Fibre Channel Storage area Network

An effective information management solution must provide

→Just-in-time information to business users
→Integration of information infrastructure with business processes:
→Flexible and resilient storage architecture

Components of SAN

A SAN consists of three basic components: servers, network infrastructure, and storage. These components can be further broken down into the following key elements: node ports, cabling, interconnecting devices (such as FC switches or hubs), storage arrays, and SAN management software.

1. **Node port**

   In fibre channel, devices such as hosts, storage and tape libraries are all referred to as nodes.

Each node is a source or destination of information for one or more nodes. Each node requires one or more ports to provide a physical interface for communicating with other nodes. These ports are integral components of an HBA and the storage front-end adapters. A port operates in full-duplex data transmission mode with a
transmit (Tx) and a receive (Rx)link.

2. **Cabling**

SAN implementations use optical fiber cabling. Copper can be used for shorter

distances for back-end connectivity, as it provides a better 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, multi-mode and single-mode.
Multi-mode fiber (MMF) cable carries multiple beams of light projected at different angles simultaneously onto the core of the cable (see Figure
6-4 (a)).

Based on the bandwidth, multi-mode fibers are classified as OM1 (62.5µm), OM2
(50µm) and laser optimized OM3 (50µm). 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 usually used for distances of up to 500
meters 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 6-4 (b)). These cables are available in diameters of 7–11 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 limits modal dispersion. Among all types of fibre 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, limited only by the power of the laser at the transmitter and sensitivity of the receiver.

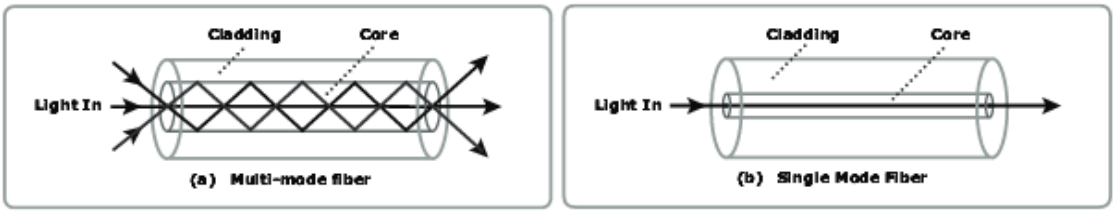

Figure 6-4: Multi-mode fiber and single-mode fiber

Connector

A Standard connector (SC) (see Figure 6-5 (a)) and a Lucent connector (LC) are two commonly used connectors for fiber optic cables. An
SC is used for data transmission speeds up to 1 Gb/s, whereas an LC is used for speeds up to 4 Gb/s. Figure
6-6 depicts a Lucent connector and a Standard
connector.A
Straight Tip (ST)
is a fiber optic connector with a plug and a socket that is
locked with a half-twisted bayonet lock (see Figure
6-5 (c)). In the early days
of FC deployment, fiber optic cabling predominantly used ST connectors. This

connector is often used with Fibre Channel patch panels.

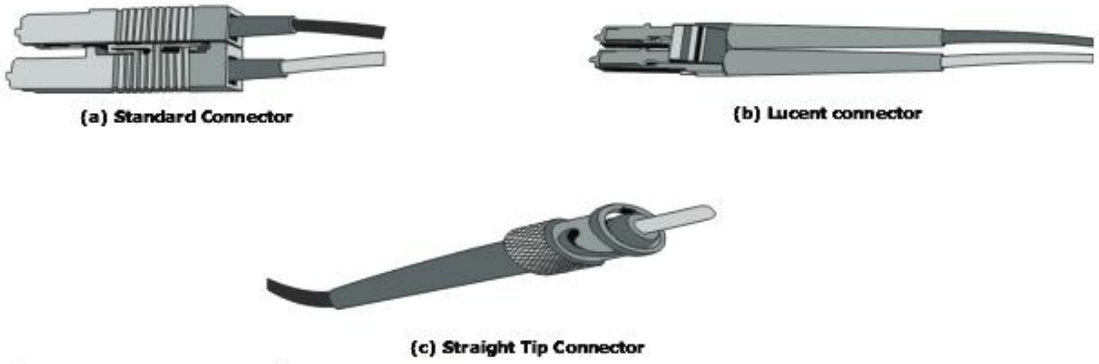

Figure 6-5: SC, LC, and ST connectors

Interconnect Devices

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

Hubs are used as communication devices in FC-AL implementations. Hubs physically connect nodes in a logical loop or a physical star topology.

