# Cisco MDS 9000 Family – The Enterprise-Class SAN Infrastructure for Unified I/O

## Contents

**What You Will Learn** ..... **2**
**Overview** ..... **2**
**Principles of SAN Design** ..... **2**
High Availability and Dual (A+B) Fabric ..... 2
Oversubscription and Fan-out Considerations ..... 3
Topologies of Individual Fabrics ..... 4
**Cisco MDS 9000 Family Optimization of Cisco Nexus 5000 Series SAN Connectivity** ..... **4**
Fabric Virtualization: VSANs ..... 5
Network Port Virtualization Mode ..... 5
Network Edge in NPV Mode ..... 5
NPV Scalability Considerations ..... 7
Sharing Devices Across VSANs ..... 8
Network-Level High Availability: PortChannel ..... 9
Transparent Fabric Services: Zoning and Aliases Services ..... 10
Distributed Device Alias Service ..... 11
Homogeneous Authentication, Authorization, and Accounting ..... 11
Unified Fabric Access Security ..... 11
Configuration Synchronization with Cisco Fabric Services ..... 12
Quality of Service ..... 13
Simplified Operations and Training ..... 13
Multiprotocol Options: Fibre Channel, iSCSI, and FCoE ..... 14
Small Computer System Interface over IP ..... 14
SAN Extension and Data Replication ..... 15
Metropolitan Area: Native Fibre Channel Transport ..... 15
Long-Distance SAN Extension: FCIP ..... 16
Cisco MDS 9000 Family Solutions for SAN Extension ..... 16
**Conclusion** ..... **17**
**For More Information** ..... **17**
**Appendix A: FCoE Step-by-Step Configuration** ..... **17**
Cisco Nexus 5000 Series in Switch Mode ..... 19
Cisco Nexus 5000 Series in NPV Mode: Prior to Cisco NX-OS 4.2(1)N1(1) ..... 21
Cisco Nexus 5000 Series in NPV Mode: Cisco NX-OS 4.2(1)N1(1) or Later ..... 22

## What You Will Learn

The Cisco® MDS 9000 Family and the Cisco Nexus® 5000 Series Switches are uniquely positioned to enable transparent coexistence of FCoE and Fibre Channel in both new and existing SAN environments.

## Overview

Customers have deployed SANs to allow multiple servers to access storage over a dedicated network. Fibre Channel is the technology of choice for SANs because it provides high-bandwidth (1/2/4/8-Gbps), no-drop, highly available, resilient capabilities for mission-critical environments. Though traditionally the Fibre Channel SAN and the Ethernet LAN have been deployed on separate physical infrastructures, 10 Gigabit Ethernet has allowed the combination of LAN and SAN traffic on the same physical infrastructure. Fibre Channel over Ethernet (FCoE) allows Fibre Channel to be transported over Ethernet. Unified I/O combines LAN and FCoE traffic on the same link to bring together the ubiquitous and low-cost advantages of Ethernet and the no-drop, highly available, resilient nature of Fibre Channel.

FCoE has steadily gained industry acceptance since the introduction of the Cisco Nexus 5000 Series Switches in mid 2008, with FCoE being approved as a standard by ANSI INCITS T11 on June 3, 2009. A large ecosystem of server vendors, host bus adapter (HBA) and network interface card (NIC) vendors, Fibre Channel and Ethernet switch vendors, and storage array vendors have made a commitment to FCoE and have been actively working to bring FCoE products to the marketplace.

Cisco NX-OS Software is the foundation for both the Cisco MDS 9000 Family and the Cisco Nexus 5000 Series of switches. A single management tool provides visibility and control over storage networking across all Cisco SANs and unified fabrics. Running a common OS managed through a single pane, the Cisco MDS 9000 Family and Cisco Nexus 5000 Series enable the transparent coexistence of FCoE and Fibre Channel in both new and existing SAN environments. In virtualized data centers, virtual machines move transparently between servers attached to FCoE ports on a Cisco Nexus 5000 Series Switch and servers attached to Fibre Channel ports on a Cisco MDS 9000 Family switch. SAN administrators can configure services such as name servers, zone servers, security, authentication, PortChannels, statistics, and quality of service (QoS) across the Cisco Nexus 5000 Series and the Cisco MDS 9000 Family.

Customers can connect to the SAN environment using FCoE or Fibre Channel and enable storage services such as Cisco MDS 9000 Data Mobility Manager (DMM), Cisco MDS 9000 Storage Media Encryption (SME), and Cisco MDS 9000 I/O Accelerator (IOA).

