



Section 2: Storage Networking Technologies

This section contains four modules:

- Module 5: Fibre Channel SAN (FC SAN)
- Module 6: IP SAN and FCoE
- Module 7: Network-Attached Storage (NAS)
- Module 8: Object-based and Unified Storage

Welcome to Section 2: Storage Networking Technologies. This section covers four modules, namely Fibre Channel SAN or FC SAN , IP SAN and FCoE, Network-Attached Storage or NAS , and Object-based and Unified Storage.



Module 5: Fibre Channel Storage Area Network (FC SAN)

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

- Describe FC SAN and its components
- Describe FC architecture
- Describe FC SAN topologies and zoning
- Describe virtualization in SAN environment

This module focuses on FC SAN components, FC interconnectivity options, and FC architecture. This module also focuses on virtualization in SAN environment.

Module 5: Fibre Channel Storage Area Network (FC SAN)

Lesson 1: Overview of FC SAN

During this lesson the following topics are covered:

- Evolution of FC SAN
- Components of FC SAN
- FC interconnectivity options
- FC port types

This lesson covers evolution of FC SAN, its components, and three FC interconnectivity options. This lesson also covers various FC port types.

Business Needs and Technology Challenges

- An effective information management solution must provide:
 - ▶ Just-in-time information to business users
 - ▶ Flexible and resilient storage infrastructure
- Information management challenges in DAS environment:
 - ▶ Explosive growth of information storage that remains isolated and underutilized
 - ▶ Proliferation of new servers and applications
 - ▶ Complexity in sharing storage resources across multiple servers
 - ▶ High cost of managing information
- Storage area network (SAN) addresses these challenges

Organizations are experiencing an explosive growth in information. This information needs to be stored, protected, optimized, and managed efficiently. Data center managers are burdened with the challenging task of providing low-cost, high-performance information management solutions. An effective information management solution must provide the following:

Just-in-time information to business users: Information must be available to business users when they need it. 24 x 7 data availability is becoming one of the key requirements of today's storage infrastructure. The explosive growth in storage, proliferation of new servers and applications, and the spread of mission-critical data throughout enterprises are some of the challenges that need to be addressed to provide information availability in real time.

Flexible and resilient storage infrastructure: The storage infrastructure must provide flexibility and resilience that aligns with changing business requirements. Storage should scale without compromising the performance requirements of applications and, at the same time, the total cost of managing information must be low.

Direct-attached storage (DAS) is often referred to as a stovepiped storage environment. Hosts "own" the storage, and it is difficult to manage and share resources on these isolated storage devices. Efforts to organize this dispersed data led to the emergence of the storage area network (SAN).

What is a SAN?

SAN

It is a high-speed, dedicated network of servers and shared storage devices.


- Centralizes storage and management
- Enables sharing of storage resources across multiple servers at block level
- Meets increasing storage demands efficiently with better economies of scale
- Common SAN deployments are:
 - ▶ Fibre Channel (FC) SAN: uses FC protocol for communication
 - ▶ IP SAN: uses IP-based protocols for communication

SAN is a high-speed, dedicated network of servers and shared storage devices. It enables storage consolidation and enables storage to be shared across multiple servers. This improves the utilization of storage resources compared to direct-attached storage architecture and reduces the total amount of storage an organization needs to purchase and manage. With consolidation, storage management becomes centralized and less complex, which further reduces the cost of managing information. SAN also enables organizations to connect geographically dispersed servers and storage. Further, it meets the storage demands efficiently with better economies of scale and also provides effective maintenance and protection of data.

Common SAN deployments are Fibre Channel (FC) SAN and IP SAN. Fibre Channel SAN uses Fibre Channel protocol for the transport of data, commands, and status information between servers (or hosts) and storage devices. IP SAN uses IP-based protocols for communication.

Understanding Fibre Channel

- High-speed network technology
 - ▶ Latest FC implementation supports speed up to 16 Gb/s
- Highly scalable
 - ▶ Theoretically, accommodate approximately 15 million devices



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

Section 2: Storage Networking Technologies

6

The FC architecture forms the fundamental construct of the FC SAN infrastructure. Fibre Channel is a high-speed network technology that runs on high-speed optical fiber cables and serial copper cables. The FC technology was developed to meet the demand for increased speeds of data transfer between servers and mass storage systems. Technical Committee T11, which is the committee within International Committee for Information Technology Standards (INCITS), is responsible for Fibre Channel interface standards.

High data transmission speed is an important feature of the FC networking technology. In comparison with Ultra-SCSI that is commonly used in DAS environments, FC is a significant leap in storage networking technology. The latest FC implementations of 16 GFC (Fibre Channel) offers a throughput of 3200 MB/s (raw bit rates of 16 Gb/s), whereas Ultra640 SCSI is available with a throughput of 640 MB/s. Credit-based flow control mechanism in FC delivers data as fast as the destination buffer is able to receive it, without dropping frames. Also FC has very little transmission overhead. The FC architecture is highly scalable, and theoretically, a single FC network can accommodate approximately 15 million devices.

Note: FibRE refers to the protocol, whereas fibER refers to a media.


Components of FC SAN

- Node (server and storage) ports
- Cables
- Connectors
- Interconnecting devices such as FC switches and hubs
- SAN management software

FC SAN is a network of servers and shared storage devices. Servers and storage are the end points or devices in the SAN (called 'nodes'). FC SAN infrastructure consists of node ports, cables, connectors, interconnecting devices (such as FC switches or hubs), along with SAN management software.

Node Ports

- Provide physical interface for communicating with other nodes
- Exist on
 - ▶ HBA in server
 - ▶ Front-end adapters in storage
- Each port has a transmit (Tx) link and a receive (Rx) link




In a Fibre Channel network, the end devices, such as hosts, storage arrays, and tape libraries, are all referred to as nodes. Each node is a source or destination of information. Each node requires one or more ports to provide a physical interface for communicating with other nodes. These ports are integral components of host adapters, such as HBA, and storage front-end controllers or adapters. In an FC environment a port operates in full-duplex data transmission mode with a transmit (Tx) link and a receive (Rx) link

Cables

- SAN implementation uses
 - ▶ Copper cables for short distance
 - ▶ Optical fiber cables for long distance
- Two types of optical cables: single-mode and multimode

Single-mode	Multimode
Carries single beam of light	Can carry multiple beams of light simultaneously
Distance up to 10km	Used for short distance (Modal dispersion weakens signal strength after certain distance)



SAN implementations use optical fiber cabling. Copper can be used for shorter distances for back-end connectivity because it provides acceptable signal-to-noise ratio for distances up to 30 meters. Optical fiber cables carry data in the form of light. There are two types of optical cables: multimode and single-mode. Multimode fiber (MMF) cable carries multiple beams of light projected at different angles simultaneously onto the core of the cable. Based on the bandwidth, multimode fibers are classified as OM1 (62.5µm core), OM2 (50µm core), and laser-optimized OM3 (50µm core). In an MMF transmission, multiple light beams traveling inside the cable tend to disperse and collide. This collision weakens the signal strength after it travels a certain distance—a process known as modal dispersion. An MMF cable is typically used for short distances because of signal degradation (attenuation) due to modal dispersion.

Single-mode fiber (SMF) carries a single ray of light projected at the center of the core. These cables are available in core diameters of 7 to 11 microns; the most common size is 9 microns. In an SMF transmission, a single light beam travels in a straight line through the core of the fiber. The small core and the single light wave help to limit modal dispersion. Among all types of fiber cables, single-mode provides minimum signal attenuation over maximum distance (up to 10 km). A single-mode cable is used for long-distance cable runs, and distance usually depends on the power of the laser at the transmitter and sensitivity of the receiver.

MMFs are generally used within data centers for shorter distance runs, whereas SMFs are used for longer distances.

Connectors

- Attached at the end of a cable
- Enable swift connection and disconnection of the cable to and from a port
- Commonly used connectors for fiber optic cables are:
 - ▶ Standard Connector (SC)
 - ▶▶ Duplex connectors
 - ▶ Lucent Connector (LC)
 - ▶▶ Duplex connectors
 - ▶ Straight Tip (ST)
 - ▶▶ Patch panel connectors
 - ▶▶ Simplex connectors



Standard Connector



Lucent Connector



Straight Tip Connector

A connector is attached at the end of a cable to enable swift connection and disconnection of the cable to and from a port. A Standard connector (SC) and a Lucent connector (LC) are two commonly used connectors for fiber optic cables. Straight Tip (ST) is another fiber-optic connector, which is often used with fiber patch panels.

Interconnecting Devices

- Commonly used interconnecting devices in FC SAN are:
 - ▶ Hubs, switches, and directors
- Hubs provide limited connectivity and scalability
- Switches and directors are intelligent devices
 - ▶ Switches are available with fixed port count or modular design
 - ▶ Directors are always modular, and its port count can be increased by inserting additional 'line cards' or 'blades'
 - ▶ High-end switches and directors contain redundant components

FC hubs, switches, and directors are the interconnect devices commonly used in FC SAN.

Hubs are used as communication devices in FC-AL implementations. Hubs physically connect nodes in a logical loop or a physical star topology. All the nodes must share the loop because data travels through all the connection points. Because of the availability of low-cost and high-performance switches, hubs are no longer used in FC SANs.

Switches are more intelligent than hubs and directly route data from one physical port to another. Therefore, nodes do not share the data path. Instead, each node has a dedicated communication path.

Directors are high-end switches with a higher port count and better fault-tolerance capabilities.

Switches are available with a fixed port count or with modular design. In a modular switch, the port count is increased by installing additional port cards to open slots. The architecture of a director is always modular, and its port count is increased by inserting additional line cards or blades to the director's chassis. High-end switches and directors contain redundant components to provide high availability. Both switches and directors have management ports (Ethernet or serial) for connectivity to SAN management servers.

SAN Management Software

- A suite of tools used in a SAN to manage interfaces between host and storage arrays
- Provides integrated management of SAN environment
- Enables web-based management using GUI or CLI



SAN management software manages the interfaces between hosts, interconnect devices, and storage arrays. The software provides a view of the SAN environment and enables management of various resources from one central console.

It provides key management functions, including mapping of storage devices, monitoring and generating alerts for discovered devices, and zoning (discussed later in the module).


FC Interconnectivity Options

- Point-to-Point
- Fibre Channel Arbitrated Loop (FC-AL)
- Fibre Channel Switched Fabric (FC-SW)

The FC architecture supports three basic interconnectivity options: point-to-point, fibre channel arbitrated loop (FC-AL), and fibre channel switched fabric (FC-SW).

Point-to-Point Connectivity


- Enables direct connection between nodes
- Offers limited connectivity and scalability
- Used in DAS environment



Point-to-point is the simplest FC configuration—two devices are connected directly to each other, as shown in the slide. This configuration provides a dedicated connection for data transmission between nodes. However, the point-to-point configuration offers limited connectivity, because only two devices can communicate with each other at a given time. Moreover, it cannot be scaled to accommodate a large number of nodes. Standard DAS uses point-to-point connectivity.

FC-AL Connectivity

- Provides shared loop to attached nodes
 - ▶ Nodes must arbitrate to gain control
- Implemented using ring or star topology
- Limitations of FC-AL
 - ▶ Only one device can perform I/O operation at a time
 - ▶ Supports up to 126 nodes
 - ▶ Addition or removal of a node causes momentary pause in loop traffic



In the FC-AL configuration, devices are attached to a shared loop. FC-AL has the characteristics of a token ring topology and a physical star topology. In FC-AL, each device contends with other devices to perform I/O operations. Devices on the loop must “arbitrate” to gain control of the loop. At any given time, only one device can perform I/O operations on the loop.

As a loop configuration, FC-AL can be implemented without any interconnecting devices by directly connecting one device to another two devices in a ring through cables.


However, FC-AL implementations may also use hubs whereby the arbitrated loop is physically connected in a star topology.

The FC-AL configuration has the following limitations in terms of scalability:

- FC-AL shares the loop and only one device can perform I/O operations at a time. Because each device in a loop must wait for its turn to process an I/O request, overall performance in FC-AL environment is low.
- FC-AL uses only 8-bits of 24-bit Fibre Channel addressing (the remaining 16-bits are masked) and enables the assignment of 127 valid addresses to the ports. Hence, it can support up to 127 devices on a loop. One address is reserved for optionally connecting the loop to an FC switch port. Therefore, up to 126 nodes can be connected to the loop.
- Adding or removing a device results in loop re-initialization, which can cause a momentary pause in loop traffic.