Switches are more intelligent than hubs and directly route data from one physical port to another. Therefore, nodes do not share the bandwidth.

Directors are larger than switches and are deployed for data center implementations. The function of directors is similar to that of FC switches, but directors have higher port count and fault tolerance capabilities.

Storage array

SAN implementations complement the standard features of storage arrays by providing high availability and redundancy, improved performance, business continuity, and multiple host connectivity.

SAN Management software

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.

Fibre Channel Architecture

The FC architecture represents true channel/network integration with standard interconnecting devices.
Fibre Channel Protocol (FCP)is the implementation of serial SCSI-3 over an FC network.

FC-Protocol stack

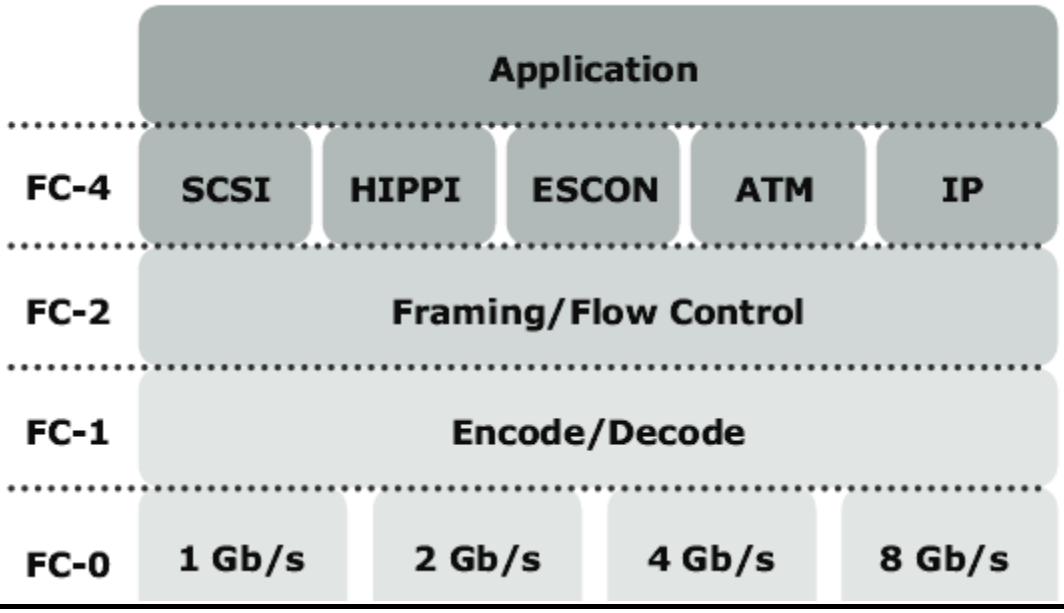

FC-4 Upper Layer Protocol

FC-4 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, HIPPI Framing Protocol, Enterprise Storage Connectivity (ESCON), ATM, and IP.

FC-2 Transport Layer

The FC-2 is the transport layer that contains the payload, addresses of the source and destination ports, and link control information. The FC-2 layer 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 Transmission Protocol

This layer defines the transmission protocol that includes serial encoding and decoding rules, special characters used, and error control. At the transmitter node, an 8-bit character is encoded into a 10-bit transmissions 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-0 Physical Interface

FC-0 is the lowest layer in the FCP stack. This layer defines the physical interface, media, and transmission of raw 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 Frame

An FC frame (Figure 6-17) consists of five parts:
start of frame (SOF)frame header datafield
cyclic redundancy check(CRC) and endofframe(EOF).
The SOF and EOF act as delimiters. In addition to this role, the SOF is a flag that indicates whether the frame is the first frame in a sequence of frames. The frame header is 24 bytes long and contains addressing information for the frame. It includes the following information: Source ID (S_ID), Destination ID (D_ID), Sequence ID (SEQ_ID), Sequence Count (SEQ_CNT), Originating Exchange ID (OX_ID), and Responder Exchange ID (RX_ID), in addition to some control fields.

