This appendix covers the following topics:

- Understanding the Concepts of Storage Networking
- Understanding the Storage Networking Protocols
- Integrating SAN in Cisco Multilayer Switched Networks
- Implementing Cisco Solutions for Storage Networking

APPENDIX C

Introduction to Storage Networking

Previous chapters described the design and configuration of Cisco multilayer switched networks; this appendix describes the basic concepts of storage networking and its integration into Cisco multilayer switched networks. In recent years, the demand for high-performance, scalable, and highly available storage has grown significantly. This growth is due to the evolution of Internet-based solutions and computer-based applications that require storage, scaling, and management of large amounts of information for business operations in data centers.

Storage networking enables storage to be extended and scaled easily and efficiently using Fibre Channel, Fibre Channel over IP (FCIP), and SCSI over IP (iSCSI). These technologies extend beyond Direct Attached Storage (DAS) to high-performance, remote storage area networks (SAN). DAS is a technology to connect host devices to storage devices via direct methods such as parallel SCSI, while SAN technology interconnect hosts to storage devices using high-speed and scalable fabric switches.

Companies that are looking for additional terabytes of storage every year, and that want to reduce costs associated with integrating and managing DAS, are migrating to SAN environments. SAN technology presents a new horizon to the old legacy storage world by providing easier management, high performance, high availability, and scalability of vast amounts of storage.

Extending SANs over long distances has become a necessity of enterprise networks. A single SAN island with a single storage system does not satisfy the redundancy and disaster-recovery requirements necessary for enterprise networks in the event of an environment disaster or critical business interruption. As a result, many enterprise engineers architect disaster-recovery sites for the purposes of data recovery and business operation continuance over long distances using FCIP or native Fibre Channel.

Cisco, first in the SAN industry, provides complete intelligent solutions for integrating all storage services into a single package, thereby lowering the cost of ownership for today's storage requirements. As a result, organizations can easily incorporate SANs into their existing multilayer switched network using Cisco multilayer switches and Cisco SAN products. This also expedites the implementation of expanding SANs over WANs. Cisco storage products include the Cisco MDS 9000 family of switches and the SN 5400 family of routers.

This appendix is not a complete discussion of SAN and focuses primarily on SAN integration into multilayer switched networks. Many enterprise networks opt to keep SANs isolated to separate, autonomous networks. This appendix includes a brief introduction of storage networking, storage networking protocols, and SAN integration into Cisco multilayer switched networks. This appendix concludes with a discussion of Cisco storage products. Specifically, this appendix covers the following topics:

- Storage networking overview
- Storage protocols
- Integration of SANs with Cisco multilayer switched networks
- Cisco storage products and their features

Upon the completion of this appendix, you will understand the need for storage networking and storage networking protocols, and their application in the existing Cisco multilayer switched networks. You will also have a handle on Cisco storage solutions.

NOTE This appendix discusses topics outside the scope of the BCMSN certification test; however, this is helpful in understanding the basic concepts of emerging storage technologies that incorporate SANs into the multilayer switched network.

Storage Networking Overview

A SAN is a network of storage devices that connect to each other back-to-back or through a Fibre Channel switch or hub, as shown in Figure C-1. A storage switch is analogous to the LAN switch, which provides connectivity between servers or clients and storage devices, such as arrays, tape drives, and just a bunch of disks (JBOD), but the devices on a SAN use the Fibre Channel protocol to transmit I/O in the form of SCSI commands instead of using IP and Ethernet protocols.

Small Computer System Interface (SCSI) is the most widely used protocol in storage environments for sending commands and responses between hosts and storage devices. SCSI uses a client/server model. In SCSI, clients are the initiators and servers are the targets. Clients (initiators) send requests of Read or Write in the form of SCSI commands to storage devices (targets). SCSI runs over many I/O interfaces, including the parallel bus used in computers. Interfaces such as the parallel bus suffer from limitations in distance, speed and attachment capabilities, and scalability. SCSI over Fibre Channel, iSCSI, and FCIP, however, overcome those limitations.

Table C-1 describes a few examples of SCSI commands and their functionality.

Table C-1 *Examples of SCSI Commands*

SCSI Command	Functionality
read	Reads blocks of data from the storage devices
write	Writes blocks of data from the initiator to the target storage devices
read capacity	Obtains the block size and number of blocks
sense data	Obtains device information such as the vendor and the serial number

Furthermore, SANs provide high-performance, centralized management, reliability, and scalability that are not easily achieved with traditional DAS networks, as shown in Figure C-1. SAN technology enables storage devices to be shared across multiple hosts so that organizations can centrally manage storage.

Figure C-1 *DAS and SAN Networks*

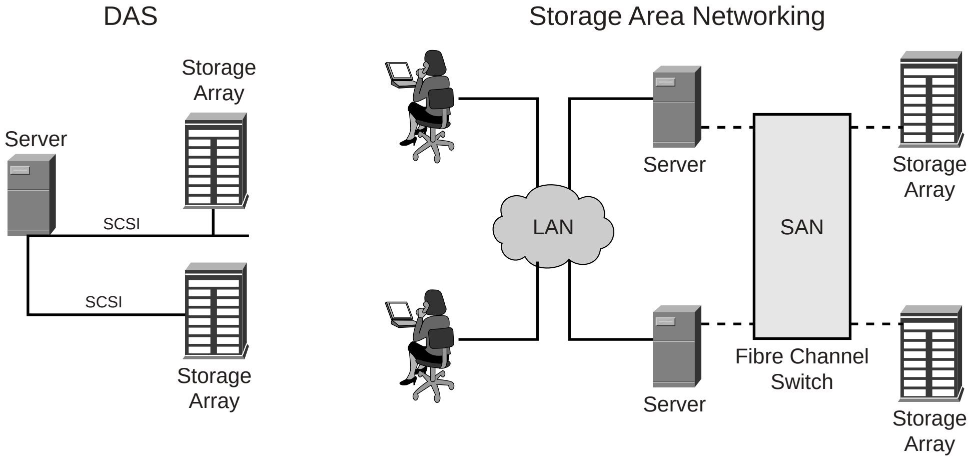

SANs easily integrate into the existing Cisco multilayer switched network via the Campus Backbone of the Enterprise Composite Network Model. By using iSCSI services, organizations can extend storage over the Campus Backbone directly to users in the Building Access submodules. In addition, iSCSI is extendable to remote storage through the Edge Distribution modules, via the WAN, to remote users and sites. FCIP enables peering between two isolated autonomous SANs over IP for enhanced utilization and integration of the dislocated storage devices.

Figure C-2 illustrates the integration of iSCSI, FCIP, and Fibre Channel services into the Campus Backbone model.