## Principles of SAN Design

The introduction of FCoE does not affect the overall principles and best practices that have been established during more than a decade of Fibre Channel–based SAN design.

Most of the familiar Fibre Channel concepts, such as Fibre Channel name services, zoning, and routing, are integral parts of an FCoE design, although some low-level mechanisms has been replaced by FCoE or Ethernet features that offer equivalent functions.

This section discusses the basic SAN design principles and best practices in the specific context of the deployment of a unified I/O fabric composed of Cisco MDS 9000 Family switches and directors and Cisco Nexus 5000 Series Switches.

### High Availability and Dual (A+B) Fabric

The application level, for instance, a file system or a database, assumes continuous and reliable access to storage. The Small Computer System Interface (SCSI) driver provides a minimum level of error recovery, reporting an operation error after a long timeout, but the error will trigger a complex application-level recovery and dramatically affect performance.

To help ensure high availability, SAN designers commonly provide each application server with at least two HBAs or converged network adaptors (CNAs), with each adaptor connected to one parallel, fully independent SAN. As a consequence, the application server can access, using a multipathing driver, a dual-port active-passive or active-active storage device over two independent paths.

This approach does not depend on a transport technology of either Fibre Channel or FCoE, but on the delivery of traffic through two independent physical fabrics (Figure 1).