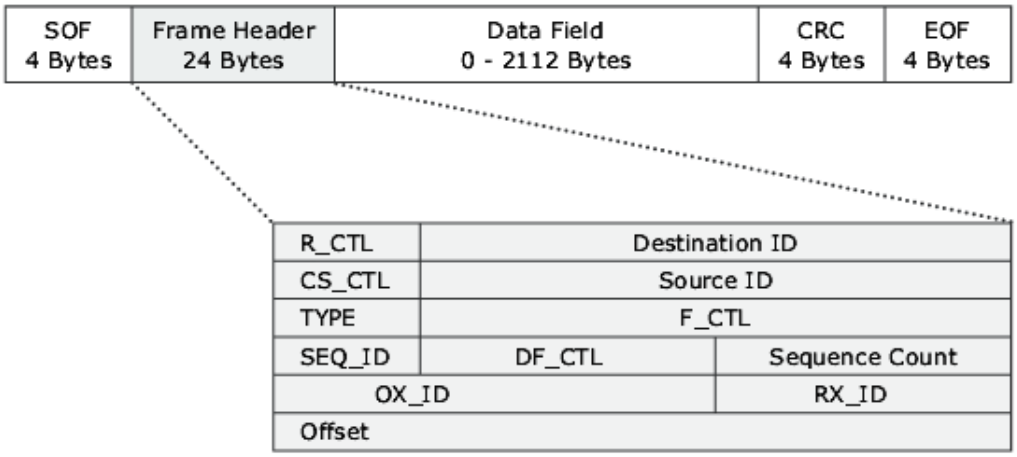

Figure 6-17: FC frame

The S_ID and D_ID are standard FC addresses for the source port and the destination port, respectively. The SEQ_ID and OX_ID identify the frame as a component of a specific sequence and exchange, respectively.
The frame header also defines the following fields:
Routing Control (R_CTL)
: This field denotes whether the frame is a link control frame or a data frame. Link control frames are nondata frames that do not carry any payload. These frames are used for setup and messaging. In contrast, data frames carry the payload and are used for data transmission.
Class Specific Control (CS_CTL)
: This field specifies link speeds for class 1 and class 4 data transmission.
TYPE
: This field describes the upper layer protocol (ULP) to be carried on the frame if it is a data frame. However, if it is a link control frame, this field is used to signal an event such as “fabric busy.” For example, if the TYPE is 08, and the frame is a data frame, it means that the SCSI will be carried on an FC.
Data Field Control (DF_CTL)
: A 1-byte field that indicates the existence of any optional headers at the beginning of the data payload. It is a mechanism to extend header information into the payload.
Frame Control (F_CTL)
: A 3-byte field that contains control information

related to frame content. For example, one of the bits in this field indicates whether this is the first sequence of the exchange.

FC topologies

Fabric design follows standard topologies to connect devices. Core-edge fabric is one of the popular topology designs. Variations of core-edge fabric and mesh topologies are most commonly deployed in SAN implementations.

1.Core-Edge Fabric

In the core-edge fabric topology, there are two types of switch tiers in this fabric. The edge tier usually comprises switches and offers an inexpensive approach to adding more hosts in a fabric.

The core tier usually comprises enterprise directors that ensure high fabric availability.

The core tier usually comprises enterprise directors that ensure high fabric availability.