Figure C-2 *Storage Services in the Enterprise Infrastructure*

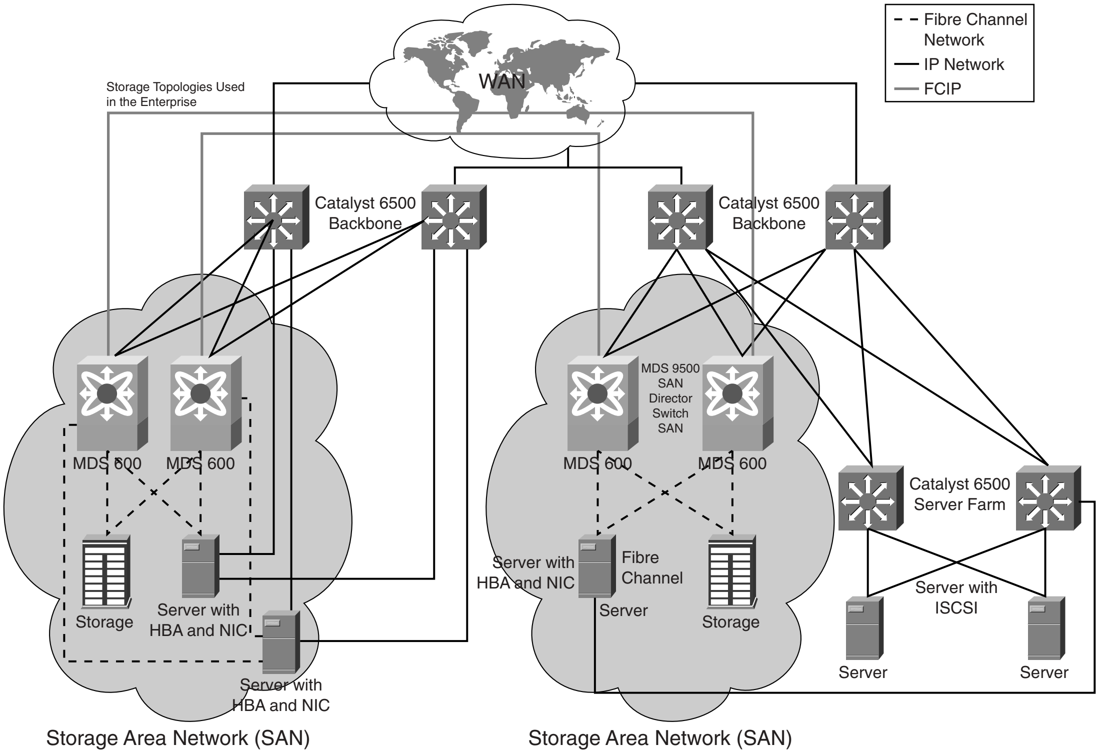

Storage Networking Protocols

The following list summarizes the common storage networking protocols supported by Cisco MDS switches:

- **Fibre Channel**—A protocol that transports SCSI commands between hosts and storage devices.
- **iSCSI**—A command protocol in storage networking used to transmit SCSI commands between hosts and storage devices over TCP/IP.
- **FCIP**—Encapsulates Fibre Channel protocol over TCP/IP to interconnect isolated SAN islands.
- **Fibre Connection (FICON)**—A high-speed, full-duplex I/O technology developed by IBM to connect mainframes to storage devices, at greater distances than the earlier Enterprise Systems Connection (ESCON). FICON uses a mapping layer based on the Fibre Channel technology for transporting data.

Introduction to Fibre Channel

Fibre Channel, a technology designed to transmit data between hosts and storage devices, is a widely open industry standard that broadens the ways for accessing and managing storage over long distances. Fibre Channel not only transports SCSI but is also capable of transporting other protocols, such as FICON and TCP/IP. The T11 organization defines the standards for Fibre Channel (www.t11.org).

Fibre Channel operates at speeds of 1, 2, 4, or 10 Gbps at full duplex. Although Fibre Channel standards have defined speed in excess of 10 Gbps, at the time of this book's publication, available products are limited to 4 Gbps at full duplex. Fibre Channel uses either fiber optics or electrical cables for the transmission medium. With optical fiber, Fibre Channel links can operate at distances of thousands of kilometers if connected to optical switches. The use of fiber in Fibre Channel greatly extends the operable distance of SANs.

Fibre Channel networks refer to end devices, such as servers, storage arrays, and tape drives, as nodes. Each node connects to the Fibre Channel network through an adapter called a host bus adapter (HBA). Furthermore, each node is composed of one or more Fibre Channel ports called node ports.

Just like IP, Fibre Channel is a structured protocol that has a standard architecture for addressing and network topologies. The following sections discuss in detail the following Fibre Channel topics:

- Fibre Channel architecture
- Fibre Channel addressing
- Fibre Channel topologies and port types

Fibre Channel Architecture

Like IP, Fibre Channel is based on a structured architecture (similar to the OSI model) that provides specifications from the physical layer to the upper-layer protocols. Figure C-3 shows the OSI and Fibre Channel models.

At the protocol level, to keep track of peer communications, IP uses TCP sessions, whereas Fibre Channel uses exchanges and sequences. An exchange manages the operation between the two ports in the transaction for one particular protocol type. All frames in the transaction have the same exchange ID. In Fibre Channel, each informational unit is sent using a sequence of one or more frames. An exchange can carry one or more sequences. If the information is not able to fit in one frame, it is broken into multiple frames with the same sequence number. Figure C-4 shows the components of the Fibre Channel protocol.