**Figure 1.** Dual Fabric in a Unified I/O Deployment

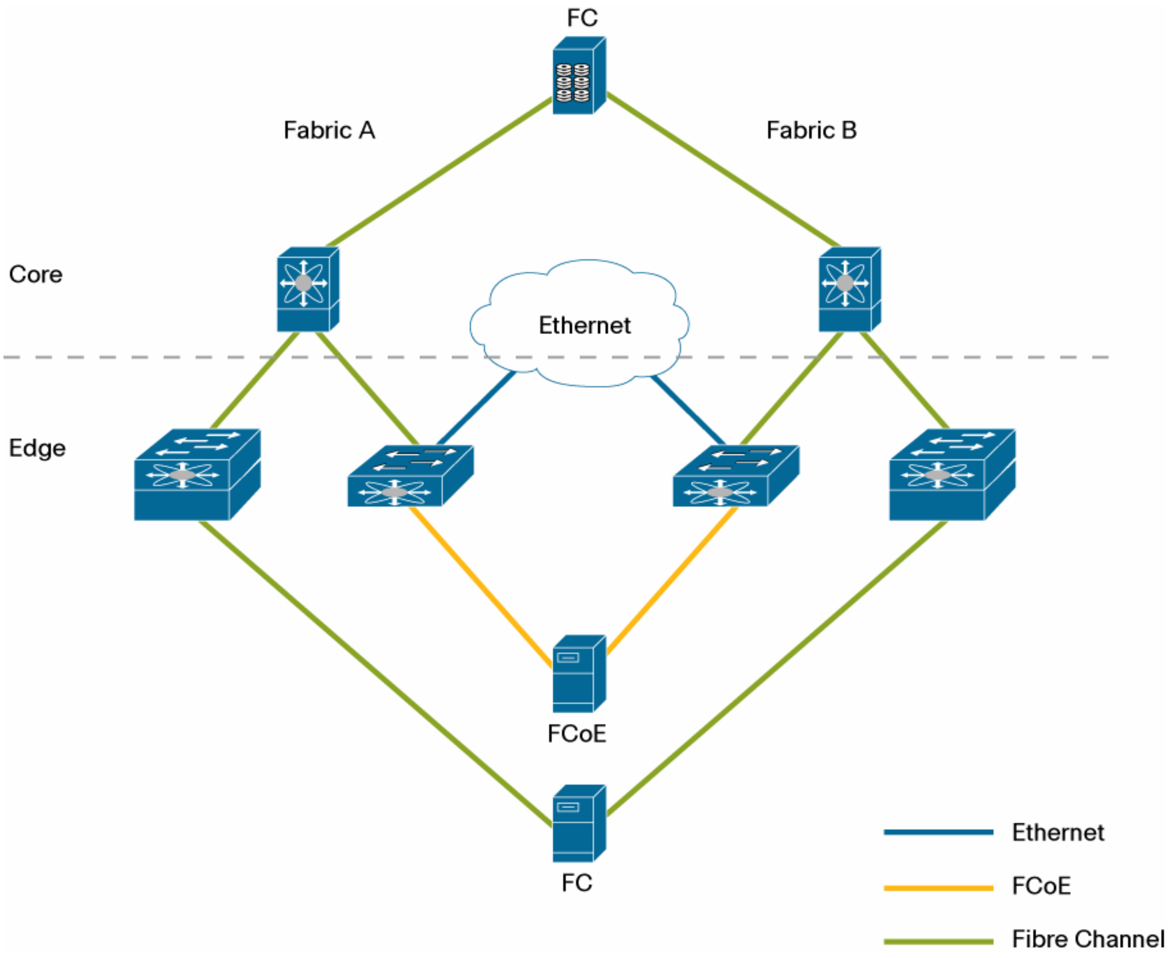

## Oversubscription and Fan-out Considerations

The oversubscription ratio and storage-to-server fan-out configuration considerations for a traditional Fibre Channel network design are the same for an FCoE network design, though CNA settings may influence the actual bandwidth that the 10 Gigabit Ethernet interface can dedicate to FCoE traffic rather than to Ethernet traffic. For example, a 10 Gigabit Ethernet CNA can generate up to 4 Gbps of storage traffic or up to the full 10 Gbps, depending on the CNA internal architecture or the bandwidth allocation settings; these considerations must be taken into account to determine the effective storage bandwidth and the consequent actual oversubscription and fan-out estimates.

Keep in mind that Fibre Channel bandwidth is computed at the physical layer, but Ethernet bandwidth is computed at the data-link layer. As a consequence, 4-Gbps Fibre Channel corresponds to 3.2 Gbps of Ethernet, and 8-Gbps

Fibre Channel corresponds to 6.4 Gbps of Ethernet; according to this convention, a 10 Gigabit Ethernet link could carry up to 12.5-Gbps Fibre Channel.

### Topologies of Individual Fabrics

Multiple design choices are available for multiswitch fabrics, but the most common are the core-edge and the edge-core-edge designs (Figure 2). If the number of Fibre Channel ports is not too high, you can use a collapsed core-edge design, in which a single chassis accommodates different types of line cards, with some serving as core line cards and others as edge line cards.

- **Core-edge topology:** In this design, storage is always put in the core, and hosts are always attached at the edge. This design is effective because SAN traffic flows are typically not peer to peer but instead many to one (hosts to storage).
- **Edge-core-edge topology:** This common design (storage edge to core to host edge) is used when a core-edge design provides insufficient scalability and an additional edge tier is needed to handle the large number of devices.

