IT8001
Information Storage and Management

Professional Elective

Topic

Network Attached Storage

Varun C.M.
Assistant Professor, Department of Information Technology,
St. Xavier's Catholic College of Engineering

Expected Outcome

CO2

Evaluate storage architectures, including storage subsystems, DAS, SAN, NAS, and CAS.

Network Attached Storage

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.

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.

It enables the movement of files across NAS devices, even if the files are being accessed.

Benefits of NAS

- Comprehensive access to information
- Improved efficiency
- Improved flexibility
- Centralized storage
- Simplified management
- Scalability
- High availability
- Security
- Low cost
- Ease of deployment

A file system is a structured way to store and organize data files. Many file systems maintain a file access table to simplify the process of searching and accessing files.

Accessing a File System

The operating system mounts a local file system during the boot process. The mount process creates a link between the file system on the NAS and the operating system on the client.

When mounting a file system, the operating system organizes files and directories in a tree-like structure and grants the privilege to the user to access this structure. The tree is rooted at a mount point. The mount point is named using operating system conventions.

Users and applications can traverse the entire tree from the root to the leaf nodes as file system permissions allow. Files are located at leaf nodes, and directories and subdirectories are located at intermediate roots.

Network File Sharing

Network file sharing refers to storing and accessing files over a network.
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

File Transfer Protocol (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.

Standard ***client-server file sharing protocols***, such as NFS and CIFS, enable the owner of a fi le 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 fi le servers.