FC-SW Connectivity

- Creates a logical space (called fabric) in which all nodes communicate with one another using switches
 - ▶ Interswitch links (ISLs) enable switches to be connected together
- Provides dedicated path between nodes
- Addition/removal of node does not affect traffic of other nodes



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

Section 2: Storage Networking Technologies

16


FC-SW is also referred to as fabric connect. A fabric is a logical space in which all nodes communicate with one another in a network. This virtual space can be created with a switch or a network of switches. Each switch in a fabric contains a unique domain identifier, which is part of the fabric's addressing scheme. In FC-SW, nodes do not share a loop; instead, data is transferred through a dedicated path between the nodes. Each port in a fabric has a unique 24-bit Fibre Channel address for communication.

In a switched fabric, the link between any two switches is called an interswitch link (ISL). ISLs enable switches to be connected together to form a single, larger fabric. ISLs are used to transfer host-to-storage data and fabric management traffic from one switch to another. By using ISLs, a switched fabric can be expanded to connect a large number of nodes.

FC-SW uses switches that are intelligent devices. They can switch data traffic between nodes directly through switch ports. Frames are routed between source and destination by the fabric.

Unlike a loop configuration, a FC-SW network provides dedicated path and scalability. The addition or removal of a device in a switched fabric is minimally disruptive; it does not affect the ongoing traffic between other devices.

Port Types in Switched Fabric



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

Section 2: Storage Networking Technologies

17

Ports in a switched fabric can be one of the following types:

- **N_Port:** An end point in the fabric. This port is also known as the node port. Typically, it is a host port (HBA) or a storage array port that is connected to a switch in a switched fabric.
- **E_Port:** A port that forms the connection between two FC switches. This port is also known as the expansion port. The E_Port on an FC switch connects to the E_Port of another FC switch in the fabric ISLs.
- **F_Port:** A port on a switch that connects an N_Port. It is also known as a fabric port.
- **G_Port:** A generic port on a switch that can operate as an E_Port or an F_Port and determines its functionality automatically during initialization.

Module 5: Fibre Channel Storage Area Network (FC SAN)

Lesson 2: Fibre Channel (FC) Architecture

During this lesson the following topics are covered:

- FC protocol stack
- FC addressing
- WWN addressing
- Structure and organization of FC data
- Fabric services
- Fabric login types

This lesson covers FC protocol stack, FC and WWN addressing, and structure and organization of FC data. This lesson also covers fabric services and login types.

FC Architecture Overview

- Provides benefits of both channel and network technologies
 - ▶ Provides high performance with low protocol overheads
 - ▶ Provides high scalability with long distance capability
- Implements SCSI over FC network
 - ▶ Transports SCSI data through FC network
- Storage devices, attached to SAN, appear as local storage devices to host operating system

Traditionally, host computer operating systems have communicated with peripheral devices over channel connections, such as ESCON and SCSI. Channel technologies provide high levels of performance with low protocol overheads. Such performance is achievable due to the static nature of channels and the high level of hardware and software integration provided by the channel technologies. However, these technologies suffer from inherent limitations in terms of the number of devices that can be connected and the distance between these devices.


In contrast to channel technology, network technologies are more flexible and provide greater distance capabilities. Network connectivity provides greater scalability and uses shared bandwidth for communication. This flexibility results in greater protocol overhead and reduced performance.

The FC architecture represents true channel/network integration and captures some of the benefits of both channel and network technology. FC SAN uses the Fibre Channel Protocol (FCP) that provides both channel speed for data transfer with low protocol overhead and scalability of network technology.

FCP forms the fundamental construct of the FC SAN infrastructure. Fibre Channel provides a serial data transfer interface that operates over copper wire and optical fiber. FCP is the implementation of SCSI over an FC network. In FCP architecture, all external and remote storage devices attached to the SAN appear as local devices to the host operating system. The key advantages of FCP are as follows:

- Sustained transmission bandwidth over long distances.
- Support for a larger number of addressable devices over a network. Theoretically, FC can support more than 15 million device addresses on a network.
- Support speeds up to 16 Gbps (16 GFC).

Fibre Channel Protocol Stack



FC Layer	Function	Features Specified by FC Layer
FC-4	Mapping interface	Mapping upper layer protocol (e.g. SCSI) to lower FC layers
FC-3	Common services	Not implemented
FC-2	Routing, flow control	Frame structure, FC addressing, flow control
FC-1	Encode/decode	8b/10b or 64b/66b encoding, bit and frame synchronization
FC-0	Physical layer	Media, cables, connector

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

Section 2: Storage Networking Technologies

20

It is easier to understand a communication protocol by viewing it as a structure of independent layers. FCP defines the communication protocol in five layers: FC-0 through FC-4 (except FC-3 layer, which is not implemented).

FC-4 Layer: It is the uppermost layer in the FCP stack. This layer defines the application interfaces and the way Upper Layer Protocols (ULPs) are mapped to the lower FC layers. The FC standard defines several protocols that can operate on the FC-4 layer. Some of the protocols include SCSI, High Performance Parallel Interface (HIPPI) Framing Protocol, Enterprise Storage Connectivity (ESCON), Asynchronous Transfer Mode (ATM), and IP.

FC-2 Layer: It provides Fibre Channel addressing, structure, and organization of data (frames, sequences, and exchanges). It also defines fabric services, classes of service, flow control, and routing.

FC-1 Layer: It defines how data is encoded prior to transmission and decoded upon receipt. At the transmitter node, an 8-bit character is encoded into a 10-bit transmission character. This character is then transmitted to the receiver node. At the receiver node, the 10-bit character is passed to the FC-1 layer, which decodes the 10-bit character into the original 8-bit character. FC links with speed 10 Gbps and above use 64-bit to 66-bit encoding algorithm. This layer also defines the transmission words such as FC frame delimiters, which identify the start and end of a frame and primitive signals that indicate events at a transmitting port. In addition to these, the FC-1 layer performs link initialization and error recovery.

FC-0 Layer : It is the lowest layer in the FCP stack. This layer defines the physical interface, media, and transmission of bits. The FC-0 specification includes cables, connectors, and optical and electrical parameters for a variety of data rates. The FC transmission can use both electrical and optical media.

FC Addressing in Switched Fabric

- FC Address is assigned to nodes during fabric login
 - ▶ Used for communication between nodes within FC SAN
- Address format



- Domain ID is a unique number provided to each switch in the fabric
 - ▶ 239 addresses are available for domain ID
- Maximum possible number of node ports in a switched fabric:
 - ▶ 239 domains X 256 areas X 256 ports = 15,663,104

An FC address is dynamically assigned when a node port logs on to the fabric. The FC address has a distinct format, as shown in the slide. The first field of the FC address contains the domain ID of the switch. A Domain ID is a unique number provided to each switch in the fabric. Although this is an 8-bit field, there are only 239 available addresses for domain ID because some addresses are deemed special and reserved for fabric management services. For example, FFFFFC is reserved for the name server, and FFFFFE is reserved for the fabric login service. The area ID is used to identify a group of switch ports used for connecting nodes. An example of a group of ports with common area ID is a port card on the switch. The last field, the port ID, identifies the port within the group.

Therefore, the maximum possible number of node ports in a switched fabric is calculated as:
 $239 \text{ domains} \times 256 \text{ areas} \times 256 \text{ ports} = 15,663,104$

World Wide Name (WWN)

- Unique 64 bit identifier
- Static to node ports on an FC network
 - ▶ Similar to MAC address of NIC
 - ▶ WWNN and WWPN are used to uniquely identify nodes and ports respectively

World Wide Name - Array															
5	0	0	6	0	1	6	0	0	0	6	0	0	1	B	2
0101	0000	0000	0110	0000	0001	0110	0000	0000	0000	0110	0000	0000	0001	1011	0010
Format Type	Company ID 24 bits						Port	Model Seed 32 bits							

World Wide Name - HBA															
1	0	0	0	0	0	0	0	c	9	2	0	d	c	4	0
Format Type	Reserved 12 bits				Company ID 24 bits						Company Specific 24 bits				

Each device in the FC environment is assigned a 64-bit unique identifier called the World Wide Name (WWN). The Fibre Channel environment uses two types of WWNs: World Wide Node Name (WWNN) and World Wide Port Name (WWPN). Unlike an FC address, which is assigned dynamically, a WWN is a static name for each node on an FC network. WWNs are similar to the Media Access Control (MAC) addresses used in IP networking. WWNs are burned into the hardware or assigned through software. Several configuration definitions in a SAN use WWN for identifying storage devices and HBAs. The name server in an FC environment keeps the association of WWNs to the dynamically created FC addresses for nodes. Figure in the slide illustrates the WWN structure examples for an array and an HBA.

Structure and Organization of FC Data

- FC data is organized as Exchange, Sequence, and Frame

FC Data Structure	Description
Exchange	- Enables two N_Ports to identify and manage a set of information units - Information unit: upper layer protocol-specific information that is sent to another port to perform certain operation - Each information unit maps to a sequence - Includes one or more sequences
Sequence	Contiguous set of frames that correspond to an information unit
Frame	- Fundamental unit of data transfer - Each frame consists of five parts: SOF, frame header, data field, CRC, and EOF



FC Frame

In an FC network, data transport is analogous to a conversation between two people, whereby a frame represents a word, a sequence represents a sentence, and an exchange represents a conversation.

Exchange: An exchange operation enables two node ports to identify and manage a set of information units. Each upper layer protocol has its protocol-specific information that must be sent to another port to perform certain operations. This protocol-specific information is called an information unit. The structure of these information units is defined in the FC-4 layer. This unit maps to a sequence. An exchange is composed of one or more sequences.

Sequence: A sequence refers to a contiguous set of frames that are sent from one port to another. A sequence corresponds to an information unit, as defined by the ULP.

Frame: A frame is the fundamental unit of data transfer at Layer 2. An FC frame consists of five parts: start of frame (SOF), frame header, data field, cyclic redundancy check (CRC), and end of frame (EOF). The SOF and EOF act as delimiters. The frame header is 24 bytes long and contains addressing information for the frame. The data field in an FC frame contains the data payload, up to 2,112 bytes of actual data—in most cases, the SCSI data. The CRC checksum facilitates error detection for the content of the frame. This checksum verifies data integrity by checking whether the content of the frames was received correctly. The CRC checksum is calculated by the sender before encoding at the FC-1 layer. Similarly, it is calculated by the receiver after decoding at the FC-1 layer.

Fabric Services

- FC switches provide fabric services as defined in FC standards

Fabric Services	Description
Fabric Login Server	• Used during the initial part of the node's fabric login process • Located at pre-defined address of FFFFFE
Name Server	• Responsible for name registration and management of node ports • Located at pre-defined address FFFFFC
Fabric Controller	• Responsible for managing and distributing Registered State Change Notifications (RSCNs) to attached node ports • Responsible for distributing SW-RSCNs to every other switch • SW-RSCNs keep the name server up-to-date on all switches • Located at pre-defined address FFFFFD
Management Server	• Enables FC SAN management using fabric management software • Located at pre-defined address FFFFA

All FC switches, regardless of the manufacturer, provide a common set of services as defined in the Fibre Channel standards. These services are available at certain predefined addresses. Some of these services are Fabric Login Server, Fabric Controller, Name Server, and Management Server.

The Fabric Login Server is located at the predefined address of FFFFFE and is used during the initial part of the node's fabric login process.

The Name Server (formally known as Distributed Name Server) is located at the predefined address FFFFFC and is responsible for name registration and management of node ports. Each switch exchanges its Name Server information with other switches in the fabric to maintain a synchronized, distributed name service.

Each switch has a Fabric Controller located at the predefined address FFFFFD. The Fabric Controller provides services to both node ports and other switches. The Fabric Controller is responsible for managing and distributing Registered State Change Notifications (RSCNs) to the node ports registered with the Fabric Controller. If there is a change in the fabric, RSCNs are sent out by a switch to the attached node ports. The Fabric Controller also generates Switch Registered State Change Notifications (SW-RSCNs) to every other domain (switch) in the fabric. These RSCNs keep the name server up-to-date on all switches in the fabric.