**Note:** Edge switches can be set to operate in NPV mode as described later in this document.

**Figure 2.** Core-Edge and Edge-Core-Edge Topologies

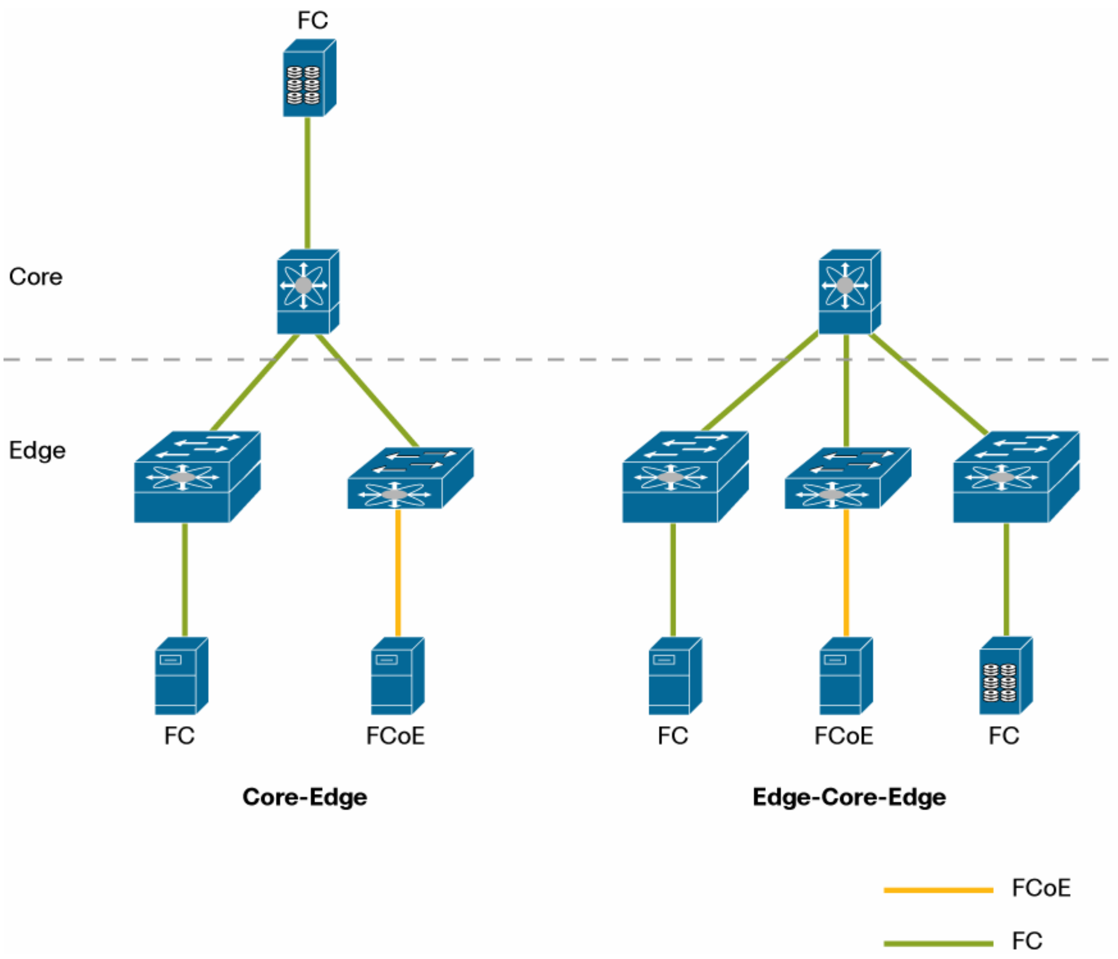

## Cisco MDS 9000 Family Optimization of Cisco Nexus 5000 Series SAN Connectivity

The operating system homogeneity across the Cisco MDS 9000 Family and the Cisco Nexus 5000 Series conserves and extends the capabilities that are specific of a SAN based on the Cisco MDS 9000 Family.

### Fabric Virtualization: VSANs

Virtual SANs (VSANs) are a Cisco innovation, now part of the ANSI INCITS T11 Fibre Channel standards under “Virtual Fabrics.” VSANs provide greater security and stability in Fibre Channel and FCoE fabrics by helping ensure isolation among devices that are physically connected to the same consolidated network infrastructure. VSAN trunking enables frames to be transmitted and received in more than one VSAN over the same physical link by using the Enhanced Inter-Switch Link (EISL) frame format. VSAN trunking is supported on native Fibre Channel interfaces on both the Cisco MDS 9000 Family and Cisco Nexus 5000 Series, but not on virtual Fibre Channel interfaces configured on the 10 Gigabit Ethernet interfaces that provide FCoE access in the Cisco Nexus 5000 Series. By default, the VSAN trunking protocol is enabled. If the trunking protocol is disabled, a port will transport only the traffic belonging to the VSAN that has been statically assigned to that port (Figure 3).