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

Components of NAS

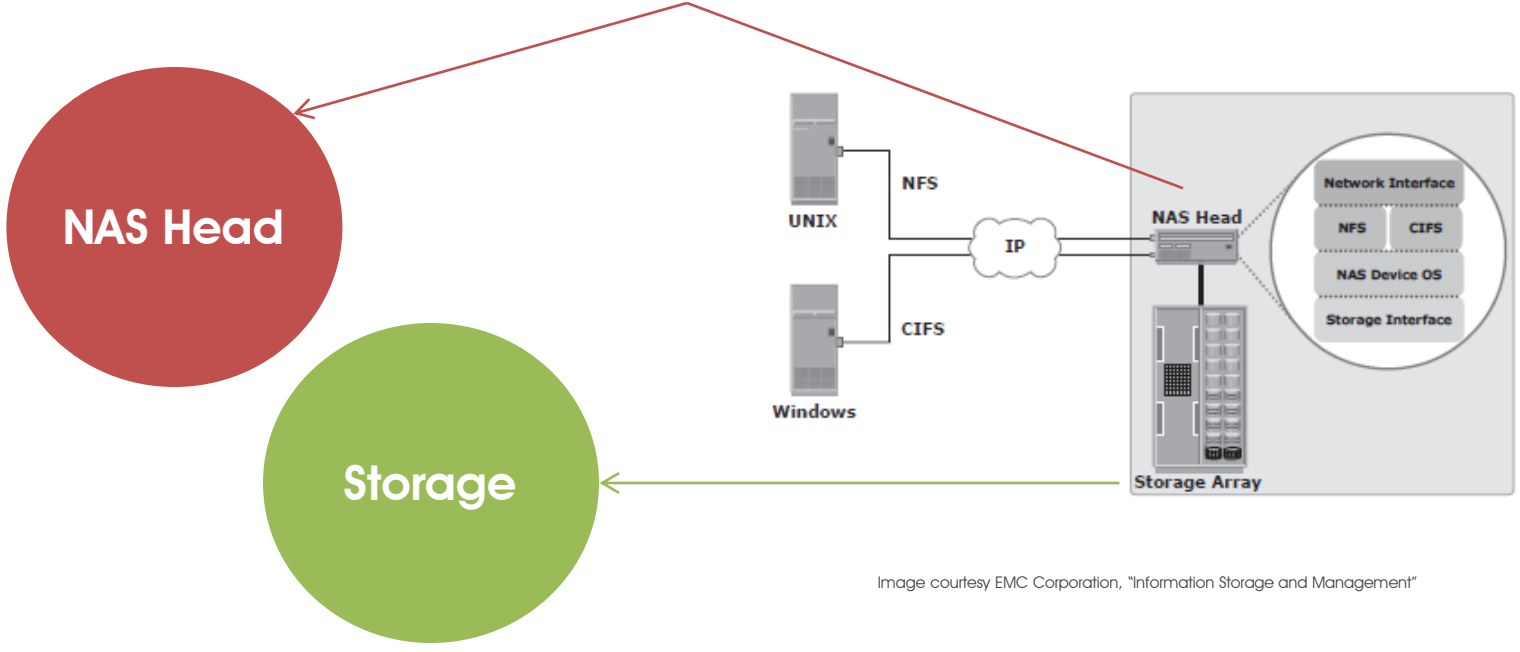

Image courtesy EMC Corporation, "Information Storage and Management"

Components of NAS contd...