Figure C-3 *OSI and Fibre Channel Models*

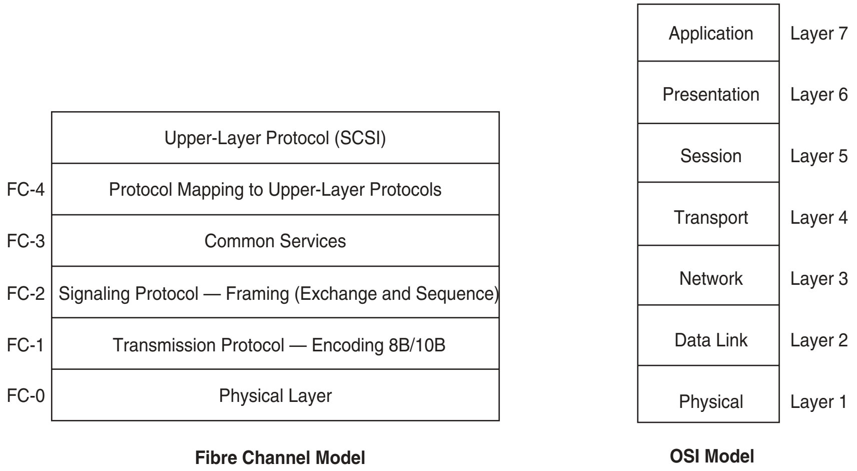

Figure C-4 *Components of the Fibre Channel Protocol*

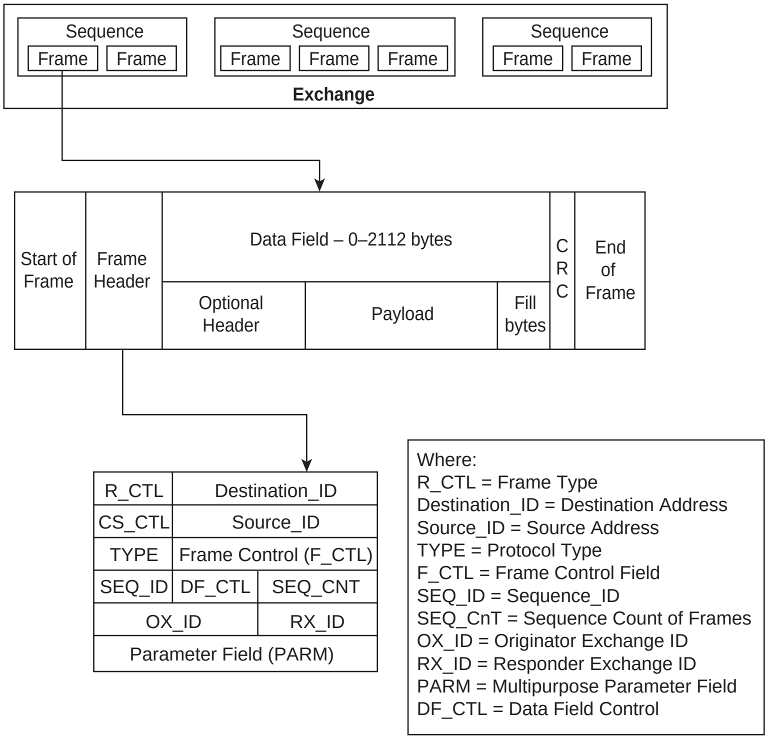

Where:
R_CTL = Frame Type
Destination_ID = Destination Address
Source_ID = Source Address
TYPE = Protocol Type
F_CTL = Frame Control Field
SEQ_ID = Sequence_ID
SEQ_CnT = Sequence Count of Frames
OX_ID = Originator Exchange ID
RX_ID = Responder Exchange ID
PARM = Multipurpose Parameter Field
DF_CTL = Data Field Control

NOTE Fibre Channel is basically a connection-oriented protocol that needs confirmation of delivery before it sends subsequent frames. Furthermore, Fibre Channel uses several transmission words in between the frames to perform control and signaling. The most common words are IDLE, Receive_Ready (R_Rdy), Start-of-Frame (SOF), and End-of-Frame (EOF). IDLE is usually sent when there is no information to send on the link. R_Rdy indicates that the receiver has emptied the buffer and is ready to receive another frame. SOF and EOF are sent to inform the receiver of the beginning and end of the frame. In addition, Fibre Channel also has defined multiple classes of service to transmit information such as Class 1, Class 2, Class 3, Class 4, and Class 6. Discussion of these classes is out of the scope of this appendix.

Fibre Channel Addressing

Each node in the Fibre Channel protocol has a unique node name represented by a 64-bit identifier referred to as the node World Wide name (nWWn), which is analogous to the MAC address in the LAN switching environment. Every vendor has a unique code prefix similar to a MAC address prefix. Similar to nWWn, each port is also represented by a 64-bit identifier called the port World Wide name (pWWn) for identification and authentication purposes. As mentioned earlier, each node connects to one or more ports, so each node may have more than one pWWn.

Fibre Channel uses 24-bit addresses similar to IP addresses, known as the Fibre Channel_ID (FCID), to route frames across the network. Because of the 24-bit addressing, the number of FCIDs in Fibre Channel networks is limited to 2^{24}. Each node port receives an FCID during port initialization. If a node moves to a different location in the network, the node may receive a new FCID. This appendix discusses FCID in more detail in the section titled "Switched Fabric."

Fibre Channel Topologies and Port Types

Fibre Channel uses one of three types of topologies as a medium. Each topology has its own bandwidth and performance considerations. The topology itself defines how the node port connects to the Fibre Channel network. Fibre Channel also describes different types of port types based on these topologies and functionalities. Fibre Channel supports the following topologies:

- Point-to-point
- Arbitrated loop
- Switched fabric

Table C-2 summarizes the different port types and their properties.

Table C-2 *Fibre Channel Port Types*

Fibre Channel Technologies	Port Type and Abbreviation	Scalability	Performance
Point-to-point	Node port (N)	2	Limited to the type of HBA
Arbitrated loop	Arbitrated loop node port (NL)	127	Limited to the loop, number, and type of devices in the loop; approximately 93% of loop speed with one port
Switched fabric (SF)	Node port that supports both point- to-point and arbitrated loops (Nx) Port that connects two switches (E) Fabric switch port that represents both F and FL ports (Fx) Fabric switch port connected to N port (F) Fabric switch port connected to arbitrated loop NL port (FL) Port that connects two switches over an EISL trunk (TE)	2^{24}	Line rate

Point-to-Point

In a point-to-point topology, as illustrated in Figure C-5, the two node devices connect back-to-back. This type of topology provides for guaranteed in-order delivery and bandwidth but is not scalable because there is only back-to-back, two-way communication. Adding multiple HBAs does not expand well for adding storage devices in this topology.

Figure C-5 *Point-to-Point Topology*

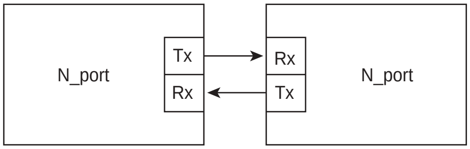

Arbitrated Loop

In the arbitrated loop topology, as many as 126 devices are daisy-chained together to form a loop, as shown in Figure C-6. Because all the devices reside on a shared medium, each node port (represented as NL) arbitrates access to the loop, thus reducing performance slightly compared to point-to-point topologies.

Due to shared bandwidth of the medium, this topology is best suited for the Fibre Channel disk arrays and small Fibre Channel topologies. Most new HBAs and disk arrays support both point-to-point and arbitrated loop modes.

Figure C-6 *Arbitrated Loop Topology*

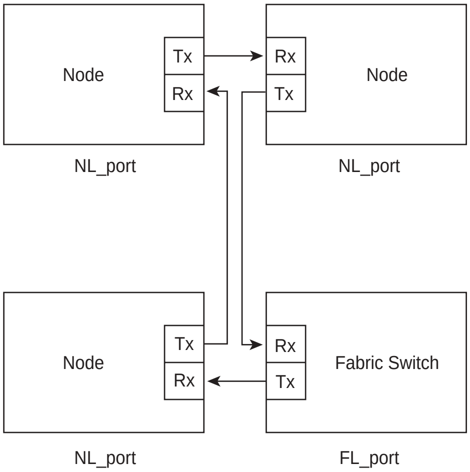

Switched Fabric

The switched fabric topology is the most commonly used topology for Fibre Channel. This topology is a collection of Nx ports, as illustrated in Figure C-7. The host and storage devices are connected to the switch using a switched fabric to communicate with each other. This topology provides for increased flexibility and security, higher bandwidth, higher availability, and scalability when compared to the point-to-point and arbitrated loop technologies.

A storage switch is analogous to the Ethernet switch. With a storage switch, the Fibre Channel hosts generally connect directly to switch interfaces. The ports of a Fibre Channel switch are called fabric ports (Fx ports, where F*x* represents F or FL). When fabric ports connect to point-to-point N ports, the fabric port operates in the F port mode. When fabric ports connect to arbitrated loops, the fabric ports operate in the FL port mode. A switch port connected to another switch port is represented as an expansion port (E_port), as shown in Figure C-7. The Fibre Channel switched topology is a distributed function topology, where each switch runs its own functions. To manage and control the distributed functions, each switch uses a Domain_ID, which is 8 bits long and the most significant byte of the FCID. A domain usually refers to a single switch. Each domain has a primary domain controller that manages and controls the communication among all the individual switch domain controllers in the fabric. The domain controllers are also responsible for assigning a dynamic FCID to ports during login.