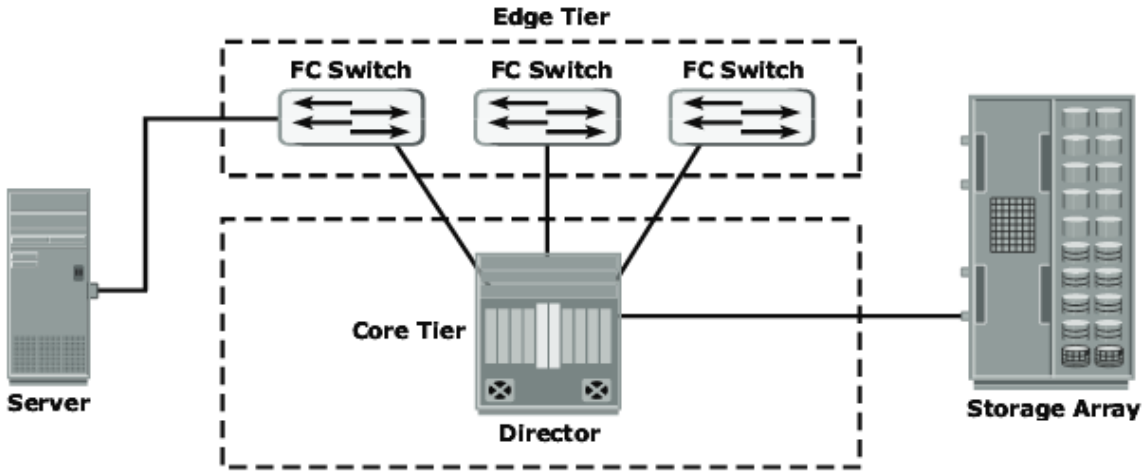

Figure 6-21: Single core topology

A dual-core topology can be expanded to include more core switches. However, to maintain the topology, it is essential that new ISLs are created to connect each edge switch to the new core switch that is added.

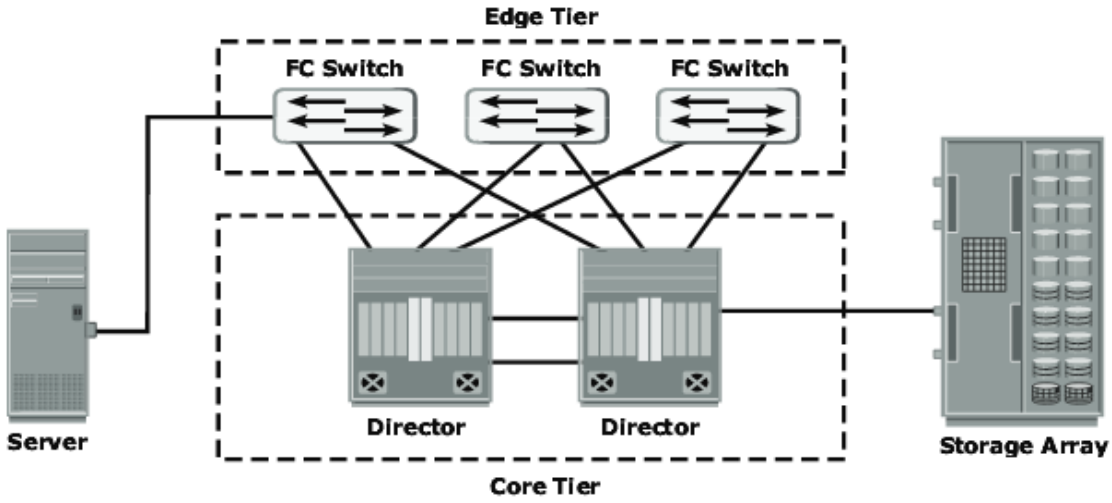

Figure 6-22: Dual-core topology

Mesh Topology

In a mesh topology each switch is directly connected to other switches by using ISLs. This topology promotes enhanced connectivity within the SAN. When the number of ports on a network increases, the number of nodes that can participate and communicate also increases.

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. 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.

In a partial mesh topology, several hops or ISLs may be required for the traffic to reach its destination. Hosts and storage can be located anywhere in the fabric, and storage can be localized to a director or a switch in both mesh topologies. A full mesh topology with a symmetric design results in an even number of switches, whereas a partial mesh has an asymmetric design and may result in an odd number of switches.

