is added, the Autobalance feature of OneFS automatically moves data onto this new node via the Infiniband based internal network. This automatic rebalancing ensures that the new node does not become a hot spot for new data.

OneFS includes a core technology, called FlexProtect, to provide data protection. FlexProtect provides protection for up to four simultaneous failures of either nodes or individual drives per stripe. FlexProtect provides file-specific protection capabilities. Different protection levels can be assigned to individual files, directories, or to portions of a file system.

EMC VNX Gateway

- Gateway NAS solution
- Provides multi-protocol network file system access, dynamic expansion of file systems, high availability, and high performance
- Comprises one or more NAS heads, called 'X-Blades' that run VNX operating environment
- Includes 'Control Station' that provides a single point for configuring X-Blades

The VNX Series Gateway contains one or more NAS heads, called X-Blades, that access external storage arrays, such as Symmetrix and block-based VNX via SAN. X-Blades run the VNX operating environment that is optimized for high-performance and multiprotocol network file system access. Each X-Blade consists of processors, redundant data paths, power supplies, Gigabit Ethernet, and 10-Gigabit Ethernet optical ports. All the X-Blades in a VNX gateway system are managed by Control Station, which provides a single point for configuring VNX Gateway. The VNX Gateway supports both pNFS and EMC patented Multi-Path File System (MPFS) protocols, which further improves the VNX Gateway performance. VNX Series Gateway offers two models: VG2 and VG8. VG8 supports up to eight X-Blades, whereas VG2 supports up to two. X-Blades may be configured as either primary or standby. A primary X-Blade is the operating NAS head, whereas a standby X-Blade becomes operational if the primary X-Blade fails. The Control Station handles an X-Blade failover. The Control Station also provides other high-availability features, such as fault monitoring, fault reporting, call home, and remote diagnostics.

Module 7: Summary

Key points covered in this module:

- NAS benefits
- NAS components
- NAS file sharing protocols
- NAS implementations
- File-level virtualization

This module covered the benefits and components of NAS. The key components of NAS are NAS head and storage. This module covered two common NAS file sharing protocols – CIFS and NFS – that enabled clients to access files located on a server. It also detailed on three common NAS implementations such as Unified, Gateway, and Scale-out that provides file sharing environment. Finally, it covered file-level virtualization that provides simple, nondisruptive file mobility solution.

Check Your Knowledge – 1

- Which component of a NAS head translates file-level requests into block-storage requests?
 - A. Front-end ports
 - B. Optimized operating system
 - C. CIFS and NFS
 - D. Network Interface Card
- Which is a feature of scale-out NAS?
 - A. Uses general purpose operating system for file serving
 - B. Creates multiple file systems on each node in the cluster
 - C. Uses external and independently-managed nodes
 - D. Enables pooling of nodes that work as a single NAS device

Check Your Knowledge – 2

- Which is a feature of gateway NAS?
 - A. Uses dedicated storage for each NAS head
 - B. NAS head and storage are managed independently
 - C. Creates a single file system that runs on all NAS heads
 - D. Provides connectivity to iSCSI and FC hosts

- Which NAS implementation consolidates file-based and block-based access on a single storage platform?
 - A. Scale-out
 - B. Gateway
 - C. Unified
 - D. Both gateway and scale-out

Check Your Knowledge – 3

- Which is a benefit of file-level virtualization?
 - A. Enables users to use physical path, rather than logical path, to access files
 - B. Translates file-level request into block-storage request non disruptively
 - C. Consolidates NAS-based and SAN-based access on a single storage platform
 - D. Eliminates dependency between data accessed at file level and file location