Figure C-7 *Switched Fabric Topology*

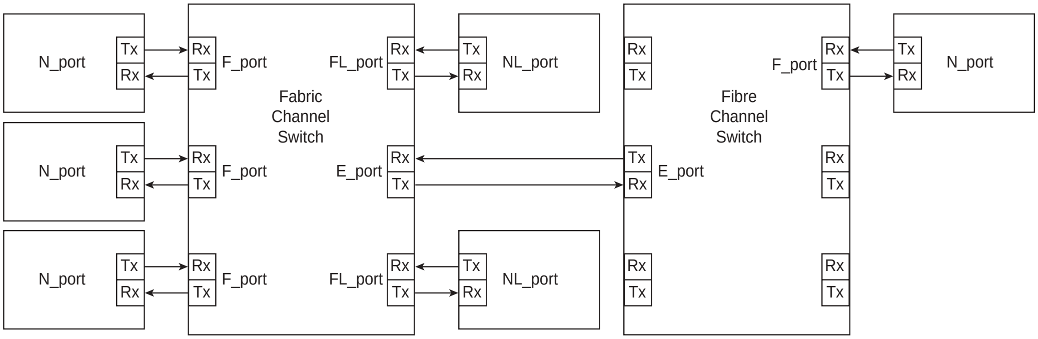

To keep track of the devices located in the fabric network and access the resources of the fabric, every node in the switch fabric network needs to log in with its pWWn, nWWn, and service parameters to the switch fabric. The login method is called Fabric Login (FLOGI). Once the node logs in, the fabric assigns a dynamic FCID based on the Domain_ID. For example, if the Domain_ID of the switch is 0x76, valid FCIDs are 0x760002, 0x760101, and so on. Furthermore, anytime a node establishes a connection to another node port, it must log in to the other port (i.e., Port Login [PLOGI]).

A complex Fibre Channel network may consist of between 1 and 239 switches, due to an address limitation. In multiswitch storage networks, the switches can be located over long distances, similar to multilayer switched networks. Fibre Channel uses the destination FCID to route the frame in the switch fabric. Similar to IP, Fibre Channel uses the Fabric Shortest Path First protocol (FSPF) routing protocol to compute the path to other switches. FSPF is a topology-based protocol similar to Open Shortest Path First (OSPF) in IP networking. If the target is on the same switch as the source, then the switch forwards the frame. If the destination is on a different switch, then the switch uses the route table to forward the frame to the correct destination.

In short, Fibre Channel is quite similar to the IP protocol in the switched fabric topology. Table C-3 compares Fibre Channel with the switched fabric topology and IP networking to show the similarities between the two protocols.

Table C-3 *Comparison of IP and Fibre Channel*

Feature	IP	Fibre Channel
Model	Follows OSI model	Follows Fibre Channel model
Connection-orientated protocol	Uses TCP sessions	Uses exchanges and sequences to keep track of sessions

Table C-3 *Comparison of IP and Fibre Channel (Continued)*

Feature	IP	Fibre Channel
Network devices	Layer 3 switches and routers	Fibre Channel switches and hub
Addressing mechanism	Uses 32-bit addressing for routing	Uses 24-bit addressing for routing
Mechanism to route frame	Uses destination IP address to route frame	Uses destination FCID address to route frame
Routing protocol	OSPF, BGP, RIP, etc.	FSPF

Introduction to iSCSI

iSCSI is a connection-orientated command and response protocol, which is actually a transport layer for the SCSI protocol over TCP/IP to connect IP hosts to reach storage devices on the Fibre Channel network. The main purpose of implementing iSCSI is to take advantage of existing IP networks to facilitate and extend SANs, as shown in Figure C-8, a sample iSCSI implementation in an IP network.

Figure C-8 *iSCSI Implementation*

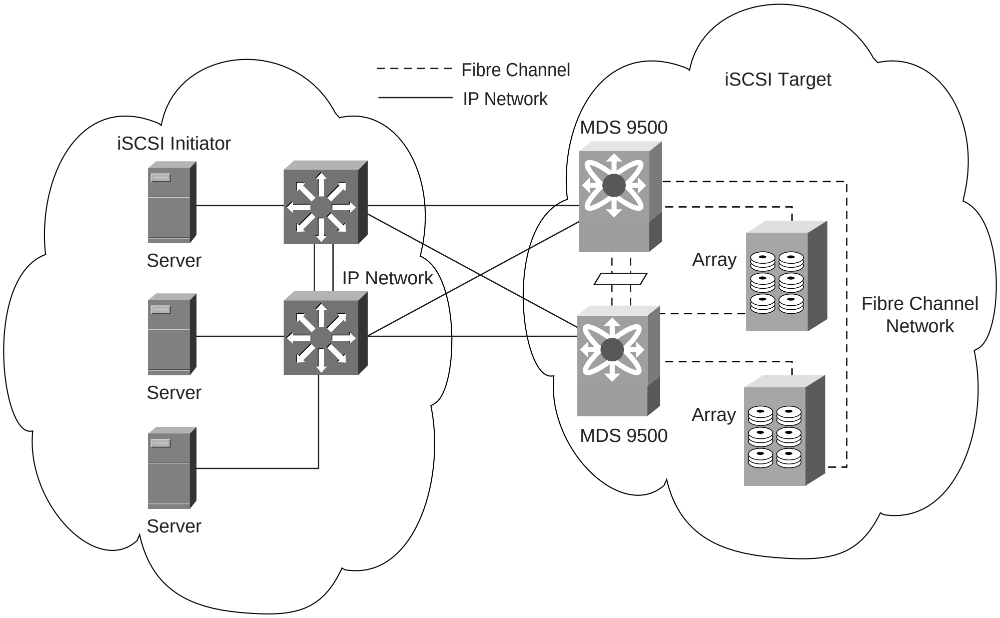

Because iSCSI runs over TCP/IP, it uses the traditional IP features such as compression, traffic shaping, security, and QoS techniques to optimize performance, scalability, and availability. Using iSCSI, a local host easily connects to the remote storage devices for block-level access similar to directly attached devices.

The iSCSI protocol, just like SCSI, uses the concept of an initiator and a target. In iSCSI, the initiator and targets are known as iSCSI nodes. The initiator, usually a server or desktop computer, initiates the session, while the target, usually a storage device (for example, JBOD or storage arrays) receives the commands.

IP gateway devices are responsible for transforming the iSCSI frame to Fibre Channel and vice versa. Figure C-9 logically depicts an iSCSI IP gateway, which interconnects the IP network and Fibre Channel networks. The IP gateway is a network entity that consists of one or more network portals. A network portal is a component for implementing the TCP/IP protocol stack. For an initiator, a network portal consists of an IP address; for a target, the IP address and listening port constitute the portal, as shown in Figure C-9. The initiator and target portal establish a TCP connection to communicate in iSCSI. When multiple TCP connections exist to the same target, the initiator forms a session to communicate to a specific target.

Figure C-9 *iSCSI IP Gateway*

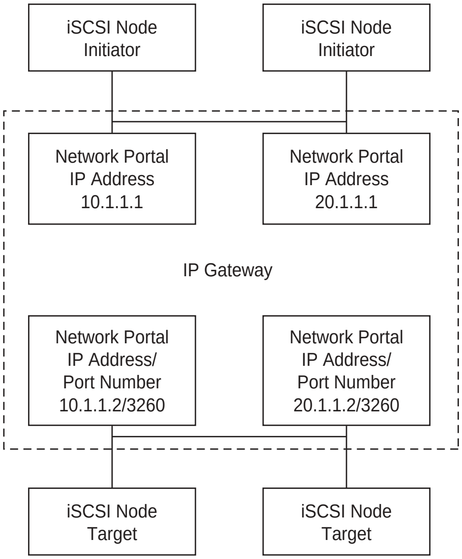

During this session, the initiator may add or remove TCP connections but only between the same initiator and target.

iSCSI uses two types of sessions: a normal session to send SCSI commands and a discovery session to discover targets.