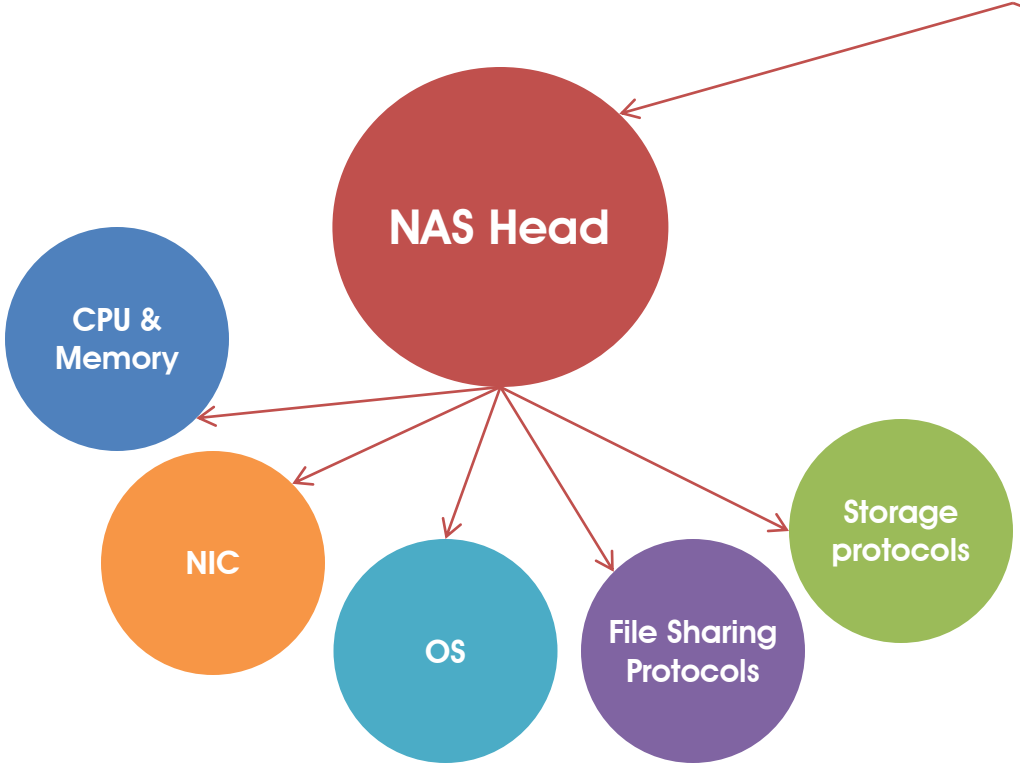

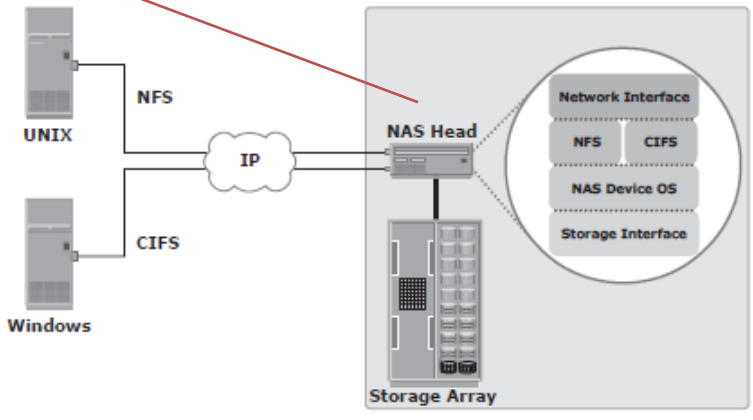

Image courtesy EMC Corporation, "Information Storage and Management"

Components of NAS contd...

Network Interface Cards (NIC)

- provides connectivity to the client network

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.

File Sharing Protocols

- NFS, CIFS, and other protocols for fi le sharing

Storage Protocols

- Industry-standard storage protocols and ports to connect and manage physical disk resources

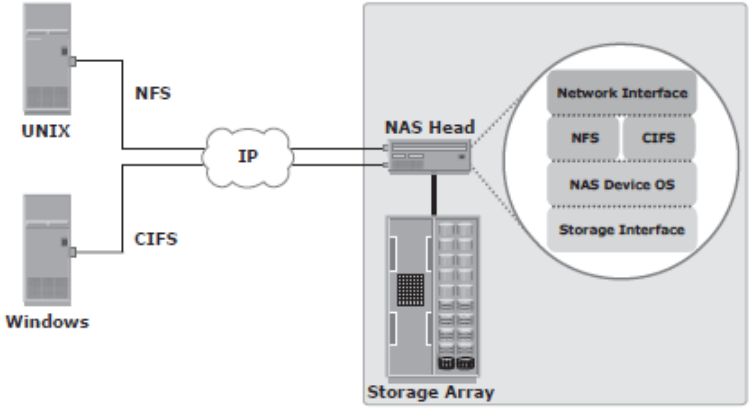

Image courtesy EMC Corporation, "Information Storage and Management"

NAS Implementations

NAS Implementations contd...

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 fi le access and iSCSI and FC protocols for block level access.
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.

Unified NAS Connectivity

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 Implementations contd...

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 use 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.

Gateway NAS Connectivity

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.

Implementation of both unified and gateway solutions requires analysis of the SAN environment. This analysis is required to determine the feasibility of combining the NAS workload with the SAN workload.

Analyze the SAN to determine whether the workload is primarily read or write, and if it is random or sequential. Also determine the predominant I/O size in use.

NAS Implementations contd...

Scale-out NAS enables grouping multiple nodes together to construct a clustered NAS system. 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.

Scale-Out NAS Connectivity

Scale-out NAS clusters use separate internal and external networks for back-end and front-end connectivity, respectively. An internal network provides connections for intra cluster 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. 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.

NAS File Sharing Protocols

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.

NFS

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.

NFS

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

NFS creates a connection between the client and the remote system to transfer data.
NFS 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.

NFS

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 re-fetching.

NFS version 4 (NFSv4)

Uses TCP and is based on a state-full 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.

CIFS

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

CIFS provides the following features to ensure data integrity:

- It uses fi le and record locking to prevent users from overwriting the work of another user on a fi le 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 state-full 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.

NAS I/O Operations

NAS provides file-level data access to its clients. File I/O is a high-level request that specifies the file to be accessed.
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.

NAS I/O Operations contd...

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 fi le protocol response.
4. The NAS device packages this response into TCP/IP again and forwards it to the client through the network.