**Figure 3.** Trunking Expansion Port in Switch Mode

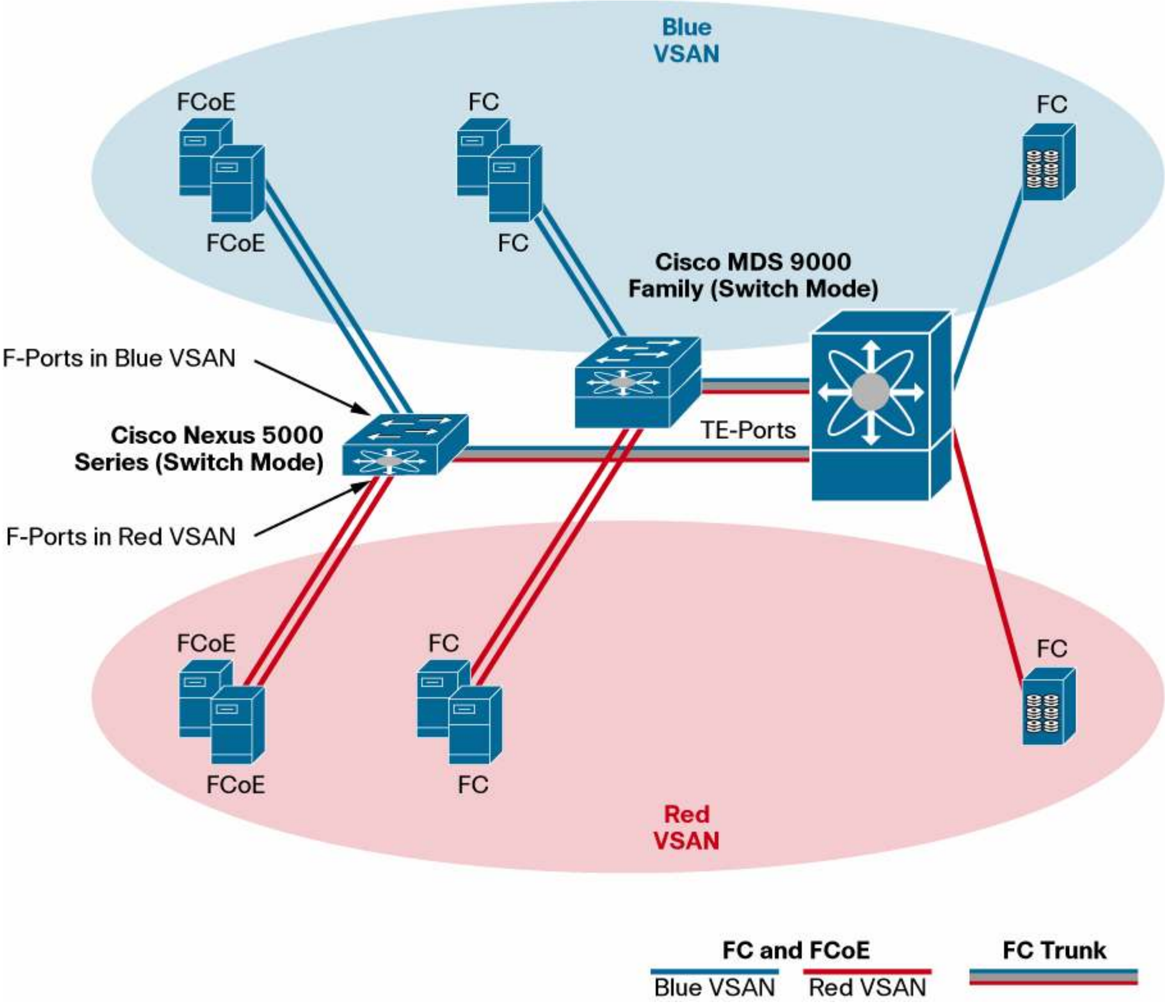

### Network Port Virtualization Mode

The Cisco Nexus 5000 Series network port virtualization (NPV) implementation is designed to integrate the Cisco Nexus 5000 Series access switch with a Cisco MDS 9000 Family core switch supporting multiple VSANs to extend the benefits of the VSAN infrastructure to the end devices connected through FCoE interfaces.

#### Network Edge in NPV Mode

To optimize scalability, streamline management, and simplify maintenance, the servers can optionally be connected to the core Cisco MDS 9000 Family fabric through a Cisco Nexus 5000 Series access switch in NPV mode.

This mode, compliant with the relevant ANSI INCITS T11 Fibre Channel standards, requires the core Cisco MDS 9000 Family switch to be configured to support N-Port ID Virtualization (NPIV).

In switch mode, each switch that joins a SAN is assigned a domain ID. Each ANSI INCITS T11 standards-compliant SAN (or VSAN) supports a maximum of 239 domain IDs, so the SAN has a limit of 239 switches. In a SAN topology with a large number of edge switches, the SAN may need to grow beyond this limit, or beyond the much lower limit (for example, 50) supported by Original Storage Manufacturers (OSMs).

NPV alleviates the domain ID limit problem by sharing the domain ID of the core switch among multiple edge switches. The Cisco Nexus 5000 Series access switch emulates an HBA node port (N-port) without being identified as a switch part of the core fabric. The Cisco MDS 9000 Family core switch or director is connected to each Cisco Nexus 5000 Series Switch, and ultimately to each server, using standard fabric ports (F-ports) instead of expansion ports (E-ports). In this configuration, the Cisco MDS 9000 Family core switch provides critical F-port functions, such as login and port security, and all the Fibre Channel switching capabilities (Figure 4).