FFFFFA is the Fibre Channel address for the Management Server. The Management Server is distributed to every switch within the fabric. The Management Server enables the FC SAN management software to retrieve information and administer the fabric.

Login Types in Switched Fabric

- Fabric login (FLOGI)
 - ▶ Occurs between an N_Port and an F_Port
 - ▶ Node sends a FLOGI frame with WWN to Fabric Login Server on switch
 - ▶ Node obtains FC address from switch
 - ▶ Immediately after FLOGI, N_Port registers with Name Server on switch, indicating its WWN, port type, assigned FC address, etc.
 - ▶ N_Port queries name server about all other logged in ports
- Port login (PLOGI)
 - ▶ Occurs between two N_Ports to establish a session
 - ▶ Exchange service parameters relevant to the session
- Process login (PRLI)
 - ▶ Occurs between two N_Ports to exchange ULP related parameters

Fabric services define three login types:

- Fabric login (FLOGI): Performed between an N_Port and an F_Port. To log on to the fabric, a node sends a FLOGI frame with the WWNN and WWPNN parameters to the login service at the predefined FC address FFFFFE (Fabric Login Server). In turn, the switch accepts the login and returns an Accept (ACC) frame with the assigned FC address for the node. Immediately after the FLOGI, the N_Port registers itself with the local Name Server on the switch, indicating its WWNN, WWPNN, port type, class of service, assigned FC address and so on. After the N_Port has logged in, it can query the name server database for information about all other logged in ports.
- Port login (PLOGI): Performed between two N_Ports to establish a session. The initiator N_Port sends a PLOGI request frame to the target N_Port, which accepts it. The target N_Port returns an ACC to the initiator N_Port. Next, the N_Ports exchange service parameters relevant to the session.
- Process login (PRLI): Also performed between two N_Ports. This login relates to the FC-4 ULPs, such as SCSI. If the ULP is SCSI, N_Ports exchange SCSI-related service parameters.

Module 5: Fibre Channel Storage Area Network (FC SAN)

Lesson 3: FC SAN Topologies and Zoning


During this lesson the following topics are covered:

- Mesh and core-edge topologies
- Benefits of zoning
- Types of zoning

This lesson covers FC SAN topologies such as mesh and core-edge. This lesson also covers zoning and its benefits, components, and types.

Mesh Topology

- Full mesh
 - ▶ Each switch is connected to every other switch
 - ▶ Maximum of one ISL or hop is required for host-to-storage traffic
 - ▶ Host and storage can be connected to any switch
- Partial mesh
 - ▶ Not all the switches are connected to every other switch




A mesh topology may be one of the two types: full mesh or partial mesh. In a full mesh, every switch is connected to every other switch in the topology. A full mesh topology may be appropriate when the number of switches involved is small. A typical deployment would involve up to four switches or directors, with each of them servicing highly localized host-to-storage traffic. In a full mesh topology, a maximum of one ISL or hop is required for host-to-storage traffic. However, with the increase in the number of switches, the number of switch ports used for ISL also increases. This reduces the available switch ports for node connectivity.

In a partial mesh topology, several hops or ISLs may be required for the traffic to reach its destination. Partial mesh offers more scalability than full mesh topology. However, without proper placement of host and storage devices, traffic management in a partial mesh fabric might be complicated and ISLs could become overloaded due to excessive traffic aggregation.

Core-edge Topology

- Consists of edge and core switch tiers
- Network traffic traverses core tier or terminate at core tier
- Storage is usually connected to the core tier
- Benefits
 - ▶ High availability
 - ▶ Medium scalability
 - ▶ Medium to maximum connectivity



The core-edge fabric topology has two types of switch tiers. The edge tier is usually composed of switches and offers an inexpensive approach to adding more hosts in a fabric. Each switch at the edge tier is attached to a switch at the core tier through ISLs.

The core tier is usually composed of directors that ensure high fabric availability. In addition, typically all traffic must either traverse this tier or terminate at this tier. In this configuration, all storage devices are connected to the core tier, enabling host-to-storage traffic to traverse only one ISL. Hosts that require high performance may be connected directly to the core tier and consequently avoid ISL delays.


In core-edge topology, the edge-tier switches are not connected to each other. The core-edge fabric topology increases connectivity within the SAN while conserving the overall port utilization. If fabric expansion is required, additional edge switches are connected to the core. The core of the fabric is also extended by adding more switches or directors at the core tier. Based on the number of core-tier switches, this topology has different variations, such as, single-core topology and dual-core topology. To transform a single-core topology to dual-core, new ISLs are created to connect each edge switch to the new core switch in the fabric.

Zoning

Zoning

It is an FC switch function that enables node ports within the fabric to be logically segmented into groups, and communicate with each other within the group.

- Zone set comprises zones
- Each zone comprises zone members (HBA and array ports)
- Benefits
 - ▶ Restricts RSCN traffic
 - ▶ Provides access control




Zoning is an FC switch function that enables node ports within the fabric to be logically segmented into groups and communicate with each other within the group.

Whenever a change takes place in the name server database, the fabric controller sends a Registered State Change Notification (RSCN) to all the nodes impacted by the change. If zoning is not configured, the fabric controller sends an RSCN to all the nodes in the fabric. Involving the nodes that are not impacted by the change results in increased fabric-management traffic. For a large fabric, the amount of FC traffic generated due to this process can be significant and might impact the host-to-storage data traffic. Zoning helps to limit the number of RSCNs in a fabric. In the presence of zoning, a fabric sends the RSCN to only those nodes in a zone where the change has occurred.

Zoning also provides access control, along with other access control mechanisms, such as LUN masking. Zoning provides control by allowing only the members in the same zone to establish communication with each other.

Zone members, zones, and zone sets form the hierarchy defined in the zoning process. A zone set is composed of a group of zones that can be activated or deactivated as a single entity in a fabric. Multiple zone sets may be defined in a fabric, but only one zone set can be active at a time. Members are nodes within the SAN that can be included in a zone. Switch ports, HBA ports, and storage device ports can be members of a zone. A port or node can be a member of multiple zones. Nodes distributed across multiple switches in a switched fabric may also be grouped into the same zone. Zone sets are also referred to as zone configurations.

Types of Zoning



Zoning can be categorized into three types:

- **Port zoning:** Uses the physical address of switch ports to define zones. In port zoning, access to node is determined by the physical switch port to which a node is connected. The zone members are the port identifier (switch domain ID and port number) to which HBA and its targets (storage devices) are connected. If a node is moved to another switch port in the fabric, port zoning must be modified to allow the node, in its new port, to participate in its original zone. However, if an HBA or storage device port fails, an administrator just has to replace the failed device without changing the zoning configuration.
- **WWN zoning:** Uses World Wide Names to define zones. The zone members are the unique WWN addresses of the HBA and its targets (storage devices). A major advantage of WWN zoning is its flexibility. WWN zoning allows nodes to be moved to another switch port in the fabric and maintain connectivity to its zone partners without having to modify the zone configuration. This is possible because the WWN is static to the node port.
- **Mixed zoning:** Combines the qualities of both WWN zoning and port zoning. Using mixed zoning enables a specific node port to be tied to the WWN of another node.

Figure in the slide shows the three types of zoning on an FC network.

Module 5: Fibre Channel Storage Area Network (FC SAN)

Lesson 4: Virtualization in SAN


During this lesson the following topics are covered:

- Block-level storage virtualization
- Virtual SAN

This lesson covers block-level storage virtualization and virtual SAN.

Block-level Storage Virtualization

- Provides a virtualization layer in SAN
- Abstracts block storage devices and creates a storage pool by aggregating LUNs
- Virtual volumes are created from storage pool and assigned to hosts
 - ▶ Virtualization layer maps virtual volumes to LUNs
- Benefits
 - ▶ Online volume expansion
 - ▶ Nondisruptive migration



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

Section 2: Storage Networking Technologies

32

Block-level storage virtualization aggregates block storage devices (LUNs) and enables provisioning of virtual storage volumes, independent of the underlying physical storage. A virtualization layer, which exists at the SAN, abstracts the identity of physical storage devices and creates a storage pool from heterogeneous storage devices. Virtual volumes are created from the storage pool and assigned to the hosts. Instead of being directed to the LUNs on the individual storage arrays, the hosts are directed to the virtual volumes provided by the virtualization layer. For hosts and storage arrays, the virtualization layer appears as the target and initiator devices, respectively. The virtualization layer maps the virtual volumes to the LUNs on the individual arrays. The hosts remain unaware of the mapping operation and access the virtual volumes as if they were accessing the physical storage attached to them. Typically, the virtualization layer is managed via a dedicated virtualization appliance to which the hosts and the storage arrays are connected.


Figure in the slide illustrates a virtualized environment. It shows two physical servers, each of which has one virtual volume assigned. These virtual volumes are used by the servers. These virtual volumes are mapped to the LUNs in the storage arrays. When an I/O is sent to a virtual volume, it is redirected through the virtualization layer at the storage network to the mapped LUNs. Depending on the capabilities of the virtualization appliance, the architecture may allow for more complex mapping between array LUNs and virtual volumes.

cont..

Block-level storage virtualization enables extending the storage volumes online to meet application growth requirements. It consolidates heterogeneous storage arrays and enables transparent volume access.

Block-level storage virtualization also provides the advantage of nondisruptive data migration. In a traditional SAN environment, LUN migration from one array to another is an offline event because the hosts needed to be updated to reflect the new array configuration. In other instances, host CPU cycles were required to migrate data from one array to the other, especially in a multivendor environment. With a block-level virtualization solution in place, the virtualization layer handles the back-end migration of data, which enables LUNs to remain online and accessible while data is migrating. No physical changes are required because the host still points to the same virtual targets on the virtualization layer. However, the mappings information on the virtualization layer should be changed. These changes can be executed dynamically and are transparent to the end user.

Use Case: Block-level Storage Virtualization across Data Centers



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

Section 2: Storage Networking Technologies

34


Previously, block-level storage virtualization provided nondisruptive data migration only within a data center. The new generation of block-level storage virtualization enables nondisruptive data migration both within and between data centers. It provides the capability to connect the virtualization layers at multiple data centers. The connected virtualization layers are managed centrally and work as a single virtualization layer stretched across data centers. This enables the federation of block-storage resources both within and across data centers. The virtual volumes are created from the federated storage resources.

Virtual SAN (VSAN)/Virtual Fabric

VSAN

It is a logical fabric on an FC SAN, enabling communication among a group of nodes, regardless of their physical location in the fabric.

- Each VSAN has its own fabric services (name server, zoning), configuration, and set of FC addresses
- VSANs improve SAN security, scalability, availability, and manageability



Virtual SAN (also called virtual fabric) is a logical fabric on an FC SAN, which enables communication among a group of nodes regardless of their physical location in the fabric. In a VSAN, a group of hosts or storage ports communicate with each other using a virtual topology defined on the physical SAN. Multiple VSANs may be created on a single physical SAN. Each VSAN acts as an independent fabric with its own set of fabric services, such as name server, and zoning. Fabric-related configurations in one VSAN do not affect the traffic in another.

VSANs improve SAN security, scalability, availability, and manageability. VSANs provide enhanced security by isolating the sensitive data in a VSAN and by restricting access to the resources located within that VSAN. The same Fibre Channel address can be assigned to nodes in different VSANs, thus increasing the fabric scalability. Events causing traffic disruptions in one VSAN are contained within that VSAN and are not propagated to other VSANs. VSANs facilitate an easy, flexible, and less expensive way to manage networks. Configuring VSANs is easier and quicker compared to building separate physical FC SANs for various node groups. To regroup nodes, an administrator simply changes the VSAN configurations without moving nodes and recabling.

Module 5: Fibre Channel Storage Area Network (FC SAN)

Concept in Practice:

- EMC Connectrix
- EMC VPLEX

The Concept in Practice section covers EMC Connectrix and VPLEX.

EMC Connectrix

- Connectrix family includes networked storage connectivity products
 - ▶ Offers high-speed FC connectivity, highly resilient switching technology, intelligent IP storage networking, and I/O consolidation with Fibre Channel over Ethernet
- Connectrix family consist of enterprise directors, departmental switches, and multi-purpose switches



Enterprise Director



Departmental Switch



Multi-purpose Switch

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

Section 2: Storage Networking Technologies

37

The EMC Connectrix family represents the industry's most extensive selection of networked storage connectivity products. Connectrix integrates high-speed FC connectivity, highly resilient switching technology, options for intelligent IP storage networking, and I/O consolidation with products that support Fibre Channel over Ethernet (FCoE). The connectivity products offered under the Connectrix brand are: Enterprise directors, departmental switches, and multi-purpose switches .