The initiator uses the discovery process to identify targets available via the IP gateway. iSCSI uses the following methods to learn or discover targets:

- **Manual configuration**—With the manual process, a network administrator manually configures the IP address and the target name in the iSCSI driver of the initiator. The manual process is simple and useful for a small number of iSCSI targets.

- **iSCSI discovery process**—The iSCSI discovery process uses a **send target** command to retrieve a list of targets from an iSCSI target or IP gateway. This method is common for storage routers and gateways with changing targets or those with more than a few targets.
- **SLP-based discovery**—For larger storage environments, iSCSI designs typically use the Service Location Protocol (SLP), which uses multicast to discover servers that contain lists of iSCSI target names and IP addresses.
- **Internet Simple Name Service (iSNS)**—iSNS is another method for locating targets; it discovers and maintains a list of targets on centralized servers. In iSNS, both the initiator and target must support iSNS and communicate with the server to update the lists. For convenience, DNS names are usable in iSCSI.

Each iSCSI session goes through the following three phases:

- Login phase
- Operational negotiation phase
- Full-feature phase

When the initiator establishes a TCP connection to the target, the initiator starts the login phase. During the login process, the initiator and target authenticate with each other or only the initiator authenticates with the target, using various security protocols. After the login process is completed, the initiator may start the optional negotiation process to negotiate the other TCP session parameters. Afterward, the initiator moves to the final full-feature phase. When the login process is completed, the session may send SCSI commands.

In iSCSI, similar to the login sequence, the initiator is responsible for sending the logout command to gracefully close the session. In addition, a target may send the logout request to end the session or if the target gets an internal error in an asynchronous iSCSI message. In any case, the initiator sends the logout request to finish the session and the target responds with the logout response to complete the logout process.

Example C-1 shows an excerpt from a sample configuration for iSCSI. In this example, an MDS 9509 Fibre Channel switch connects directly to a Catalyst 6500 on interface GigabitEthernet 2/1. The iSCSI gateway IP address, as shown in the example, is 10.1.1.1. All the clients are attached to the Catalyst 6500 running Cisco IOS in this case and can establish iSCSI sessions to the switch. There are two initiators, which are trying to connect to two targets.

Example C-1 *Sample iSCSI Configuration for the MDS 9000 Family of Switches*

```
MDS9500:
!
(text deleted)
!
iscsi initiator name iqn.1987-05.com.cisco:02.96ea87ea4403.san-h31
nWWN 20:04:00:0b:fd:06:37:42
pWWN 20:07:00:0b:fd:06:37:42
iscsi initiator name iqn.1987-05.com.cisco:02.c49171c8e2a6.san-w2k3
```

continues

Example C-1 *Sample iSCSI Configuration for the MDS 9000 Family of Switches (Continued)*

```
nWWN 20:04:00:0b:fd:44:66:42
pWWN 20:03:00:0b:fd:06:37:42
!
(text deleted)
!
iscsi virtual-target name iSCSI-2
pWWN 21:00:00:04:cf:67:d5:f6
advertise interface GigabitEthernet2/1
all-initiator-permit

iscsi virtual-target name iSCSI-3
pWWN 21:00:00:04:cf:bd:56:59
advertise interface GigabitEthernet2/1
all-initiator-permit
!
(text deleted)
!
interface GigabitEthernet2/1
ip address 10.1.1.1 255.255.255.0
iscsi authentication none
no shutdown
!
(text deleted)
!
interface iscsi2/1  (matching GiGE int is 2/1)
no shutdown
!
(text deleted)
!
```

```
Cat6500:
!
(text deleted)
!
interface vlan 40
ip address 10.1.1.2 255.255.255.0
no shutdown
!
(text deleted)
!
```

Introduction to FCIP

Fibre Channel over IP is used to interconnect isolated SAN islands over elongated distances, as shown in Figure C-10. Also, Figure C-10 shows the Eastside SAN interconnecting with the Westside SAN through FCIP over an IP network to replicate data.

FCIP encapsulates Fibre Channel frames into TCP/IP for tunneling those frames over an IP network for isolated SANs. In this manner, the IP network is transparent to the isolated SANs. Organizations can better manage and scale the storage performance, scalability, and availability of hosts and targets of a Fibre Channel network by using FCIP to interconnect various isolated SANs in an existing IP infrastructure.

Figure C-10 *FCIP Implementation*

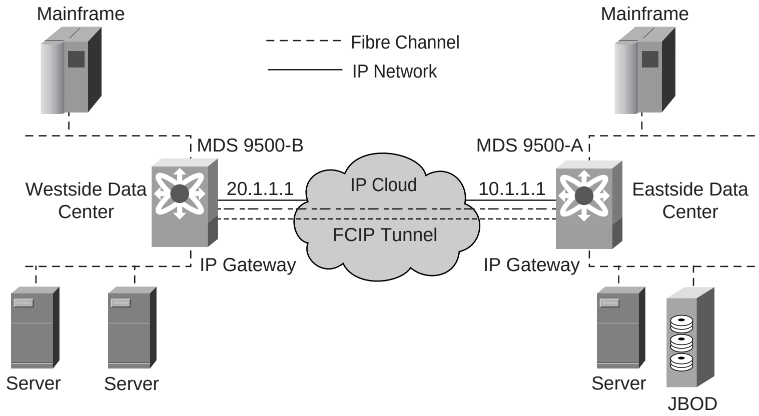

As disaster-recovery solutions extend their reach to distant data centers, replicating, copying, migrating, and vaulting data over TCP/IP from a main site to a remote site has become a valuable solution. FCIP is suited for long distances with Async replication where host-delay is not dependent on the WAN infrastructure. With Cisco MDS switches, enterprises can design their data-replication, data copy, or data migration solutions using the Cisco FCIP solution without building separate transport networks for Fibre Channel. Furthermore, due to buffer constraints and requirements for a dedicated transport, the feasibility of extending SANs well beyond several hundred kilometers is limited with native Fibre Channel. In addition, FCIP solutions offer the ability to share existing bandwidth with Ethernet/IP solutions and to reduce the costs of having dedicated transports, such as optical transports.

Example C-2 shows a sample configuration of FCIP for two MDS 9500 switches connecting over an IP network. The MDS 9500A switch is using the IP address of 10.1.1.1 on interface GigabitEthernet 2/1 to tunnel through the IP network to the MDS 9500B switch using the IP address of 20.1.1.1 on interface GigabitEthernet 2/1. Figure C-10 illustrates this topology.