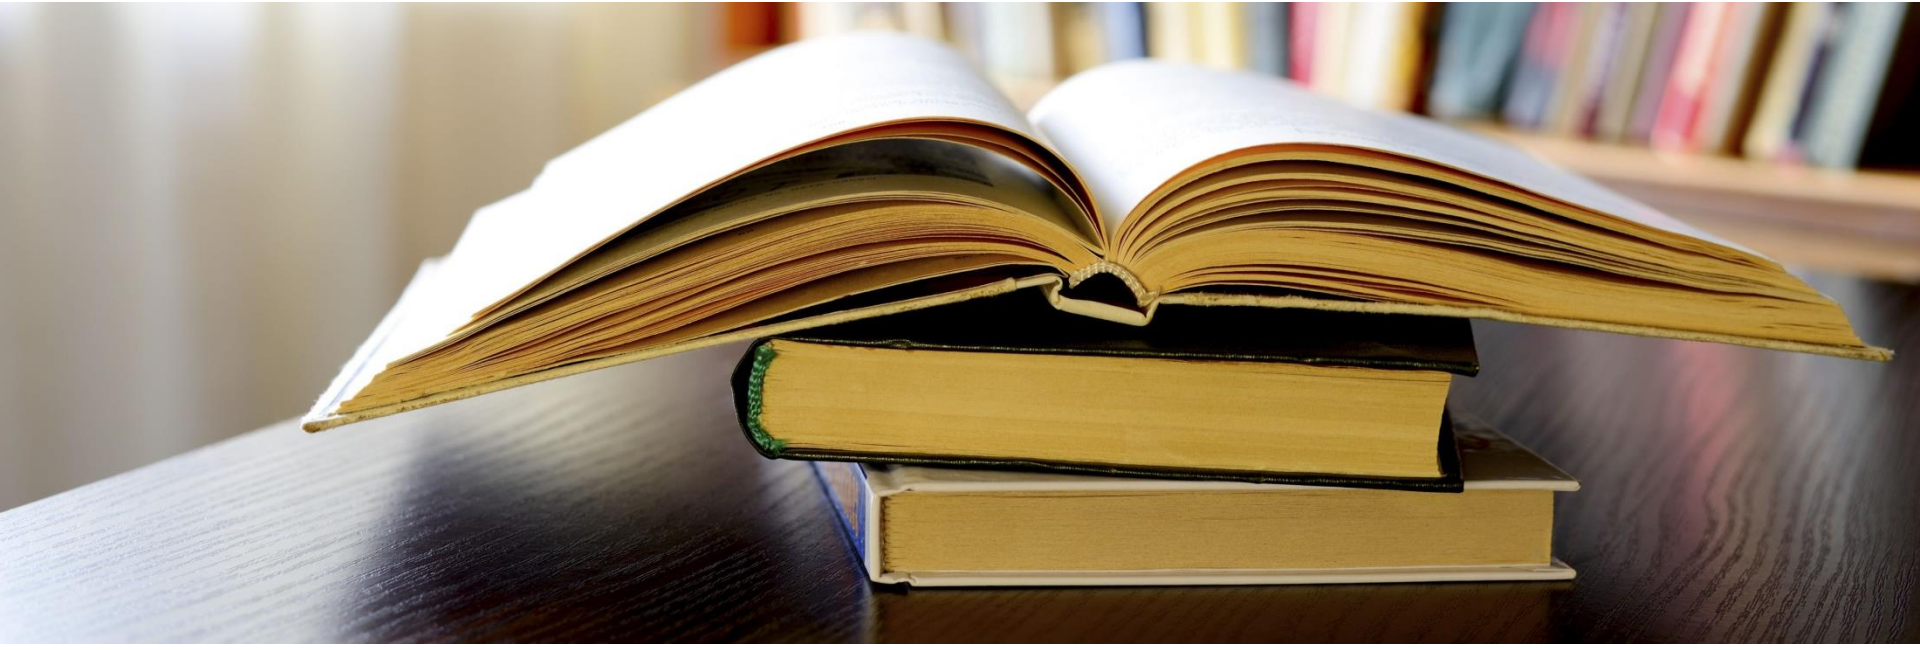

References

EMC Corporation, "Information Storage and Management" , Wiley, India.
Robert Spalding, "Storage Networks: The Complete Reference", Tata McGraw Hill, Osborne, 2003.

Varun C.M.
Assistant Professor, Department of Information Technology, St. Xavier's Catholic College of Engineering