Enterprise directors offer high port density and high component redundancy. They are deployed in high-availability or large-scale environments. Connectrix directors offer several hundred ports per domain. Departmental switches are best suited for workgroup, mid-tier environments. Multi-purpose switches support various protocols such as iSCSI, FCIP, FCoE, FICON, in addition to FC protocol. In addition to FC ports, Connectrix switches and directors have Ethernet ports and serial ports for communication and switch management functions..

EMC ControlCenter SAN Manager provides a single interface for managing a SAN. With SAN Manager, an administrator can discover, monitor, manage, and configure complex heterogeneous SAN environments. It streamlines and centralizes SAN management operations across multivendor storage networks and storage devices. It enables storage administrators to manage SAN zones and LUN masking consistently across multivendor SAN arrays and switches. EMC ControlCenter SAN Manager also supports virtual environments, including VMware, and virtual SANs.

EMC ProSphere is a newly launched tool with additional features specifically for the cloud computing environment. A future release of EMC ProSphere will include all the functionalities of EMC ControlCenter.

EMC VPLEX

- Provides solution for block-level storage virtualization and data mobility – both within and across data centers
- Enables multiple hosts located across two locations to access a single copy of data
- Provides capability to mirror a virtual volume – both within and across locations
 - ▶ Allows hosts at different data centers to simultaneously access cache-coherent copies of the same virtual volume
- VPLEX family consists of three products
 - ▶ VPLEX Local
 - ▶ VPLEX Metro
 - ▶ VPLEX Geo

EMC VPLEX is the next-generation solution for block-level virtualization and data mobility both within and across datacenters. The VPLEX appliance resides between the servers and heterogeneous storage devices. It forms a pool of distributed block storage resources and enables creating virtual storage volumes from the pool. These virtual volumes are then allocated to the servers. The virtual-to-physical-storage mapping remains hidden to the servers.

VPLEX provides nondisruptive data mobility among physical storage devices to balance the application workload and to enable both local and remote data access. The mapping of virtual volumes to physical volumes can be changed dynamically by the administrator.

VPLEX uses a unique clustering architecture and distributed cache coherency that enable multiple hosts located across two locations to access a single copy of data. VPLEX also provides the capability to mirror data of a virtual volume both within and across locations. This enables hosts at different data centers to access cache-coherent copies of the same virtual volume. To avoid application downtime due to outage at a data center, the workload can be moved quickly to another data center. Applications continue accessing the same virtual volume and remain uninterrupted by the data mobility.

The VPLEX family consists of three products: VPLEX Local, VPLEX Metro, and VPLEX Geo.

EMC VPLEX Local delivers local federation, which provides simplified management and nondisruptive data mobility across heterogeneous arrays within a data center. EMC VPLEX Metro delivers distributed federation, which provides data access and mobility between two VPLEX clusters within synchronous distances that support round-trip latency up to 5 ms. EMC VPLEX Geo delivers data access and mobility between two VPLEX clusters within asynchronous distances (that support round-trip latency up to 50 ms).

Module 5: Summary

Key points covered in this module:

- FC SAN components and connectivity options
- FC protocol stack and addressing
- Structure and organization of FC data
- Fabric services
- Fabric topologies
- Types of zoning
- Block-level storage virtualization and virtual SAN

This module covered FC SAN components – node port, cable, connector, interconnecting devices, and SAN management software; FC connectivity options – point-to-point, FC-AL, FC-SW; and fabric port types such as N_Port, E_Port, F_Port, and G_Port. It includes FC protocol stack and addressing, structure and organization of FC data, and fabric services and login types. This module also covered fabric topologies – core-edge and mesh; types of zoning – port, WWN, and mixed; block-level storage virtualization; and virtual SAN.

Check Your Knowledge – 1

- Which cable type provides minimum signal attenuation over long distance?
 - A. Twisted-pair copper
 - B. Coaxial copper
 - C. Single-mode optical
 - D. Multimode optical
- What is an F_Port in FC SAN?
 - A. Switch port that connects an E_Port
 - B. Switch port that connects an N_Port
 - C. Node port that connects an N_Port
 - D. Node port that connects an E_Port

Check Your Knowledge – 2

- Which type of fabric login enables the exchange of upper layer protocol-related parameters between N_Ports?
 - A. Fabric login
 - B. Port login
 - C. Process login
 - D. ULP login
- Which is a benefit of zoning in FC SAN?
 - A. Isolates fabric services
 - B. Restricts RSCN traffic
 - C. Enables online volume expansion
 - D. Provides non-disruptive data migration

Check Your Knowledge – 3

- Which is a benefit of VSAN?
 - A. Eliminates need for fabric login process
 - B. Provides higher network bandwidth
 - C. Improves security by isolating traffic between VSANs
 - D. Enables VSANs to share fabric zoning service

Exercise: FC SAN

- Current situation
 - ▶ An organization's IT infrastructure consists of three storage arrays direct-attached to a heterogeneous mix of 45 servers
 - ▶ Servers are dual-attached to the arrays to ensure high availability
 - ▶ Each storage array has 32 front-end ports, which could support maximum of 16 servers
 - ▶ Each storage array has the disk capacity to support a maximum of 32 servers

Current situation:

The IT infrastructure of an organization consists of three storage arrays direct-attached to a heterogeneous mix of 45 servers. All servers are dual-attached to the arrays for high availability. Because each storage array has 32 front-end ports, each could support a maximum of 16 servers. However, each existing storage array has the disk capacity to support a maximum of 32 servers.

Exercise: FC SAN (contd.)

- Organization's challenges/requirements
 - ▶ Organization needs to deploy additional 45 servers to meet growth requirements
 - ▶ Existing storage arrays are poorly utilized and adding of new servers require to purchase new arrays
 - ▶ Organization wants to implement FC SAN to overcome the scalability and utilization challenges
 - ▶ Hop count must be minimized to meet performance requirement
- Task
 - ▶ Propose a FC switched fabric topology to address organization's challenges/requirements and justify your choice
 - ▶ If 72-port FC switches are available, determine the minimum number of switches required in the fabric

Organization's challenges/requirements:

The organization plans to purchase 45 more servers to meet its growth requirements. If it continues using direct-attached storage, the organization needs to purchase additional storage arrays to connect these new servers. The organization realizes that its existing storage arrays are poorly utilized; therefore, it plans to implement FC SAN to overcome the scalability and utilization challenges. The organization uses high-performance applications; therefore, it wants to minimize the hop count for the server's access to storage.

Task:

Propose a FC switched fabric topology to address organization's challenges and requirements. Justify your choice of the fabric topology.

If 72-port switches are available for FC SAN implementation, determine the minimum number of switches required in the fabric.

Module – 6

IP SAN and FCoE



Module 6: IP SAN and FCoE

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

- Describe IP SAN protocols, components, and topology
- Describe FCoE protocol, components, and topology

This module focuses on IP SAN protocols such as Internet SCSI (iSCSI) and Fibre Channel over IP (FCIP), infrastructure components, and topology. It also focuses on Fibre Channel over Ethernet (FCoE) protocol, infrastructure components, and topology.

Module 6: IP SAN and FCoE

Lesson 1: IP SAN

During this lesson the following topics are covered:

- Drivers for IP SAN
- IP SAN Protocols: iSCSI and FCIP
- Components, topologies, and protocol stack for iSCSI and FCIP

Two primary protocols that leverage IP as the transport mechanism are Internet SCSI (iSCSI) and Fibre Channel over IP (FCIP). This lesson covers the drivers for IP SAN and iSCSI components, topologies, protocol stack, and discovery methods. It also covers FCIP protocol stack and topology.

Drivers for IP SAN

- IP SAN transports block-level data over IP network
- IP is being positioned as a storage networking option because:
 - ▶ Existing network infrastructure can be leveraged
 - ▶ Reduced cost compared to investing in new FC SAN hardware and software
 - ▶ Many long-distance disaster recovery solutions already leverage IP-based network
 - ▶ Many robust and mature security options are available for IP network

Traditional SAN enables the transfer of block I/O over Fibre Channel and provides high performance and scalability. These advantages of FC SAN come with the additional cost of buying FC components, such as FC HBA and switches. Organizations typically have an existing Internet Protocol (IP)-based infrastructure, which could be leveraged for storage networking. Advancements in technology have enabled IP to be used for transporting block I/O over the IP network. This technology of transporting block I/Os over an IP is referred to as IP SAN. IP is a mature technology, and using IP as a storage networking option provides several advantages. When block I/O is run over IP, the existing network infrastructure can be leveraged, which is more economical than investing in a new SAN infrastructure. In addition, many robust and mature security options are now available for IP networks. Many long-distance, disaster recovery (DR) solutions are already leveraging IP-based networks. With IP SAN, organizations can extend the geographical reach of their storage infrastructure.


IP SAN Protocol: iSCSI

- IP based protocol that is used to connect host and storage
- Encapsulates SCSI commands and data into an IP packet and transports them using TCP/IP

iSCSI is encapsulation of SCSI I/O over IP. iSCSI is an IP based protocol that establishes and manages connections between host and storage over IP. iSCSI encapsulates SCSI commands and data into an IP packet and transports them using TCP/IP. iSCSI is widely adopted for connecting servers to storage because it is relatively inexpensive and easy to implement, especially environments in which an FC SAN does not exist.

Components of iSCSI

- iSCSI initiator
 - ▶ Example: iSCSI HBA
- iSCSI target
 - ▶ Storage array with iSCSI port
 - ▶ iSCSI gateway – enables communication with FC storage array
- IP network



An initiator (host), target (storage or iSCSI gateway), and an IP-based network are the key iSCSI components. If an iSCSI-capable storage array is deployed, then a host with the iSCSI initiator can directly communicate with the storage array over an IP network. However, in an implementation that uses an existing FC array for iSCSI communication, an iSCSI gateway is used. These devices perform the translation of IP packets to FC frames and vice versa, thereby bridging the connectivity between the IP and FC environments.

iSCSI Host Connectivity Options

- Standard NIC with software iSCSI initiator
 - ▶ NIC provides network interface
 - ▶ Software initiator provides iSCSI functionality
 - ▶ Requires host CPU cycles for iSCSI and TCP/IP processing
- TCP Offload Engine (TOE) NIC with software iSCSI initiator
 - ▶ Moves TCP processing load off the host CPU onto the NIC card
 - ▶ Software initiator provides iSCSI functionality
 - ▶ Requires host CPU cycles for iSCSI processing
- iSCSI HBA
 - ▶ Offloads both iSCSI and TCP/IP processing from host CPU
 - ▶ Simplest option for boot from SAN


A standard NIC with software iSCSI initiator, a TCP offload engine (TOE) NIC with software iSCSI initiator, and an iSCSI HBA are the three iSCSI host connectivity options. The function of the iSCSI initiator is to route the SCSI commands over an IP network.

A standard NIC with a software iSCSI initiator is the simplest and least expensive connectivity option. It is easy to implement because most servers come with at least one, and in many cases two, embedded NICs. It requires only a software initiator for iSCSI functionality. Because NICs provide standard IP function, encapsulation of SCSI into IP packets and decapsulation are carried out by the host CPU. This places additional overhead on the host CPU. If a standard NIC is used in heavy I/O load situations, the host CPU might become a bottleneck. TOE NIC helps alleviate this burden. A TOE NIC offloads TCP management functions from the host and leaves only the iSCSI functionality to the host processor. The host passes the iSCSI information to the TOE card, and the TOE card sends the information to the destination using TCP/IP. Although this solution improves performance, the iSCSI functionality is still handled by a software initiator that requires host CPU cycles.

An iSCSI HBA is capable of providing performance benefits because it offloads the entire iSCSI and TCP/IP processing from the host processor. The use of an iSCSI HBA is also the simplest way to boot hosts from a SAN environment via iSCSI. If there is no iSCSI HBA, modifications must be made to the basic operating system to boot a host from the storage devices because the NIC needs to obtain an IP address before the operating system loads. The functionality of an iSCSI HBA is similar to the functionality of an FC HBA.

iSCSI Topologies: Native iSCSI

- iSCSI initiators are either directly attached to storage array or connected through IP network
 - ▶ No FC component
- Storage array has iSCSI port
- Each iSCSI port is configured with an IP address




Two topologies of iSCSI implementations are native and bridged. Native topology does not have FC components. The initiators may be either directly attached to targets or connected through the IP network.