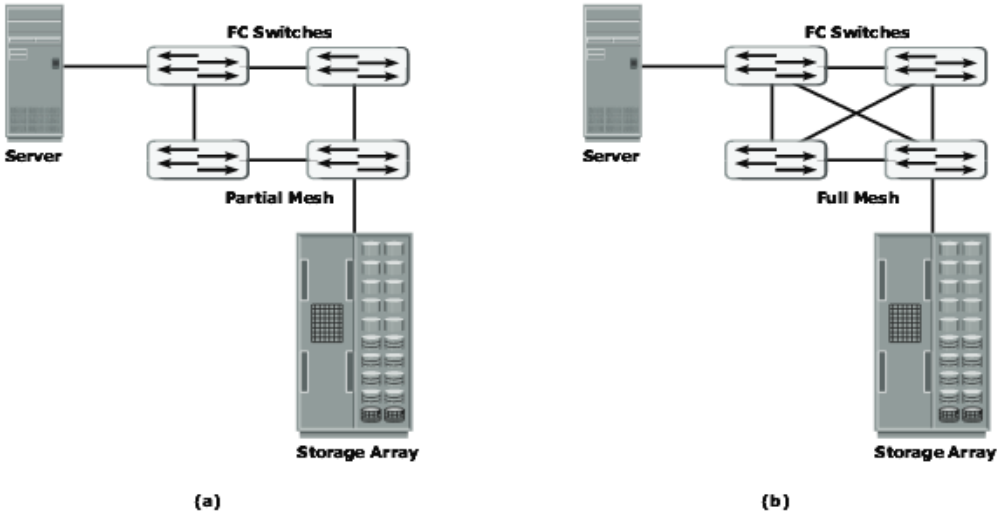

Figure 6-23: Partial mesh and full mesh topologies

NETWORK ATTACHED STORAGE

Network-attached storage (NAS)is an IP-based file-sharing device attached to a local area network. NAS provides the advantages of server consolidation by eliminating the need for multiple file servers. It provides storage consolidation through file-level data access and sharing. NAS is a preferred storage solution that enables clients to share files quickly and directly with minimum storage management overhead. NAS also helps to eliminate bottlenecks that users face when accessing files from a general-purpose server.

NAS uses network and file-sharing protocols to perform filing and storage functions. These protocols include TCP/IP for data transfer and CIFS and NFS for remote 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, 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.

NAS benefits

NAS offers the following benefits:
Supports comprehensive access to information:
→Enables efficient file
sharing and supports many-to-one and one-to-many configurations. The many-to-one configuration enables a NAS device to serve many clients simultaneously. The one-to-many configuration enables one client to connect with many NAS devices simultaneously.
→Improved efficiency:
Eliminates bottlenecks that occur during file access
from a general-purpose file server because NAS uses an operating system specialized for file serving. It improves the utilization of general-purpose servers by relieving them of file-server operations.
→Improved flexibility:
Compatible for 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, simplify data management, and ensures
greater data protection.
→Simplified management:
Provides a centralized console that makes it
possible to manage file systems efficiently.
→Scalability:
Scales well in accordance 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 networking components that provide maximum connectivity options. A NAS device can use clustering technology for failover.
→Security:
Ensures security, user authentication, and file locking in conjunction with industry-standard security schemas.

Components of NAS

A NAS device has the following components

NAS head (CPU and Memory)

One or more network interface cards (NICs), which provide connectivity to the network. Examples of NICs include Gigabit Ethernet, Fast Ethernet, ATM, and Fiber Distributed Data Interface (FDDI).