**Figure 4.** Edge Switch in NPV Mode with, Node Proxy Port (NP-Port)

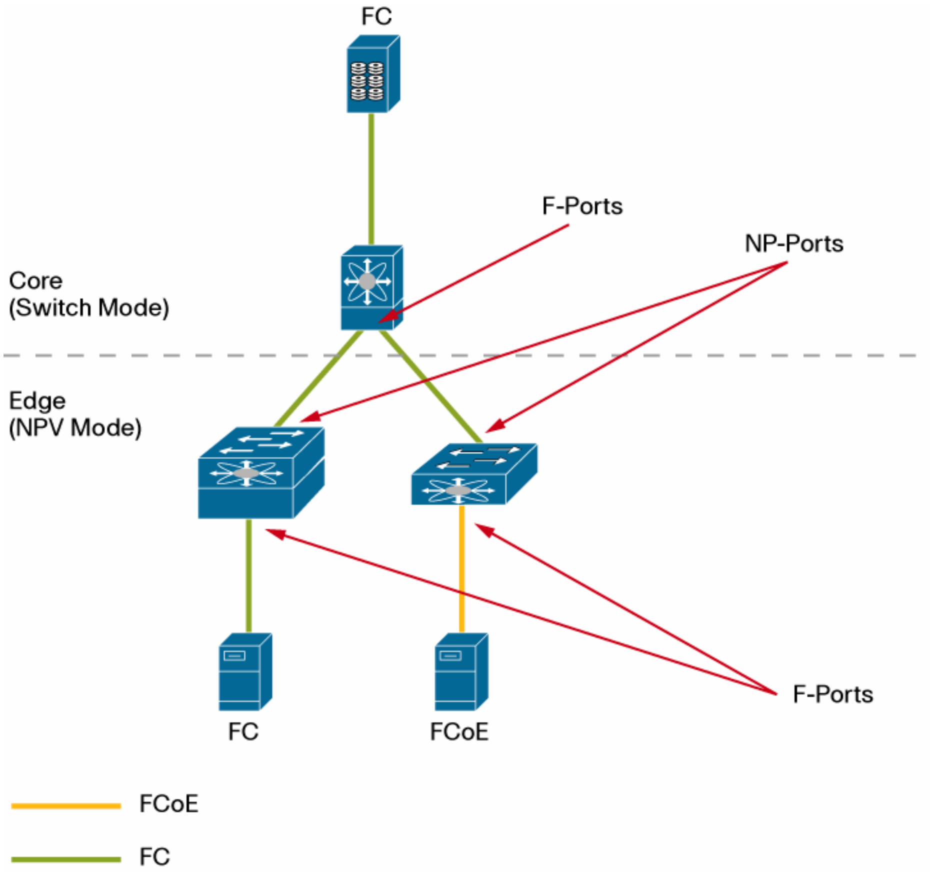

Starting from Cisco NX-OS Software Release 4.2(1)N1(1), Cisco Nexus 5000 Series Switches provide the F-port trunking and channeling modes. For the Cisco MDS 9000 Family, this is a consolidated technology. Connecting the Cisco Nexus 5000 Series Switch to a Cisco MDS 9000 Family switch using these features enables efficient sharing of physical links in between multiple VSANs, and dynamic load sharing across multiple physical links between the NPV-mode edge switch and the core switch. Please refer to Figure 8 for a detailed description of Port Channeling and Trunking with the Nexus 5000 in NPV and Switch mode, highlighting common advantages and differences.

For versions earlier than Cisco NX-OS 4.2(1)N1(1), the Cisco Nexus 5000 Series does not support PortChannels and trunking to the Cisco MDS 9000 Family. Each Cisco Nexus 5000 Series Fibre Channel uplink in NPV mode is connected to a core port on the Cisco MDS 9000 Family core switch, and the uplink belongs to the VSAN of the core port. Each FCoE initiator connected to a Cisco Nexus 5000 Series virtual interface is assigned to a specific VSAN. The traffic from the given initiator is automatically switched by the Cisco Nexus 5000 Series Switch to the NP-port uplink belonging to the same VSAN and delivered to the core Cisco MDS 9000 Family F-port in the same VSAN (Figure 5). Note that in NPV mode, the maximum number of VSANs is equal to the number of native Fibre Channel interfaces on the Cisco Nexus 5000 Series Switch connected to the Cisco MDS 9000 Family core.