Example C-2 *Sample FCIP Configuration for the MDS 9500 Family of Switches*

```
MDS9500_A:
!
(text deleted)
!
fcip profile 1
ip address 10.1.1.1
interface fcip1
no shutdown
!
(text deleted)
!
use-profile 1
peer-info ipaddr 20.1.1.1
```

continues

Example C-2 *Sample FCIP Configuration for the MDS 9500 Family of Switches (Continued)*

```
!
(text deleted)
!
interface GigabitEthernet2/1
ip address 10.1.1.1 255.255.255.0
-
no shutdown
!
(text deleted)
!
```

```
MDS9500_B:
!
(text deleted)
!
fcip profile 1
ip address 20.1.1.1
!
(text deleted)
!
interface fcip1
no shutdown
use-profile 1
peer-info ipaddr 10.1.1.1
!
(text deleted)
!

interface GigabitEthernet2/1
ip address 20.1.1.1 255.255.255.0

no shutdown
!
(text deleted)
!
```

Storage Network Integration

To meet today’s storage requirements for performance, scalability, and availability, network designs incorporate SANs as isolated SANs or as networks connected to the Campus Backbone submodule of the Enterprise Composite Network Model as a Server Farm module. When connecting into Ethernet networks, the MDS 9500 switches use IP Storage (IPS) line modules. When connected to the Campus Backbone submodule, SANs may provide iSCSI services and connect with remote SANs over WANs and metro Ethernet networks using FCIP. As shown in Figure C-11, iSCSI helps extend the SAN to low- to midrange servers while FCIP helps connect SAN islands over IP.

Figure C-11 *Campus Network Integration*

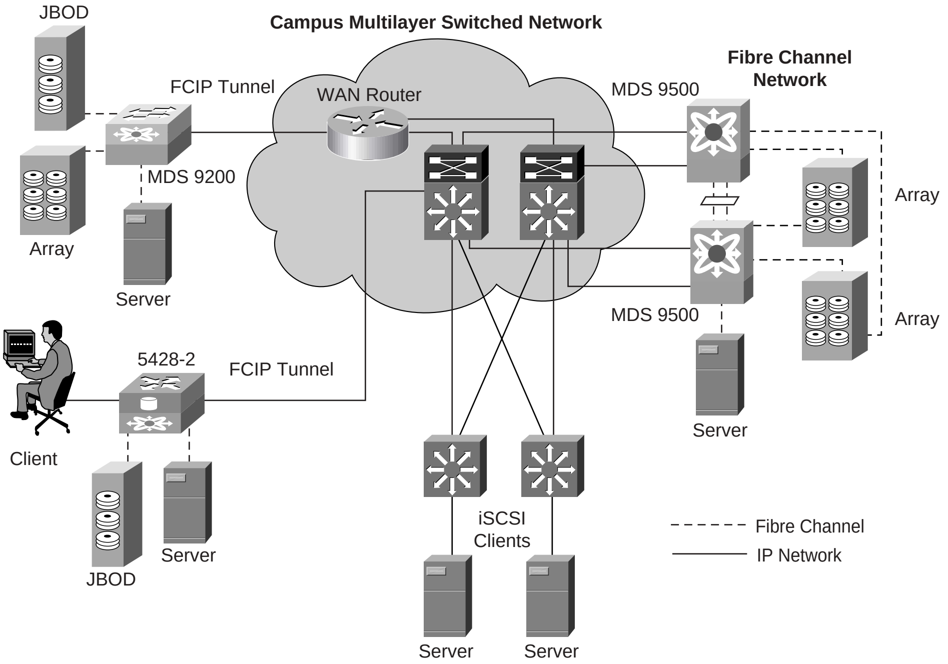

Cisco Storage Solutions

With the rapid increase in storage demand, organizations are looking for reliable solutions that are easily scalable and manageable. Cisco offers various storage products that allow enterprises to deploy manageable, reliable, and cost-effective storage networking solutions by providing multiple services in a single chassis.

Cisco offers the following two storage product families:

- Cisco MDS 9000 family of switches
- Cisco SN 5400 iSCSI storage routers

Cisco MDS 9000 Multilayer Switches

The Cisco MDS 9000 family of modular, multilayer directors and fabric switches are the industry’s first fully integrated, multiprotocol, storage networking platform. The Cisco MDS 9000 family of multiprotocol switches integrates iSCSI, FCIP, and Fibre Channel solutions into a single chassis to lower the cost of ownership. At the time of publication, the Cisco MDS 9000 family of switches consists of the following three models:

- **Cisco MDS 9500**—The switches in this series are high-performance, multilayer, multiprotocol, director-class switches, optimized for core enterprise storage networking. All members of the Cisco MDS 9500 family of switches provide high speed, high availability, security, scalability, and ease of management. The Cisco MDS 9500 family of switches supports up to 256, 1- or 2-Gbps, auto-sensing Fibre Channel ports in a single chassis with a nonblocking backplane fabric of 1.44 Tbps.
- **Cisco MDS 9200**—The switches in this series are two-slot, modular, multiprotocol fabric switches. The 9200 family is available in 9216, 9216a, and 9216i models. The MDS 9200 family of switches shares the same architecture as the MDS 9500 series. The MDS 9200 family of switches has 16, 1- or 2-Gbps, auto-sensing Fibre Channel ports and the expansion slot to accommodate any MDS 9500 line cards. With the 32-port Fibre Channel line card, the MDS 9200 family of switches incorporates up to 48 ports. The main difference between the models is that the 9216i has 14 Fibre Channel ports and two Gigabit Ethernet ports for FCIP or iSCSI connectivity in addition to support for 3500 buffer-to-buffer credits on the Fibre Channel ports.
- **Cisco MDS 9100**—The MDS 9100 family of switches is available in 20- and 40-port models: the Cisco MDS 9120 and Cisco MDS 9140. Switches in the 9100 series have a fixed configuration of one rack unit supporting 1- or 2-Gbps switches. It is similar to the MDS 9200 series in all features except the multiprotocol support. Enterprises can use this model to build small and midsize SANs or to provide edge-to-core connectivity in larger SANs.

Table C-4 illustrates the feature comparison of two MDS models of switches.

Table C-4 *Feature Comparison Between MDS 9200 and MDS 9500*

Feature	MDS 9200	MDS 9500
Available slots	1	9, 6
System bandwidth	80 Gbps	1.44 Tbps
Redundant Supervisors	No	Yes
Max ports	48	224
Max iSCSI/FCIP ports	8	48

The Cisco MDS 9500 switches are director-class switches positioned for the core, while the MDS 9200 and 9100 switches are midrange fabric switches positioned for the distribution and access layers.

In brief, the MDS 9000 family of switches supports the following value-added features compared to other vendor fabric switches:

- Virtual Storage Area Network (VSAN)
- Port channeling

- Trunking
- QoS and congestion control
- Multipath forwarding
- Zoning
- Buffer-to-buffer credit
- SPAN
- iSCSI/FCIP
- IVR

The features are described in the sections that follow.

VSAN

A VSAN in a SAN environment is analogous to the VLAN in the LAN switching environment. A VSAN has the behavior and the property of a single SAN. Each VSAN can have up to 239 switches with an identical address space and FCIDs.

Figure C-12 shows the implementation of VSANs.