An optimized operating system for managing NAS functionality

NFS and CIFS protocols for file sharing

Industry-standard storage protocols

to connect and manage physical disk resources, such as ATA, SCSI, or FC

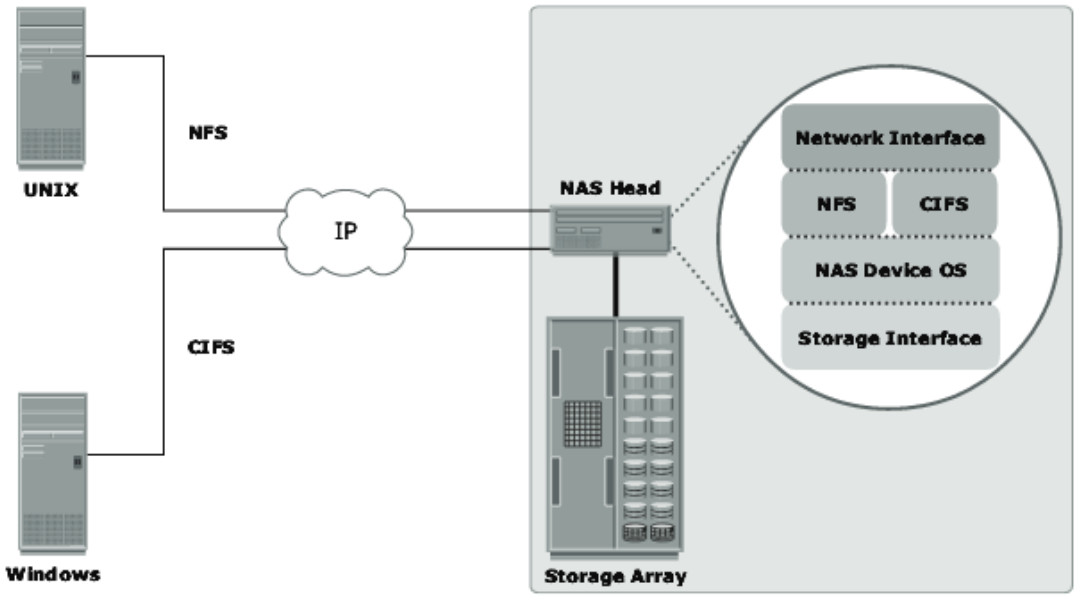

Figure 7-3: Components of NAS

IP –SAN

Traditional SAN environments allow block I/O over Fibre Channel, whereas NAS environments allow file I/O over IP-based networks. Organizations need the performance and scalability of SAN plus the ease of use and lower TCO of NAS solutions. The emergence of
IP technology that supports block I/O over IP has positioned IP for storage solutions.

Two primary protocols that leverage IP as the transport mechanism are iSCSI and Fibre Channel over IP (FCIP).