**Figure 5.** Automatic VSAN Selection on the Uplink

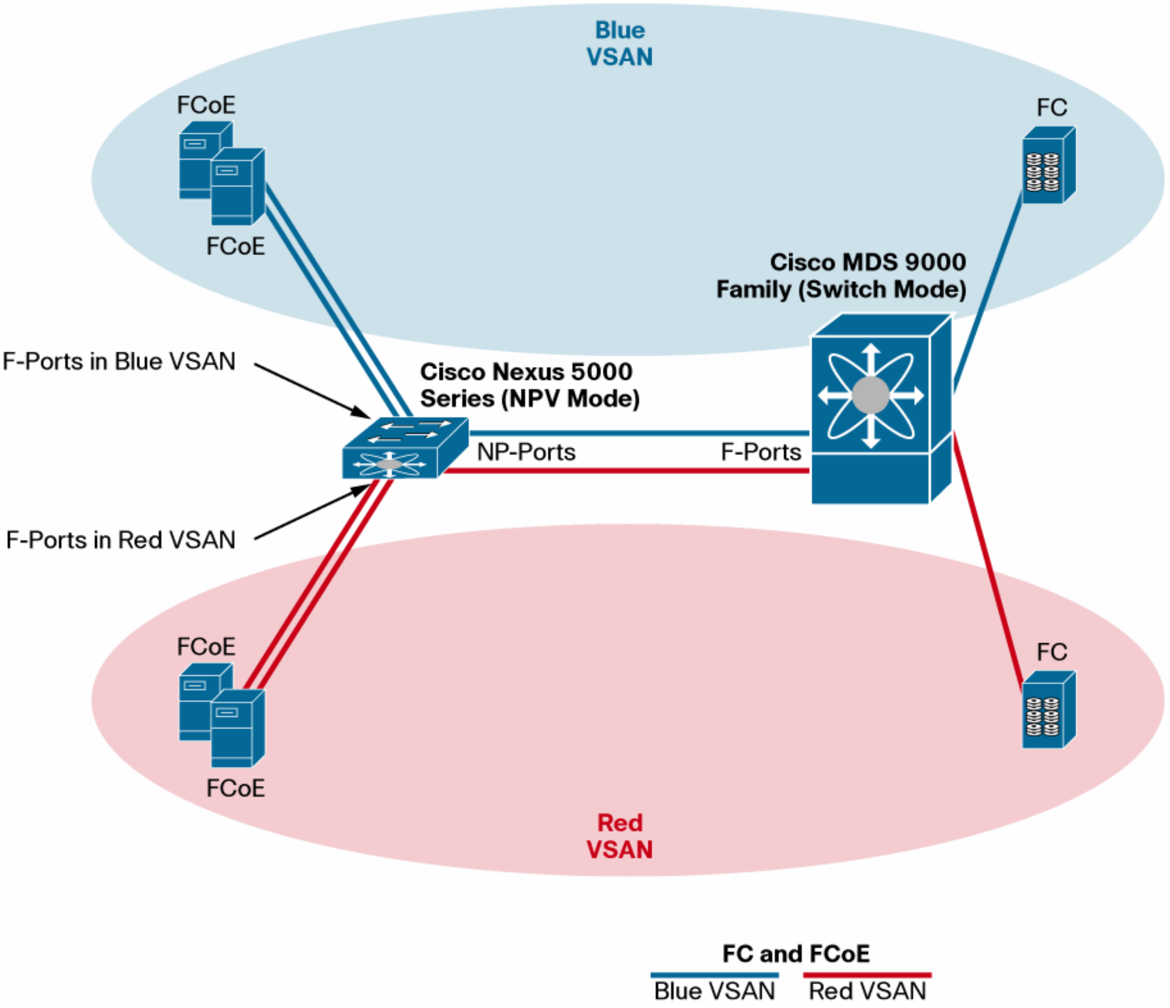

## NPV Scalability Considerations

When connected to multiple Cisco Nexus 5000 Series Switches in NPV mode, the Cisco MDS 9000 Family core switch directly processes the login of each individual server (or virtual machine, if the server implements NPIV). Notice that the Cisco MDS 9000 Family supports a maximum number of logins for the interface, the line card, and the switch, and this number must not be exceeded when aggregating a large number of servers with NPV (Figure 6).