FC components are not required for iSCSI connectivity if an iSCSI-enabled array is deployed. In figure in the slide, the array has one or more iSCSI ports configured with an IP address and connected to a standard Ethernet switch. After an initiator is logged on to the network, it can access the available LUNs on the storage array. A single array port can service multiple hosts or initiators as long as the array port can handle the amount of storage traffic that the hosts generate.

iSCSI Topologies: Bridged iSCSI

- iSCSI gateway is used to enable communication between iSCSI host and FC storage
- iSCSI gateway works as bridge between FC and IP network
 - ▶ Converts IP packets to FC frames and vice versa
- iSCSI initiator is configured with gateway's IP address as its target
- iSCSI gateway is configured as FC initiator to storage array



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


Section 2: Storage Networking Technologies

53

Bridged topology enables the coexistence of FC with IP by providing iSCSI-to-FC bridging functionality. Figure in the slide illustrates an iSCSI host connectivity to an FC storage array. In this case, the array does not have any iSCSI ports. Therefore, an external device, called a gateway or a multiprotocol router, must be used to facilitate the communication between the iSCSI host and FC storage. The gateway converts IP packets to FC frames and vice versa. The bridge devices contain both FC and Ethernet ports to facilitate the communication between the FC and IP environments. In bridged iSCSI implementation, the iSCSI initiator is configured with the gateway's IP address as its target destination. On the other side, the gateway is configured as an FC initiator to the storage array.

Combining FC and Native iSCSI Connectivity

- Array provides both FC and iSCSI ports
 - ▶ Enable iSCSI and FC connectivity in the same environment
 - ▶ No bridge devices needed



The most common topology is a combination of FC and native iSCSI. Typically, a storage array comes with both FC and iSCSI ports that enable iSCSI and FC connectivity in the same environment, as shown in the slide.




Figure in the slide displays a model of the iSCSI protocol layers and depicts the encapsulation order of the SCSI commands for their delivery through a physical carrier.

SCSI is the command protocol that works at the application layer of the Open System Interconnection (OSI) model. The initiators and targets use SCSI commands and responses to talk to each other. The SCSI command descriptor blocks, data, and status messages are encapsulated into TCP/IP and transmitted across the network between the initiators and targets.

iSCSI is the session-layer protocol that initiates a reliable session between devices that recognize SCSI commands and TCP/IP. The iSCSI session-layer interface is responsible for handling login, authentication, target discovery, and session management. TCP is used with iSCSI at the transport layer to provide reliable transmission.

TCP controls message flow, windowing, error recovery, and retransmission. It relies upon the network layer of the OSI model to provide global addressing and connectivity. The Layer 2 protocols at the data link layer of this model enable node-to-node communication through a physical network.

iSCSI Discovery

- For iSCSI communication, initiator must discover location and name of target on a network
- iSCSI discovery takes place in two ways:
 - ▶ SendTargets discovery
 - ▶▶ Initiator is manually configured with the target's network portal
 - ▶▶ Initiator issues SendTargets command; target responds with required parameters
 - ▶ Internet Storage Name Service (iSNS)
 - ▶▶ Initiators and targets register themselves with iSNS server
 - ▶▶ Initiator can query iSNS server for a list of available targets

An initiator must discover the location of its targets on the network and the names of the targets available to it before it can establish a session. This discovery can take place in two ways: SendTargets discovery or internet Storage Name Service (iSNS).

In SendTargets discovery, the initiator is manually configured with the target's network portal to establish a discovery session. The initiator issues the SendTargets command, and the target network portal responds with the required parameters of the targets available to the host.

iSNS enables automatic discovery of iSCSI devices on an IP network. The initiators and targets can be configured to automatically register themselves with the iSNS server. Whenever an initiator wants to know the targets that it can access, it can query the iSNS server for a list of available targets.

iSCSI Name

- iSCSI name is a unique iSCSI identifier that is used to identify initiators and targets within an iSCSI network
- Two common types of iSCSI names are:
 - ▶ iqn: iSCSI Qualified Name
 - ▶▶ iqn.2008-02.com.example:optional_string
 - ▶ eui: Extended Unique Identifier
 - ▶▶ eui.0300732A32598D26

A unique worldwide iSCSI identifier, known as an iSCSI name, is used to identify the initiators and targets within an iSCSI network to facilitate communication. The unique identifier can be a combination of the names of the department, application, or manufacturer, serial number, asset number, or any tag that can be used to recognize and manage the devices. Following are two types of iSCSI names commonly used:

- **iSCSI Qualified Name (IQN):** An organization must own a registered domain name to generate iSCSI Qualified Names. This domain name does not need to be active or resolve to an address. It just needs to be reserved to prevent other organizations from using the same domain name to generate iSCSI names. A date is included in the name to avoid potential conflicts caused by the transfer of domain names. An example of an IQN is `iqn.2008-02.com.example:optional_string`. The *optional_string* provides a serial number, an asset number, or any other device identifiers. An iSCSI Qualified Name enables storage administrators to assign meaningful names to iSCSI devices, and therefore, manage those devices more easily.
- **Extended Unique Identifier (EUI):** An EUI is a globally unique identifier based on the IEEE EUI-64 naming standard. An EUI is composed of the eui prefix followed by a 16-character hexadecimal name, such as `eui.0300732A32598D26`.

In either format, the allowed special characters are dots, dashes, and blank spaces.

IP SAN Protocol: FCIP


- IP-based protocol that is used to connect distributed FC SAN islands
- Creates virtual FC links over existing IP network that is used to transport FC data between different FC SANs
- Encapsulates FC frames onto IP packet
- Provides disaster recovery solution

FC SAN provides a high-performance infrastructure for localized data movement. Organizations are now looking for ways to transport data over a long distance between their disparate SANs at multiple geographic locations. One of the best ways to achieve this goal is to interconnect geographically dispersed SANs through reliable, high-speed links. This approach involves transporting the FC block data over the IP infrastructure. FCIP is a tunneling protocol that enables distributed FC SAN islands to be interconnected over the existing IP-based networks.

FCIP is a protocol in which FCIP entity such as FCIP gateway is used to tunnel FC fabrics through an IP network. In FCIP FC frames are encapsulated onto the IP payload. An FCIP implementation is capable to merge interconnected fabrics into a single fabric. Frequently, only a small subset of nodes at either end requires connectivity across fabrics. Thus, the majority of FCIP implementations today use switch-specific features such as IVR (Inter-VSAN Routing) or FCRS (Fibre Channel Routing Services) to create a tunnel. In this manner, traffic may be routed between specific nodes without actually merging the fabrics.

The FCIP standard has rapidly gained acceptance as a manageable, cost-effective way to blend the best of the two worlds: FC SAN and the proven, widely deployed IP infrastructure. As a result, organizations now have a better way to store, protect, and move their data by leveraging investments in their existing IP infrastructure. FCIP is extensively used in disaster recovery implementations in which data is duplicated to the storage located at a remote site.

FCIP Topology




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

Section 2: Storage Networking Technologies

59

In an FCIP environment, an FCIP gateway is connected to each fabric via a standard FC connection. The FCIP gateway at one end of the IP network encapsulates the FC frames into IP packets. The gateway at the other end removes the IP wrapper and sends the FC data to the layer 2 fabric. The fabric treats these gateways as layer 2 fabric switches. An IP address is assigned to the port on the gateway, which is connected to an IP network. After the IP connectivity is established, the nodes in the two independent fabrics can communicate with other.

FCIP Protocol Stack



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

Section 2: Storage Networking Technologies

60

The FCIP protocol stack is shown in the slide. Applications generate SCSI commands and data, which are processed by various layers of the protocol stack. The upper layer protocol SCSI includes the SCSI driver program that executes the read-and-write commands. Below the SCSI layer is the Fibre Channel Protocol (FCP) layer, which is simply a fibre channel frame whose payload is SCSI. The FCP layer rides on top of the Fibre Channel transport layer. This enables the FC frames to run natively within a SAN fabric environment. In addition, the FC frames can be encapsulated into the IP packet and sent to a remote SAN over the IP. The FCIP layer encapsulates the Fibre Channel frames onto the IP payload and passes them to the TCP layer. TCP and IP are used for transporting the encapsulated information across Ethernet, wireless, or other media that support the TCP/IP traffic.

Encapsulation of FC frame on to IP packet could cause the IP packet to be fragmented when the data link cannot support the maximum transmission unit (MTU) size of an IP packet. When an IP packet is fragmented, the required parts of the header must be copied by all fragments. When a TCP packet is segmented, normal TCP operations are responsible for receiving and re-sequencing the data prior to passing it on to the FC processing portion of the device.

Module 6: IP SAN and FCoE

Lesson 2: Fibre Channel over Ethernet (FCoE)

During this lesson the following topics are covered:

- Drivers for FCoE
- Components of FCoE network
- FCoE frame mapping
- Converged Enhanced Ethernet (CEE)

This lesson covers the drivers of FCoE, components of FCoE network, and FCoE frame mapping. It also covers converged enhanced ethernet (CEE).


Drivers for FCoE

- FCoE is a protocol that transports FC data over Ethernet network (Converged Enhanced Ethernet)
- FCoE is being positioned as a storage networking option because:
 - ▶ Enables consolidation of FC SAN traffic and Ethernet traffic onto a common Ethernet infrastructure
 - ▶ Reduces the number of adapters, switch ports, and cables
 - ▶ Reduces cost and eases data center management
 - ▶ Reduces power and cooling cost, and floor space

Data centers typically have multiple networks to handle various types of I/O traffic—for example, an Ethernet network for TCP/IP communication and an FC network for FC communication. TCP/IP is typically used for client-server communication, data backup, infrastructure management communication, and so on. FC is typically used for moving block-level data between storage and servers. To support multiple networks, servers in a data center are equipped with multiple redundant physical network interfaces—for example, multiple Ethernet and FC cards/adapters. In addition, to enable the communication, different types of networking switches and physical cabling infrastructure are implemented in data centers. The need for two different kinds of physical network infrastructure increases the overall cost and complexity of data center operation.

Fibre Channel over Ethernet (FCoE) protocol provides consolidation of LAN and SAN traffic over a single physical interface infrastructure. FCoE helps organizations address the challenges of having multiple discrete network infrastructures. FCoE uses the Converged Enhanced Ethernet (CEE) link (10 Gigabit Ethernet) to send FC frames over Ethernet.

Data Center Infrastructure – Before Using FCoE




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

Section 2: Storage Networking Technologies

63

Figure in the slide represents the infrastructure before FCoE deployment. Here, the storage resources are accessed using HBAs, and the IP network resources are accessed using NICs by the servers. Typically, in a data center, a server is configured with 2 to 4 NIC cards and redundant HBA cards. If the data center has hundreds of servers, it would require a large number of adapters, cables, and switches. This leads to a complex environment, which is difficult to manage and scale. The cost of power, cooling, and floor space further adds to the challenge.

Data Center Infrastructure – After Using FCoE



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

Section 2: Storage Networking Technologies

64

Figure in the slide shows the I/O consolidation with FCoE using FCoE switches and Converged Network Adapters (CNAs). A CNA replaces both HBAs and NICs in the server and consolidates both the IP and FC traffic. This reduces the requirement of multiple network adapters at the server to connect to different networks. Overall, this reduces the requirement of adapters, cables, and switches. This also considerably reduces the cost and management overhead.

Components of an FCoE Network


- Converged Network Adapter (CNA)
- Cable
- FCoE switch

The key FCoE components are:

- Converged Network Adapter (CNA)
- Cable
- FCoE switch

Converged Network Adapter (CNA)

- Provides functionality of both – a standard NIC and an FC HBA
 - ▶ Eliminates the need to deploy separate adapters and cables for FC and Ethernet communications
- Contains separate modules for 10 Gigabit Ethernet, FC, and FCoE ASICs
 - ▶ FCoE ASIC encapsulates FC frames into Ethernet frames



A CNA provides the functionality of both a standard NIC and an FC HBA in a single adapter and consolidates both types of traffic. CNA eliminates the need to deploy separate adapters and cables for FC and Ethernet communications, thereby reducing the required number of server slots and switch ports. CNA offloads the FCoE protocol processing task from the server, thereby freeing the server CPU resources for application processing. A CNA contains separate modules for 10 Gigabit Ethernet, Fibre Channel, and FCoE Application Specific Integrated Circuits (ASICs). The FCoE ASIC encapsulate FC frames into Ethernet frames. One end of this ASIC is connected to 10GbE and FC ASICs for server connectivity, while the other end provides a 10GbE interface to connect to an FCoE switch.