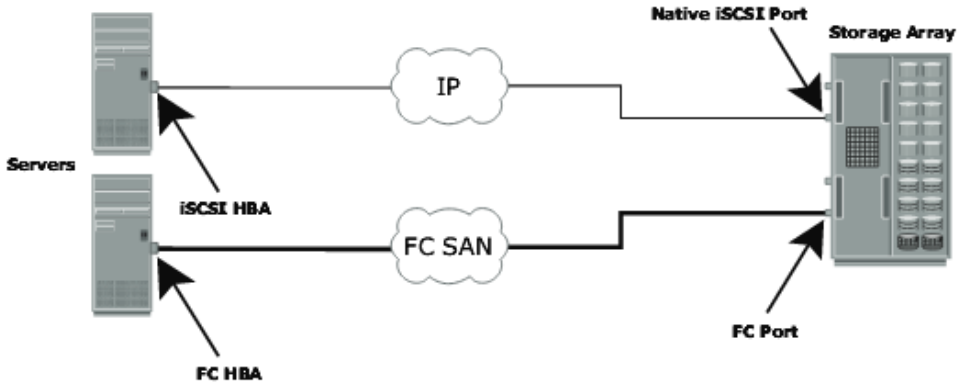

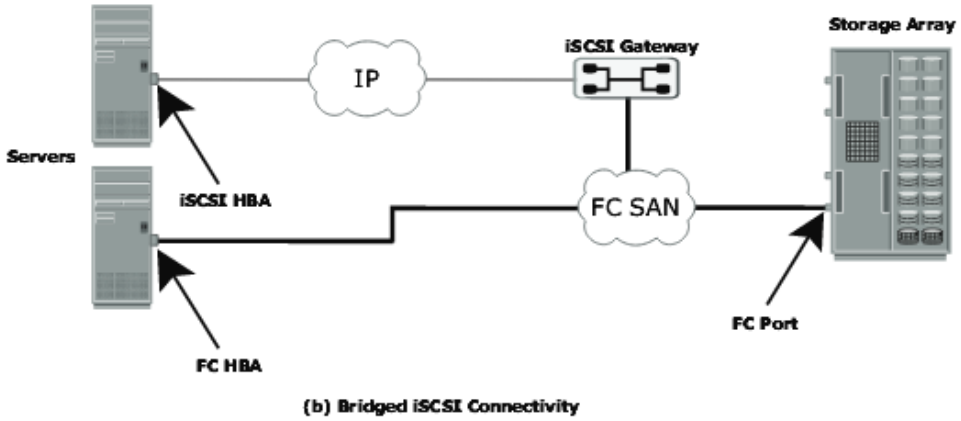

Figure 8-3: Native and bridged iSCSI connectivity

iSCSI Protocol Stack

The architecture of iSCSI is based on the client/server model. SCSI is the command protocol that works at the application layer of the 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 initiators and targets.

iSCSI is the session-layer protocol that initiates a reliable session between a device that recognizes 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 service.
TCP is used to control message flow, windowing, error recovery, and retrans mission. 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 for each hop through a separate physical network.

FCIP

Organizations are now looking for new ways to transport data throughout the enterprise, locally over the SAN as well as over longer distances, to ensure that data reaches all the users who need it. One of the best ways to achieve this goal is to interconnect geographically dispersed SANs through reliable, high-speed links. This approach involves transporting FC block data over the existing IP infrastructure used throughout the enterprise.
The FCIP standard has rapidly gained acceptance as a manageable, cost-effective way to blend the best of two worlds: FC block-data storage and the proven, widely deployed IP infrastructure. FCIP is a tunneling protocol that enables distributed FC SAN islands to be transparently interconnected over existing IP-based local, metropolitan, and wide-area networks. As a result, organizations now have a better way to protect, store, and move their data while leveraging investments in existing technology.
FCIP uses TCP/IP as its underlying protocol. In FCIP, the FC frames are encapsulated onto the IP payload. FCIP does not manipulate
FC frames (translating FC IDs for transmission).
When SAN islands are connected using FCIP, each interconnection is called an FCIP link.

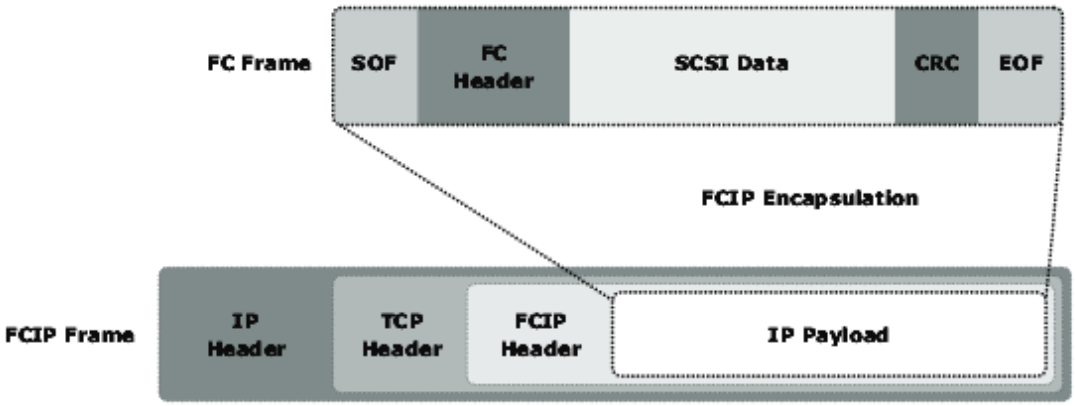

Figure 8-9: FCIP encapsulation

FCIP Topology

. An FCIP environment functions as if it is a single cohesive SAN environment.