Figure C-12 *VSAN Implementation*

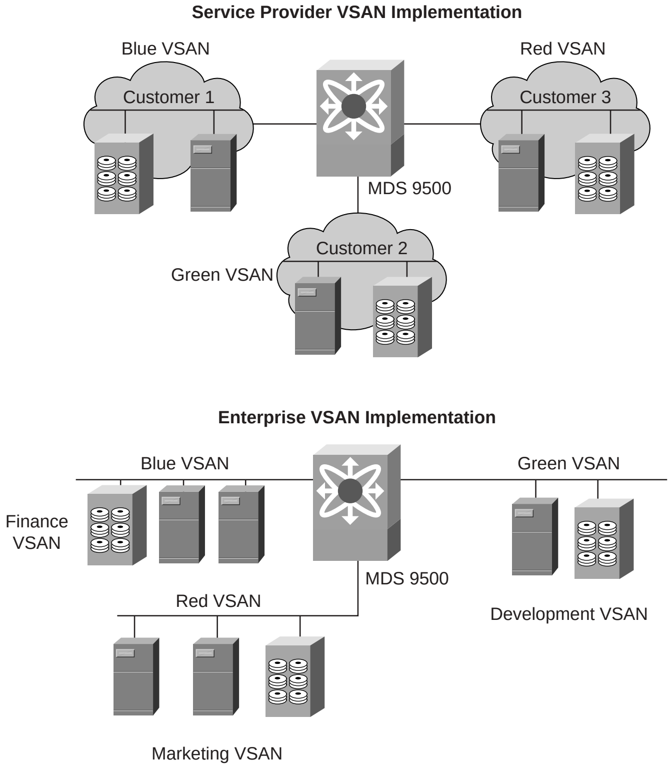

VSANs provide flexibility, security, and stability to the fabric by isolating the traffic to VSAN boundaries within one physical topology. For this reason, devices in one VSAN are not able to communicate with devices in other VSANs. The VSANs are usually created based on usage or technology. For example, in enterprise networks, one VSAN may exist for the financial department group, one for the marketing department, and one for the development team. In service provider topologies, the service provider may isolate each user or customer into its own VSAN. The MDS 9000 supports up to 256 VSANs with VSAN IDs between 1 and 4095.

Port Channel

The MDS supports port channeling to aggregate multiple physical interfaces into one logical interface. Port channeling inherently achieves load balancing for higher bandwidth and highly available link redundancy. Port channeling on the Cisco MDS platform supports 128 Port channels with up to 16 separate interfaces spanned across multiple modules in a single chassis. Port channeling is capable of the following two load-balancing mechanisms:

- **Flow based**—Load balancing based on source and destination FCID, where frames with the same source and destination use the same link.
- **Exchange based**—Load balancing based on source and destination FCID and exchange ID, where frames with the same source and destination FCID and with the same exchange ID pass along the same link. This method of load balancing is the preferred one because it preserves the transmitted order of the frame. Out-of-order frames are detrimental in Fibre Channel. This is the default load-balancing mechanism.

Trunking

SAN trunking is analogous to the ISL trunking that carries more than one VSAN over the same physical link using Extended ISL (EISL). Trunking provides the facility to interconnect switches to distribute VSANs over the fabric network. With trunking, the VSAN easily extends to multiple switches. The ports that connect the switches through EISL are known as Trunking E ports (TE ports). Cisco is the first company to introduce this feature, as it did with ISL trunking in the Catalyst switching products. Figure C-13 illustrates a trunk port carrying multiple VSANs.

Figure C-13 *Trunk Implementation*

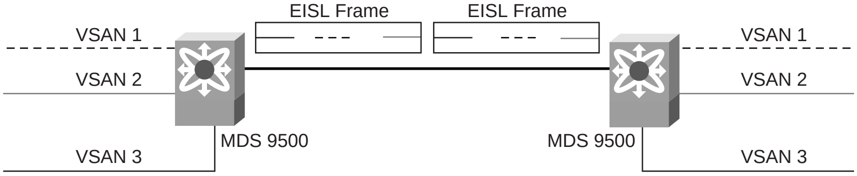

Multipath Forwarding

MDS 9000 switches can load-balance across 16 parallel paths based either on source and destination FCID or on source and destination FCID along with exchange ID. The default behavior is to load-balance based on source and destination FCID and exchange ID to preserve the order of the frame within an exchange.

Congestion Control and QoS

The MDS 9000 switch supports QoS for internally and externally generated control traffic sourced from the Supervisor and other vendor switches. QoS enables prioritization of time-critical control traffic over other types of traffic.

The MDS 9000 family of switches also supports Forward Congestion Control (FCC) to reduce congestion in the fabric without affecting other Fibre Channel traffic. When a switch detects a congestion port in the network, the switch generates an edge quench message to the sources as an alert to reduce the rate at which frames should be injected into the fabric to avoid further head-of-line blocking.

Zoning

Zoning restricts communication between devices connected to fabric ports in the same fabric. Switches use zones to restrict communications. Zones are similar to the VLAN access control lists (VACL) that restrict communication even with the same broadcast domain. There can be more than one zone in the same fabric. Because each VSAN is an individual fabric, it runs its own FSPF, domain controllers, and zone manager. The same zone names and IDs can exist in different VSANs but not in the same VSAN. Devices that belong to the same zone can communicate with members in the same zones but cannot access any members in other zones; however, one device may belong to more than one zone. Figure C-14 shows the implementation of zoning.

Zone membership is defined by any one of the following types:

- **pWWn**—The pWWn of an N port attached to the switch.
- **Fabric pWWn**—The WWn of the fabric port (switch port’s WWn), also known as port-based zoning.
- **FC ID**—The FC ID of an N port attached to the switch.
- **Symbolic node names**—Symbolic node names include iSCSI node names.

Buffer-to-Buffer Credits

Fibre Channel also has an important built-in flow control mechanism based on buffer-to-buffer credits. As mentioned earlier in the appendix, Fibre Channel protocol needs an acknowledgment for every frame it sends that may result in a longer round-trip delay, especially at larger distances. Such a delay implies that the acknowledgments from the receiver of Fibre Channel frames are delayed, and consequently delay the ability to transmit frames at the source and reduce the effective throughput. A system that supports a large number of buffer-to-buffer credits can send more frames before it has to wait for acknowledgments over distances that exceed 100 km. Cisco MDS 9000 series switches support 3500 buffer-to-buffer credits per E port with the new product line, which enables support for line rates with 2148 kB frame sizes at 2 Gbps for distances more than 3000 km.

Figure C-14 *Zone Implementation*

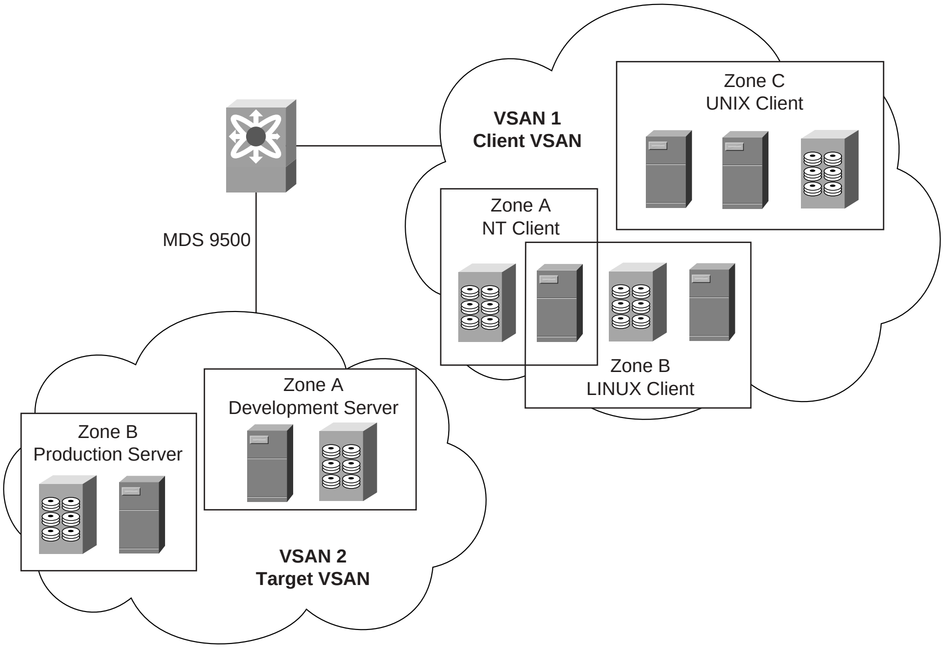

SPAN

The MDS 9000 switch also provides for SPAN functionality for ease in troubleshooting. SPAN monitors network traffic through Fibre Channel interfaces. Traffic through any Fibre Channel interface is replicated to a designated port called the SPAN destination port (SD port) for network analysis.