Cable

- Two options are available for FCoE cabling
 - ▶ Copper based Twinax cable
 - ▶ Standard fiber optical cable


Twinax Cable	Fiber Optical Cable
Suitable for shorter distances (up to 10 meters)	Can run over longer distances
Requires less power and are less expensive than fiber optical cable	Relatively more expensive than Twinax cables
Uses Small Form Factor Pluggable Plus (SFP+) connector	Uses Small Form Factor Pluggable Plus (SFP+) connector

Currently two options are available for FCoE cabling: Copper based Twinax and standard fiber optical cables. A Twinax cable is composed of two pairs of copper cables covered with a shielded casing. The Twinax cable can transmit data at the speed of 10 Gbps over shorter distances up to 10 meters. Twinax cables require less power and are less expensive than fiber optic cables.

The Small Form Factor Pluggable Plus (SFP+) connector is the primary connector used for FCoE links and can be used with both optical and copper cables.

FCoE Switch


- Provides both Ethernet and FC switch functionalities
- Consists of FCF, Ethernet bridge, and set of CEE ports and FC ports (optional)
 - ▶ FCF encapsulates and de-encapsulates FC frames
- Forwards frames based on Ethertype



An FCoE switch has both Ethernet switch and Fibre Channel switch functionalities. The FCoE switch has a Fibre Channel Forwarder (FCF), Ethernet Bridge, and set of Ethernet ports and optional FC ports. The function of the FCF is to encapsulate the FC frames, received from the FC port, into the FCoE frames and also to de-encapsulate the FCoE frames, received from the Ethernet Bridge, to the FC frames.

Upon receiving the incoming traffic, the FCoE switch inspects the Ethertype (used to indicate which protocol is encapsulated in the payload of an Ethernet frame) of the incoming frames and uses that to determine the destination. If the Ethertype of the frame is FCoE, the switch recognizes that the frame contains an FC payload and forwards it to the FCF. From there, the FC is extracted from the FCoE frame and transmitted to FC SAN over the FC ports. If the Ethertype is not FCoE, the switch handles the traffic as usual Ethernet traffic and forwards it over the Ethernet ports.

FCoE Frame Mapping



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

Section 2: Storage Networking Technologies

69

The encapsulation of the Fibre Channel frame occurs through the mapping of the FC frames onto Ethernet, as shown in the slide. Fibre Channel and traditional networks have stacks of layers where each layer in the stack represents a set of functionalities. The FC stack consists of five layers: FC-0 through FC-4. Ethernet is typically considered as a set of protocols that operates at the physical and data link layers in the seven-layer OSI stack. The FCoE protocol specification replaces the FC-0 and FC-1 layers of the FC stack with Ethernet. This provides the capability to carry the FC-2 to the FC-4 layer over the Ethernet layer.

A typical Fibre Channel data frame has a 2,112-byte payload, a 24-byte header, and an FCS. A standard Ethernet frame has a default payload capacity of 1,500 bytes. To maintain good performance, FCoE must use jumbo frames to prevent a Fibre Channel frame from being split into two Ethernet frames.

Converged Enhanced Ethernet


- Provides lossless Ethernet
- Lossless Ethernet requires following functionalities:
 - ▶ Priority-based flow control (PFC)
 - ▶ Enhanced transmission selection (ETS)
 - ▶ Congestion notification (CN)
 - ▶ Data center bridging exchange protocol(DCBX)

Conventional Ethernet is lossy in nature, which means that frames might be dropped or lost during transmission. Converged Enhanced Ethernet (CEE) or lossless Ethernet provides a new specification to the existing Ethernet standard that eliminates the lossy nature of Ethernet. This makes 10 Gb Ethernet a viable storage networking option, similar to FC. Lossless Ethernet requires certain functionalities. These functionalities are defined and maintained by the data center bridging (DCB) task group, which is a part of the IEEE 802.1 working group and they are:

- Priority-based flow control
- Enhanced transmission selection
- Congestion notification
- Data center bridging exchange protocol

Priority-based Flow Control (PFC)

- Creates eight virtual links on a single physical link
- Uses PAUSE capability of Ethernet for each virtual link
 - ▶ A virtual link can be paused and restarted independently
 - ▶ PAUSE mechanism is based on user priorities or classes of service



Traditional FC manages congestion through the use of a link-level, credit-based flow control that guarantees no loss of frames. Typical Ethernet, coupled with TCP/IP, uses a packet drop flow control mechanism. The packet drop flow control is not lossless. This challenge is eliminated by using an IEEE 802.3x Ethernet PAUSE control frame to create a lossless Ethernet. A receiver can send a PAUSE request to a sender when the receiver's buffer is filling up. Upon receiving a PAUSE frame, the sender stops transmitting frames, which guarantees no loss of frames. The downside of using the Ethernet PAUSE frame is that it operates on the entire link, which might be carrying multiple traffic flows.

PFC provides a link level flow control mechanism. PFC creates eight separate virtual links on a single physical link and allows any of these links to be paused and restarted independently. PFC enables the pause mechanism based on user priorities or classes of service. Enabling the pause based on priority allows creating lossless links for traffic, such as FCoE traffic. This PAUSE mechanism is typically implemented for FCoE while regular TCP/IP traffic continues to drop frames. Figure in the slide illustrates how a physical Ethernet link is divided into eight virtual links and allows a PAUSE for a single virtual link without affecting the traffic for the others.


Enhanced Transmission Selection (ETS)

- Allocates bandwidth to different traffic classes such as LAN, SAN, and Inter Process Communication (IPC)
- Provides available bandwidth to other classes of traffic when a particular class of traffic does not use its allocated bandwidth

Enhanced transmission selection provides a common management framework for the assignment of bandwidth to different traffic classes, such as LAN, SAN, and Inter Process Communication (IPC). When a particular class of traffic does not use its allocated bandwidth, ETS enables other traffic classes to use the available bandwidth.

Congestion Notification (CN)

- Provides a mechanism for detecting congestion and notifying the source
 - ▶ Enables a switch to send a signal to other ports that need to stop or slow down their transmissions



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

Section 2: Storage Networking Technologies

73

Congestion notification provides end-to-end congestion management for protocols, such as FCoE, that do not have built-in congestion control mechanisms. Link level congestion notification provides a mechanism for detecting congestion and notifying the source to move the traffic flow away from the congested links. Link level congestion notification enables a switch to send a signal to other ports that need to stop or slow down their transmissions. The process of congestion notification and its management is shown in the slide, which represents the communication between the nodes A (sender) and B (receiver). If congestion at the receiving end occurs, the algorithm running on the switch, generates a congestion notification (CN) message to the sending node (Node A). In response to the CN message, the sending end limits the rate of data transfer.

Data Center Bridging Exchange Protocol (DCBX)

- Enables CEE devices to convey and configure their features with other CEE devices in the network
 - ▶ Allows a switch to distribute configuration values to attached adapters
- Ensures consistent configuration across network

DCBX protocol is a discovery and capability exchange protocol, which helps Converged Enhanced Ethernet devices to convey and configure their features with the other CEE devices in the network. DCBX is used to negotiate capabilities between the switches and the adapters, and it allows the switch to distribute the configuration values to all the attached adapters. This helps to ensure consistent configuration across the entire network.

Module 6: Summary

Key points covered in this module:

- IP SAN protocols, their components, and topologies
- FCoE protocol, its components, and topology

This module covered IP SAN protocols such as iSCSI and FCIP, their components, and topologies. It also covered FCoE protocol, its components, and topology.

Check Your Knowledge – 1

- Which iSCSI host connectivity option offloads both iSCSI and TCP/IP processing from the host CPU?
 - A. Standard NIC with iSCSI initiator software
 - B. TOE NIC
 - C. iSCSI HBA
 - D. CNA
- Which type of iSCSI name requires a registered domain name to generate unique iSCSI identifier?
 - A. eui
 - B. iqn
 - C. WWN
 - D. MAC

Check Your Knowledge – 2

- Which protocol encapsulates FC frames onto IP packet?
 - A. FCoE
 - B. iSCSI
 - C. FCIP
 - D. CIFS
- Which is a feature of priority-based flow control?
 - A. A virtual link can be paused independently
 - B. All virtual links are paused together
 - C. Enables pausing virtual links based on their bandwidth
 - D. Enables pausing individual physical links based on their priority

Check Your Knowledge – 3

- Which functionality enables allocation of bandwidth to different traffic classes in an FCoE environment?
 - A. Priority-based flow control
 - B. Enhanced transmission selection
 - C. Congestion notification
 - D. Data center bridging exchange

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.

File Sharing Technology Evolution



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

Section 2: Storage Networking Technologies

83

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.


Section 2: Storage Networking Technologies

84

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.

Section 2: Storage Networking Technologies

87

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.

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 I/O Operation



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

Section 2: Storage Networking Technologies

92

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.

Section 2: Storage Networking Technologies

95

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.

Section 2: Storage Networking Technologies

97

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.

Section 2: Storage Networking Technologies


99

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.


Section 2: Storage Networking Technologies

100

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.

Section 2: Storage Networking Technologies

101

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)

Concept in Practice:

- EMC Isilon
- EMC VNX Gateway

The Concept 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

Module – 8

Object-Based and Unified Storage



Module 8: Object-based and Unified Storage

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

- Describe the object-based storage model
- List the key components of object-based storage
- Describe the storage and retrieval process in object-based storage
- Describe content-addressed storage
- List the key components of unified storage
- Describe the process of data access from unified storage

This module focuses on object-based storage device and unified storage device. It covers the object-based storage model. The module details on the function of the key components and process of storage and retrieval of object-based storage. This module also covers content-addressed storage, which is a special type of object-based storage and key features it supports.

This module also focus on the function of the key components of unified storage. Finally, this module covers the process of data access from unified storage.

Module 8: Object-based and Unified Storage

Lesson 1: Object-based Storage

During this lesson the following topics are covered:

- Comparison of hierarchical file system and flat address space
- Object-based storage model
- Key components of object-based storage
- Storage and retrieval process in object-based storage devices
- Content-addressed storage

This lesson covers a comparison of hierarchical file system with flat address space. The lesson describes the model and key components of object-based storage. This lesson also covers storage and retrieval process in the object-based storage. Finally, the lesson covers content-addressed storage.


Drivers for Object-based Storage

- More than 90% of the data being generated is unstructured
- Traditional solutions are inefficient to handle the growth
 - ▶ High overhead on NAS due to managing large number of permissions and nested directories
- These challenges demanded a smarter approach to manage unstructured data based on its content

Object-based storage is a way to store file data in the form of objects on flat address space based on its content and attributes rather than the name and location.

Recent studies have shown that more than 90 percent of data generated is unstructured. This growth of unstructured data has posed new challenges to IT administrators and storage managers. With this growth, traditional NAS, which is a dominant solution for storing unstructured data, has become inefficient. Data growth adds high overhead to the network-attached storage (NAS) in terms of managing large number of permission and nested directories. In an enterprise environment, NAS also manages large amounts of metadata generated by hosts, storage systems, and individual applications. Typically this metadata is stored as part of the file and distributed throughout the environment. This adds to the complexity and latency in searching and retrieving files. These challenges demand a smarter approach to manage unstructured data based on its content rather than metadata about its name, location, and so on. *Object-based storage* is a way to store file data in the form of objects based on its content and other attributes rather than the name and location.

Hierarchical File System vs. Flat Address Space




- Hierarchical file system organizes data in the form of files and directories
- Object-based storage devices store the data in the form of objects
 - ▶ It uses flat address space that enables storage of large number of objects
 - ▶ An object contains user data, related metadata, and other attributes
 - ▶ Each object has a unique object ID, generated using specialized algorithm

An OSD is a device that organizes and stores unstructured data, such as movies, office documents, and graphics, as objects. Object-based storage provides a scalable, self-managed, protected, and shared storage option. OSD stores data in the form of *objects*. OSD uses flat address space to store data. Therefore, there is no hierarchy of directories and files; as a result, a large number of objects can be stored in an OSD system .