Before geographically dispersed SANs are merged, a fully functional layer 2 network exists on the SANs. This layer 2 network is a standard SAN fabric. These physically independent fabrics are merged into a single fabric with an IP link between them.

An FCIP gateway router is connected to each fabric via a standard FC connection (see Figure 8-10). The fabric treats these routers like layer 2 fabric switches.

The other port on the router is connected to an IP network and an IP address is assigned to that port. This is similar to the method of assigning an IP address to an iSCSI port on a gateway. Once IP connectivity is established, the two independent fabrics are merged into a single fabric. When merging the two fabrics, all the switches and routers must have unique domain IDs, and the fabrics must contain unique zone set names. Failure to ensure these require ments will result in a segmented fabric. The FC addresses on each side of the link are exposed to the other side, and zoning or masking can be done to any entity in the new environment.

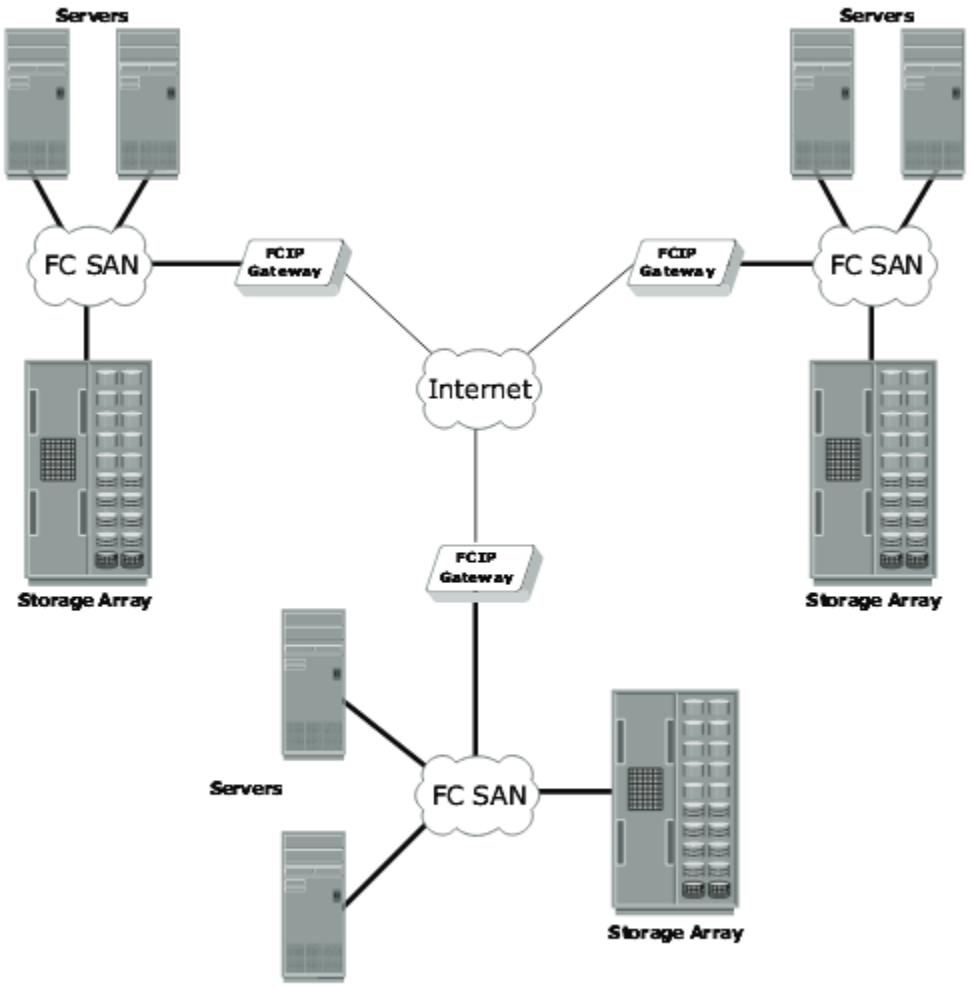

Figure 8-10: FCIP topology