iSCSI/FCIP

The Cisco MDS 9000 family of switches uses an iSCSI and FCIP module, WS-X9308-SMIP, WS-X9304-SMIP, Cisco MDS 9216i, and DS-X9302-14K9 (also known as fourteen+2 cards) to provide iSCSI and FCIP services in the same chassis with Fibre

Channel services. The following list illustrates the key features of the MDS 9000 iSCSI and FCIP module:

- Offers three FCIP tunnels between source and termination points
- Supports 32 MB of buffering to scale TCP effectively
- Offers iSCSI functionality on each interface
- Supports port channeling for fault tolerance, redundancy, and trunking for carrying Fibre Channel VSANs over long distances
- Supports the functionality to easily integrate with storage network devices such as SN 5428-2 routers and the port adapter for Cisco 7200/7400 routers
- Supports TCP extension parameters to take the benefits of full bandwidth over long distances, which is not possible with standard TCP
- Supports write acceleration and tape acceleration by spoofing the Transfer Ready, which helps the hosts to write IOs faster over long distances.

Inter-VSAN Routing (IVR)

Inter-VSAN routing (IVR) offers the facility to share resources between VSANs in a Fibre Channel Fabric without merging VSANs. Devices such as tape drives and storage systems that are part of one VSAN can be accessed across multiple VSANs using IVR. Without IVR, each of these devices has to be physically connected to each VSAN to access devices on the particular VSAN. IVR creates a path such that only devices that need to access the resources across the VSANs can communicate, maintaining VSAN isolation.

In addition, IVR provides more efficient business continuity solutions for SAN extensions over FCIP. An example of IVR is having a storage system at a local site in one VSAN, a remote storage system in another VSAN, and a third VSAN, which is called a "Transit VSAN," that carries the traffic between the VSANs. In this manner, isolating SANs into VSANs limits any disruption on either fabric to that particular VSAN where the disruption occurred. As a result, IVR provides an extra layer of resiliency by providing a control boundary between the VSANs where only related FSPF routes and name servers entries are exchanged and modified in the transient VSAN. This characteristic of IVR reduces the chance that anomalous Fabric events will spread over the long distances, as shown in Figure C-15. In Figure C-15, two storage arrays are connected to replicate data and backup servers to back up data to tape drives over FC or FCIP using transit VSANs, which provides an additional layer of resiliency to the Fabric over long distances. Refer to the configuration guide for SAN-OS 2.0 for details on configuring IVR:

> Cisco MDS Configuration Guide: http://www.cisco.com/en/US/products/hw/ps4159/ps4358/products_configuration_guide_book09186a00802e22e6.html

In addition to the preceding features, the MDS 9000 also supports many other important features such as SSH, RADIUS, SNMP, SNMP version 3, high availability via component redundancy, web management, and call home.

Figure C-15 *IVR Implementation*

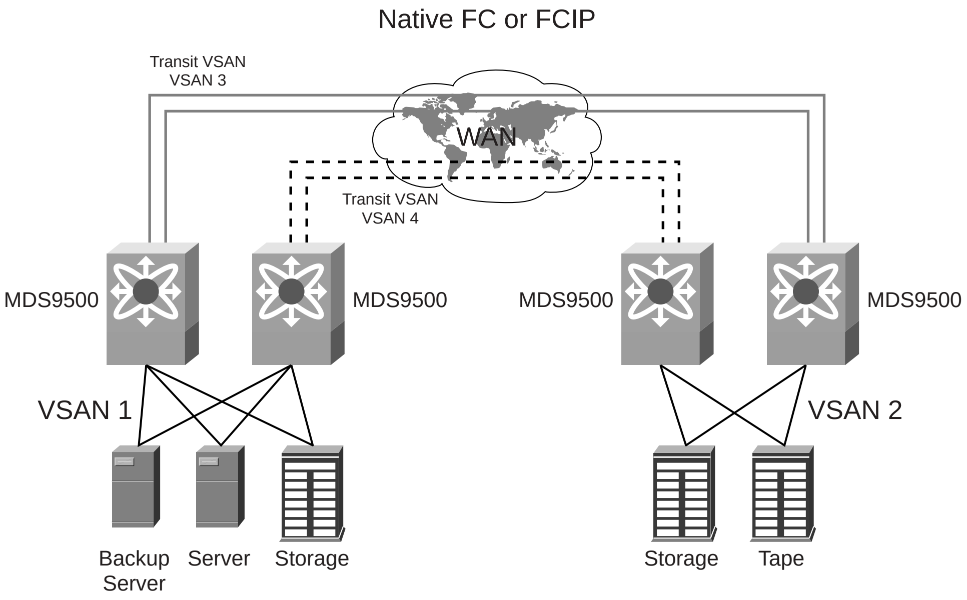

In summary, the Cisco MDS 9000 family of switches is a robust multiprotocol platform that provides Fibre Channel, iSCSI, and FCIP features in one platform. The MDS 9500 series easily integrates all future SAN protocols and technologies for investment protection.

NOTE For more information about the Cisco MDS products and features, refer to the following document on Cisco.com:

"Cisco MDS 9000 Series Multilayer Switches," http://www.cisco.com/en/US/products/hw/ps4159/ps4358/index.html

Cisco SN 542x iSCSI Routers

The Cisco SAN 5400 series routers are midrange storage solutions for cost-effective migration from direct-attached storage to SANs using iSCSI. The SN 5400 platform includes the following three models:

- SN 5420
- SN 5428
- SN 5428-2

NOTE Cisco SN 542X series routers are end-of-life, and a discussion of these routers is out of the scope of this appendix.

Summary

With the rapid increase in storage demand and the consolidated storage over the WAN, storage networking is becoming increasingly popular compared to DAS storage. Fibre Channel, iSCSI, and FCIP are important protocols of SANs. Organizations now have the capability to build Fibre Channel SANs to connect hosts to the storage devices, such as arrays and tape drives, at high speeds with scalability and highly available infrastructures.

To expand SANs over long distances, integrate low and midrange servers into a centralized SAN, and provide connectivity between multiple data centers, enterprises are choosing to use FCIP and iSCSI services. Cisco is the first in the industry to provide all the services in the single platform.

The Cisco MDS 9000 family is a robust multiprotocol switch that integrates all solutions and is designed to support future storage protocols and technologies.