An object might contain user data, related metadata (size, date, ownership, and so on), and other attributes of data (retention, access pattern, and so on). Each object stored in the system is identified by a unique ID called the *object ID*. The object ID is generated using specialized algorithms such as hash function on the data and guarantees that every object is uniquely identified.

Traditional vs. Object-based Storage Model



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

Section 2: Storage Networking Technologies


116

An I/O in the traditional block access method passes through various layers in the I/O path. The I/O generated by an application passes through the file system, the channel, or network and reaches the disk drive. When the file system receives the I/O from an application, the file system maps the incoming I/O to the disk blocks. The block interface is used for sending the I/O over the channel or network to the storage device. The I/O is then written to the block allocated on the disk drive.

The file system has two components: user component and storage component. The user component of the file system performs functions such as hierarchy management, naming, and user access control. The storage component maps the files to the physical location on the disk drive.

When an application accesses data stored in OSD, the request is sent to the file system user component. The file system user component communicates to the OSD interface, which in turn sends the request to the storage device. The storage device has the OSD storage component responsible for managing the access to the object on a storage device. After the object is stored, the OSD sends an acknowledgment to the application server. The OSD storage component manages all the required low-level storage and space management functions. It also manages security and access control functions for the objects.

Key Components of Object-based Storage Device




- OSD system typically comprises three key components:
 - ▶ OSD nodes
 - ▶ Internal network
 - ▶ Storage

The OSD system is typically composed of three key components: nodes, private network, and storage.

The OSD system is composed of one or more *nodes*. A node is a server that runs the OSD operating environment and provides services to store, retrieve, and manage data in the system. The OSD node has two key services: metadata service and storage service. The metadata service is responsible for generating the object ID from the contents (may also include other attributes of data) of a file. It also maintains the mapping of the object IDs and the file system namespace. The storage service manages a set of disks on which the user data is stored. The OSD nodes connect to the storage via an internal network. The internal network provides node-to-node connectivity and node-to-storage connectivity. The application server accesses the node to store and retrieve data over an external network. In some implementations, such as CAS, the metadata service might reside on the application server or on a separate server.


OSD typically uses low-cost and high-density disk drives to store the objects. As more capacity is required, more disk drives can be added to the system.



The process of storing objects in OSD is illustrated on the slide. The data storage process in an OSD system is as follows:

1. The application server presents the file to be stored to the OSD node.
2. The OSD node divides the file into two parts: user data and metadata.
3. The OSD node generates the object ID using a specialized algorithm. The algorithm is executed against the contents of the user data to derive an ID unique to this data.
4. For future access, the OSD node stores the metadata and object ID using the metadata service.
5. The OSD node stores the user data (objects) in the storage device using the storage service.
6. An acknowledgment is sent to the application server stating that the object is stored.

Process of Retrieving Object from OSD



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

Section 2: Storage Networking Technologies

119

After an object is stored successfully, it is available for retrieval. A user accesses the data stored on OSD by the same filename. The application server retrieves the stored content using the object ID. This process is transparent to the user.

The process of retrieving objects in OSD is illustrated on the slide. The process of data retrieval from OSD is as follows:


1. The application server sends a read request to the OSD system.
2. The metadata service retrieves the object ID for the requested file.
3. The metadata service sends the object ID to the application server.
4. The application server sends the object ID to the OSD storage service for object retrieval.
5. The OSD storage service retrieves the object from the storage device.
6. The OSD storage service sends the file to the application server.

Key Benefits of Object-based Storage

Benefits	Description
Security and reliability	• Unique object ID generated by specialized algorithms ensures data integrity and content authenticity • Request authentication is performed at storage device
Platform independence	• Because objects are abstract containers of data, it enables sharing of objects across heterogeneous platforms • This capability makes object-based storage suitable for cloud computing environment
Scalability	• Both OSD nodes and storage can be independently scaled
Manageability	• Have inherent intelligence to manage objects • Have self-healing capability • Policy based management capability enables OSD to handle routine jobs automatically

For unstructured data, object-based storage devices provide numerous benefits over traditional storage solutions. An ideal storage architecture should provide performance, scalability, security, and data sharing across multiple platforms. Traditional storage solutions, such as SAN, and NAS, do not offer all these benefits as a single solution. Object-based storage combines benefits of both the worlds. It provides platform and location independence, and at the same time, provides scalability, security and data-sharing capabilities. The key benefits of object-based storage are as follows:

- **Security and reliability:** Data integrity and content authenticity are the key features of object-based storage devices. OSD uses specialized algorithms to create objects that provide strong data encryption capability. In OSD, request authentication is performed at the storage device rather than with an external authentication mechanism.
- **Platform independence:** Objects are abstract containers of data, including metadata and attributes. This feature allows objects to be shared across heterogeneous platforms locally or remotely. This platform-independence capability makes object-based storage the best candidate for cloud computing environments.
- **Scalability:** Due to the use of flat address space, object-based storage can handle large amounts of data without impacting performance. Both storage and OSD nodes can be scaled independently in terms of performance and capacity.
- **Manageability:** Object-based storage has an inherent intelligence to manage and protect objects. It uses self-healing capability to protect and replicate objects. Policy-based management capability helps OSD to handle routine jobs automatically.



Cloud-based storage is a promising use case of OSD. OSD uses a web interface to access storage resources. OSD provides inherent security, scalability, and automated data management. It also enables data sharing across heterogeneous platforms or tenants while ensuring integrity of data. These capabilities make OSD a strong option for cloud-based storage. Cloud service providers can leverage OSD to offer storage-as-a-service. OSD supports web service access via *representational state transfer* (REST) and *simple object access protocol* (SOAP). REST and SOAP APIs can be easily integrated with business applications that access OSD over the web.

Representational State Transfer or REST is an architectural style developed for modern web applications. REST provides lightweight web services to access resources (for example, documents, blogs, and so on) on which a few basic operations can be performed, such as retrieving, modifying, creating, and deleting resources. REST-style web services are resource-oriented services. Resources can be uniquely located and identified by a Universal Resource Identifier (URI), and operations can be performed on those resources using an HTTP specification. For example, if a user accesses a blog using REST via a unique identifier, the request returns the representation of the blog in a particular format (XML or HTML). *Simple Object Access Protocol* or SOAP is a XML based protocol that enables communication between the web applications running on different OS and based on different programming languages. SOAP provides process to encode HTTP header and XML file to enable and pass information between different computers.

Cloud based storage is further discussed in module 13 Cloud Computing.

Use Case 2: Content Address Storage (CAS)

- Storage designed to store fixed content
- Stores data as objects
- Each object is assigned a globally unique identifier, known as content address (CA)
 - ▶ CA is derived from the binary representation of the data
- CAS device can be accessed via the CAS API running on the application server

A data archival solution is a promising use case for OSD. Data integrity and protection is the primary requirement for any data archiving solution. Traditional archival solutions—CD and DVD-ROM—do not provide scalability and performance. OSD stores data in the form of objects, associates them with a unique object ID, and ensures high data integrity. Along with integrity, it provides scalability and data protection. These capabilities make OSD a viable option for long term data archiving for fixed content. Content addressed storage (CAS) is a special type of object-based storage device purposely built for storing fixed content.

CAS is an object-based storage device designed for secure online storage and retrieval of fixed content. CAS stores user data and its attributes as an object. The stored object is assigned a globally unique address, known as a *content address* (CA). This address is derived from the object's binary representation. CAS provides an optimized and centrally managed storage solution. Data access in CAS differs from other OSD devices. In CAS, the application server can access the CAS device only via the CAS API running on the application server. However, the way CAS stores data is similar to the other OSD systems.

Key Features of CAS

- Content authenticity and integrity
- Location independence
- Single instance storage
- Retention enforcement
- Data protection
- Fast record retrieval
- Load balancing
- Scalability
- Self diagnosis and repair
- Audit trail and event notification


CAS provides all the features required for storing fixed content. The key features of CAS are as follows:

- *Content authenticity*: It assures the genuineness of stored content. This is achieved by generating a unique content address for each object and validating the content address for stored objects at regular intervals. Content authenticity is assured because the address assigned to each object is as unique as a fingerprint. Every time an object is read, CAS uses a hashing algorithm to recalculate the object's content address as a validation step and compares the result to its original content address. If the object fails validation, CAS rebuilds the object using a mirror or parity protection scheme.
- *Content integrity*: It provides assurance that the stored content has not been altered. CAS uses a hashing algorithm for content authenticity and integrity. If the fixed content is altered, CAS generates a new address for the altered content, rather than overwrite the original fixed content.
- *Location independence*: CAS uses a unique content address, rather than directory path names or URLs, to retrieve data. This makes the physical location of the stored data irrelevant to the application that requests the data.

Cont..

- *Single-instance storage (SIS)*: CAS uses a unique content address to guarantee the storage of only a single instance of an object. When a new object is written, the CAS system is polled to see whether an object is already available with the same content address. If the object is available in the system, it is not stored; instead, only a pointer to that object is created.
- *Retention enforcement*: Protecting and retaining objects is a core requirement of an archive storage system. After an object is stored in the CAS system and the retention policy is defined, CAS does not make the object available for deletion until the policy expires.
- *Data protection*: CAS ensures that the content stored on the CAS system is available even if a disk or a node fails. CAS provides both local and remote protection to the data objects stored on it. In the local protection option, data objects are either mirrored or parity protected. In mirror protection, two copies of the data object are stored on two different nodes in the same cluster. This decreases the total available capacity by 50 percent. In parity protection, the data object is split in multiple parts and parity is generated from them. Each part of the data and its parity are stored on a different node. This method consumes less capacity to protect the stored data, but takes slightly longer to regenerate the data if corruption of data occurs. In the remote replication option, data objects are copied to a secondary CAS at the remote location. In this case, the objects remain accessible from the secondary CAS if the primary CAS system fails.
- *Fast record retrieval*: CAS stores all objects on disks, which provides faster access to the objects compared to tapes and optical discs.
- *Load balancing*: CAS distributes objects across multiple nodes to provide maximum throughput and availability.
- *Scalability*: CAS allows the addition of more nodes to the cluster without any interruption to data access and with minimum administrative overhead.
- *Event notification*: CAS continuously monitors the state of the system and raises an alert for any event that requires the administrator's attention. The event notification is communicated to the administrator through SNMP, SMTP, or e-mail.
- *Self diagnosis and repair*: CAS automatically detects and repairs corrupted objects and alerts the administrator about the potential problem. CAS systems can be configured to alert remote support teams who can diagnose and repair the system remotely.
- *Audit trails*: CAS keeps track of management activities and any access or disposition of data. Audit trails are mandated by compliance requirements.

CAS Implementation Use Case 1: Healthcare Solution




- Each X-ray image size range from about 15MB to over 1GB
- Patient records are stored online for a period of 60 days
- Beyond 60 days patient records are archived to CAS

Large healthcare centers examine hundreds of patients every day and generate large volumes of medical records. Each record might be composed of one or more images that range in size from approximately 15 MB for a standard digital X-ray to more than 1 GB for oncology studies. The patient records are stored online for a specific period of time for immediate use by the attending physicians. Even if a patient's record is no longer needed, compliance requirements might stipulate that the records be kept in the original format for several years.

Medical image solution providers offer hospitals the capability to view medical records, such as X-ray images, with acceptable response times and resolution to enable rapid assessments of patients. Patients' records are retained on the primary storage for 60 days after which they are moved to the CAS system. CAS facilitates long-term storage and at the same time, provides immediate access to data, when needed.

CAS Implementation Use Case 2: Financial Solution



- Each check image size is about 25KB
- Check imaging service provider might process around 90 million check images per month
- Checks are stored online for a period of 60 days
- Beyond 60 days data is archived to CAS

In a typical banking scenario, images of checks, each approximately 25 KB in size, are created and sent to archive services over an IP network. A check imaging service provider might process approximately 90 million check images per month. Typically, check images are actively processed in transactional systems for about 5 days.

For the next 60 days, check images may be requested by banks or individual consumers for verification purposes and beyond 60 days, access requirements drop drastically. The check images are moved from the primary storage to the CAS system after 60 days, and can be held there for long term based on retention policy. Check imaging is one example of a financial service application that is best serviced with CAS. Customer transactions initiated by e-mail, contracts, and security transaction records might need to be kept online for 30 years; CAS is the preferred storage solution in such cases.

Module 8: Object-based and Unified Storage

Lesson 2: Unified Storage

During this lesson the following topics are covered:


- Key components of unified storage
- Data access from unified storage

This lesson describes the function of key components of unified storage. This lesson also describes the process of data access from unified storage.

Drivers for Unified Storage


- Deploying disparate storage solutions (SAN, NAS, and OSD) adds management cost, complexity, and environmental overhead
- Unified storage consolidates block, file, and object-based access within one unified platform
 - ▶ Supports multiple protocols for data access
 - ▶ Can be managed through single management interface

Due to varied application requirements, organizations have been deploying storage area networks (SANs), NAS, and object-based storage devices (OSDs) in their data centers. Deploying these disparate storage solutions adds management complexity, cost and environmental overheads. An ideal solution would be to have an integrated storage solution that supports block, file, and object access. Unified storage has emerged as a solution that consolidates block, file, and object-based access within one unified platform. It supports multiple protocols for data access and can be managed using a single management interface.



A unified storage system consists of following key components: storage controller, NAS head, OSD node, and storage.

- The *storage controller* provides block-level access to application servers through iSCSI, FC, or FCoE protocols. It contains iSCSI, FC, and FCoE front-end ports for direct block access. The storage controller is also responsible for managing the back-end storage pool in the storage system. The controller configures LUNs and presents them to application servers, NAS heads, and OSD nodes. The LUNs presented to the application server appear as local physical disks. A file system is configured on these LUNs and is made available to applications for storing data.
- A *NAS head* is a dedicated file server that provides file access to NAS clients. The NAS head is connected to the storage via the storage controller typically using a FC or FCoE connection. The system typically has two or more NAS heads for redundancy. The LUNs presented to the NAS head appear as physical disks. The NAS head configures the file systems on these disks, creates a NFS, CIFS, or mixed share, and exports the share to the NAS clients.
- The *OSD node* accesses the storage through the storage controller using a FC or FCoE connection. The LUNs assigned to the OSD node appear as physical disks. These disks are configured by the OSD nodes, enabling them to store the data from the web application servers.



In a unified storage system, block, file, and object requests to the storage travel through different I/O paths.

- **Block I/O request:** The application servers are connected to an FC, iSCSI, or FCoE port on the storage controller. The server sends a block request over an FC, iSCSI, or FCoE connection. The storage controller processes the I/O and responds to the application server.
- **File I/O request:** The NAS clients (where the NAS share is mounted or mapped) sends a file request to the NAS head using the NFS or CIFS protocol. The NAS head receives the request, converts it into a block request, and forwards it to the storage controller. Upon receiving the block data from the storage controller, the NAS head again converts the block request back to the file request and sends it to the clients.
- **Object I/O request:** The web application servers send an object request, typically using REST or SOAP protocols, to the OSD node. The OSD node receives the request, converts it into a block request, and sends it to the disk through the storage controller. The controller in turn processes the block request and responds back to the OSD node, which in turn provides the requested object to the web application server.

Module 8: Object-based and Unified Storage

Concept in Practice

- EMC Atmos
- EMC VNX
- EMC Centera

The Concept in Practice covers the product example of object-based and unified storage. It covers three products: EMC Atmos, EMC VNX, and EMC Centera.

EMC Atmos

- Massively scalable objects-based storage
- Can be deployed in two ways: purpose-built hardware appliance or virtual machine (VM)
- Key features
 - ▶ Enable policy-based management
 - ▶ Provide protection with replication and parity
 - ▶ Provide services such as compression and deduplication
 - ▶ Support web service and legacy protocols
 - ▶▶ REST, SOAP, CIFS, NFS, and Installable File System
 - ▶ Enable automated system management
 - ▶ Supports multitenancy
 - ▶ Flexible administration

EMC Atmos supports object-based storage for unstructured data, such as pictures and videos. Atmos combines massive scalability with specialized intelligence to address the cost, distribution, and management challenges associated with vast amounts of unstructured data.


Atmos can be deployed in two ways: as a purpose-built hardware appliance or as software in VMware environments, where Atmos VE can leverage the existing servers and storage. The hardware appliance is comprised of servers (nodes) connected to standard disk enclosures. The rack includes a 24-port Gigabit Ethernet switch to provide internode communication. The Atmos software is installed on each node.

Atmos VE enables users to exploit the power of Atmos in a virtualized environment. It can be deployed on a virtual machine in VMware ESXi hosts and configured with the VMware certified back-end storage.

Cont..

Following are the key features offered by EMC Atmos:

- *Policy-based management*: EMC Atmos improves operational efficiency by automatically distributing content based on business policy. The administrator-defined policies dictate how, when, and where the information resides.
- *Protection*: Atmos offers two options to protect the objects, replication and Geo Parity:
 - *Replication* ensures that the content is available and accessible by creating redundant copies of an object at multiple designated locations.
 - *Geo Parity* ensures that the content is available and accessible by dividing objects into multiple segments plus parity segments and distributing them to one or more designated locations.
- *Data services*: EMC Atmos includes the data services, such as, compression, and deduplication. These features are native to Atmos and can be managed and accessed via a policy.
- *Web services and legacy protocols*: EMC Atmos provides flexible web services access (REST/SOAP) for web-scale applications and file access (CIFS/NFS/Installable File System/Centera API) for traditional applications.
- *Automated system management*: EMC Atmos provides auto-configuring, auto-managing, and auto-healing capabilities to reduce administration and downtime.
- *Multitenancy*: EMC Atmos enables multiple applications to be served from the same infrastructure. Each application is securely partitioned and cannot access the other application's data. Multitenancy is ideal for service providers or large enterprises that want to provide cloud computing services to multiple customers or departments allowing logical and secure separation within a single infrastructure.
- *Flexible administration*: EMC Atmos can be managed via a graphical user interface (GUI) or command line interface (CLI).



EMC Centera is a simple, affordable, and secure repository for information archiving. EMC Centera is designed and optimized specifically to deal with the storage and retrieval of fixed content by meeting performance, compliance, and regulatory requirements. Compared to traditional archive storage, EMC Centera provides faster record retrieval, Single instance storage (SIS), guaranteed content authenticity, self-healing, and support for numerous industry and regulatory standards.

EMC Centera is offered in three different models to meet different types of user requirements—EMC Centera Basic, EMC Centera Governance Edition, and EMC Centera Compliance Edition Plus (CE+):

- *EMC Centera Basic*: Provides all functionalities without the enforcement of retention periods.
- *EMC Centera Governance Edition*: Provides the retention capabilities required by organizations to manage digital records in addition to the features provided by EMC Centera Basic.
- *EMC Centera Compliance Edition Plus*: Provides extensive compliance capabilities. CE+ is designed to meet the requirements of the most stringent regulated business environments for electronic storage media, as established by regulations from the Securities and Exchange Commission (SEC), or other national and international regulatory groups.

Cont..

A client accesses the Centera over a LAN. The client can access Centera only through the server that runs the Centera API (application programming interface). The Centera API is responsible for performing functions that enable an application to store and retrieve the data.

Centera architecture is a *Redundant Array of Independent Nodes* (RAIN). It contains storage nodes and access nodes that are networked as a cluster by using a private LAN. The internal LAN reconfigures automatically when it detects configuration changes, such as the addition of storage or access nodes. The application server accesses the Centera via an external LAN.

The nodes are configured with low-cost, high-capacity SATA disk drives. These nodes run CentraStar, the operating environment for Centera, which provides the features and functionalities required in a Centera system.

When the nodes are installed, they are configured with a “role” that defines the functionality provided by the node. A node can be configured as a storage node, an access node, or a dual-role node.

Storage nodes store and protect data objects. They are sometimes referred to as *back-end nodes*.

Access nodes provide connectivity to application servers through an external LAN. They establish connectivity with the storage nodes in the cluster through a private LAN. The number of access nodes is determined by the amount of throughput required from the cluster. If a node is configured solely as an “access node,” its disk space cannot be used to store data objects. Storage and retrieval requests are sent to the access node via the external LAN.

Dual-role nodes provide both storage and access-node capabilities. This configuration is more common than a pure access-node configuration.

EMC VNX

- Unified storage platform that consolidates block, file, and object accesses in one solution
 - ▶ Supports block access via storage processors
 - ▶ File access via X-Blade
 - ▶ Object access via EMC Atmos VE
- Components of VNX are:
 - ▶ Storage processor
 - ▶ X-Blade
 - ▶ Control station
 - ▶ Disk-array enclosures
 - ▶ Standby power supply



EMC VNX is a unified storage platform that consolidates block, file, and object access in one solution. It implements a modular architecture that integrates hardware components for block, file, and object access. EMC VNX delivers file access (NAS) functionality via X-Blades (Data Movers) and block access functionality via storage processors. Optionally it offers object access to the storage using EMC Atmos Virtual Edition (Atmos VE). VNX storage systems include the following components:

- *Storage processors* (SPs) support block I/O access to storage with FC, iSCSI, and FCoE protocols.
- *X-Blades* access data from the back end and provide host access with NFS, CIFS, MPFS, pNFS, and FTP protocols. The X-Blades in each array are scalable and provide redundancy to ensure no single point of failure.
- *Control Stations* provide management functions to the X-Blades. The Control Station is also responsible for X-Blade failover. The Control Station may optionally be configured with a matching secondary Control Station to ensure management redundancy on the VNX array.
- *Disk-array enclosures* (DAEs) house the drives used in the array. Different sized DAEs are available that can each hold a maximum of 15, 25, or 60 drives. More DAEs can be added to meet growing storage demands.
- *Standby power supplies* provide enough power to each storage processor and first DAE to ensure that any data in flight is stored in the vault area if a power failure occurs. This ensures that no writes are lost.

Module 8: Summary

Key points covered in this module:

- Object-based storage model
- Key components of object-based storage
- Process of storage and retrieval in object-based storage
- Content-addressed storage
- Key components of unified storage
- Process of data access from unified storage

This module covered the object-based storage model. Object-based storage has three key components such as OSD nodes, internal network, and storage. This module also covered the process of storage and retrieval in object-based storage and content-addressed storage.

Further, this module covered the key components of unified storage; storage processor, NAS head, OSD node, and storage. Finally, this module covered the process of data access from unified storage.

Check Your Knowledge – 1

- What is an advantage of a flat address space over a hierarchical address space?
 - A. Highly scalable with minimal impact on performance
 - B. Provides access to data, based on retention policies
 - C. Provides access to block, file, and object with same interface
 - D. Consumes less bandwidth on network while accessing data
- What is a role of metadata service in an OSD node?
 - A. Responsible for storing data in the form of objects
 - B. Stores unique IDs generated for objects
 - C. Stores both objects and object IDs
 - D. Controls functioning of storage devices

Check Your Knowledge – 2

- What is used to generate an object ID in a CAS system?
 - A. File metadata
 - B. Source and destination address
 - C. Binary representation of data
 - D. File system type and ownership
- What accurately describes block I/O access in a unified storage?
 - A. I/O traverse NAS head and storage controller to disk
 - B. I/O traverse OSD node and storage controller to disk
 - C. I/O traverse storage controller to disk
 - D. I/O is directly sent to the disk

Check Your Knowledge – 3

- What accurately describes unified storage?
 - A. Provides block, file, and object-based access within one platform
 - B. Provides block and file storage access using objects
 - C. Supports block and file access using flat address space
 - D. Specialized storage device purposely built for archiving

Section 2: Summary

Key points covered in this section:

- FC SAN components, connectivity options, fabric topologies, and zoning
- Block-level storage virtualization and virtual SAN
- IP SAN and FCoE protocols
- NAS components and implementations
- File-level virtualization
- Key components of object-based and unified storage
- Process of data access in object-based and unified storage

This section covered FC SAN components, connectivity options, fabric topologies, and zoning. It also covered block-level storage virtualization and virtual SAN. It covered IP SAN protocols such as iSCSI and FCIP, also FCoE protocol. Further, this section covered NAS components and three common NAS implementations along with file-level virtualization. Finally, this section covered the key components of object-based and unified storage along with the process of data access.