**Figure 6.** Fabric Scalability Analysis in NPV Mode

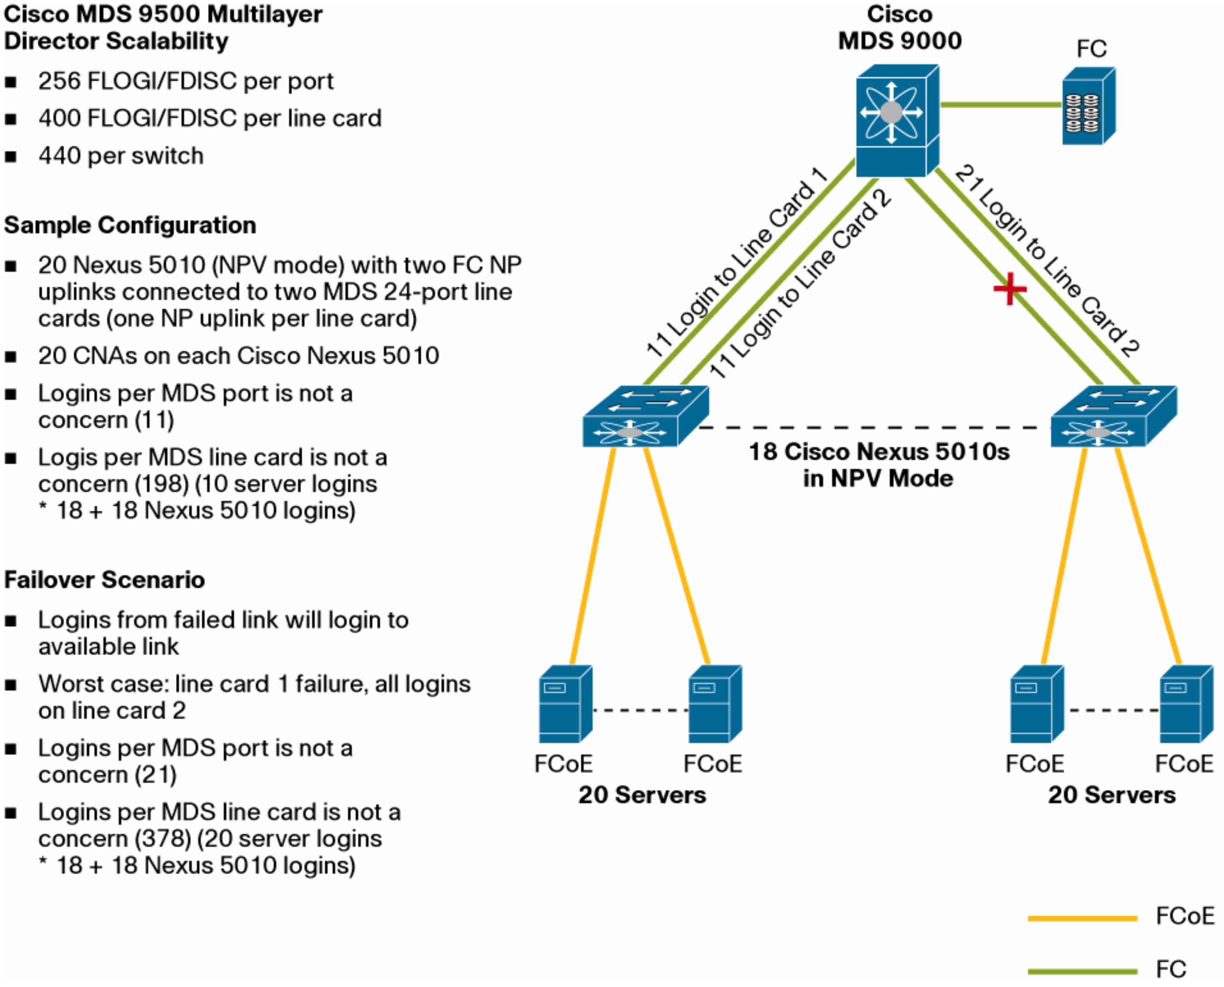

**Cisco MDS 9500 Multilayer Director Scalability**

- 256 FLOGI/FDISC per port
- 400 FLOGI/FDISC per line card
- 440 per switch

**Sample Configuration**

- 20 Nexus 5010 (NPV mode) with two FC NP uplinks connected to two MDS 24-port line cards (one NP uplink per line card)
- 20 CNAs on each Cisco Nexus 5010
- Logins per MDS port is not a concern (11)
- Logis per MDS line card is not a concern (198) (10 server logins * 18 + 18 Nexus 5010 logins)

**Failover Scenario**

- Logins from failed link will login to available link
- Worst case: line card 1 failure, all logins on line card 2
- Logins per MDS port is not a concern (21)
- Logins per MDS line card is not a concern (378) (20 server logins * 18 + 18 Nexus 5010 logins)

As shown in Figure 6, if multiple connections exist between the Cisco Nexus 5000 Series NPV switch and the Cisco MDS 9000 Family core switch, the login limits must not be exceeded even if one or more link fails and the servers log in again on the remaining links.

### Sharing Devices Across VSANs

The Cisco Inter-VSAN Routing (IVR) feature can establish tightly controlled communication between devices in different VSANs, without the need to merge the individual VSAN management scope and fault domains.

Cisco MDS 9000 IVR provides the outstanding design flexibility, allowing sharing of consolidated resources such as a tape library or an archive of boot images. Cisco MDS 9000 IVR supports Cisco Nexus 5000 Series edge switches configured in both switch mode and NPIV mode.

Cisco MDS 9000 IVR complements Cisco data center interconnect (DCI) technologies and is an essential component for SAN extension.

An application example is the virtual machine mobility over distance, using VMware VMotion. Figure 7 describes one of the possible configurations: the mobility is based on storage shared via IVR, and while the VSAN infrastructure helps ensure the desired level of data center isolation.

**Figure 7.** IVR: Enables Long-Distance Virtual Machine Mobility

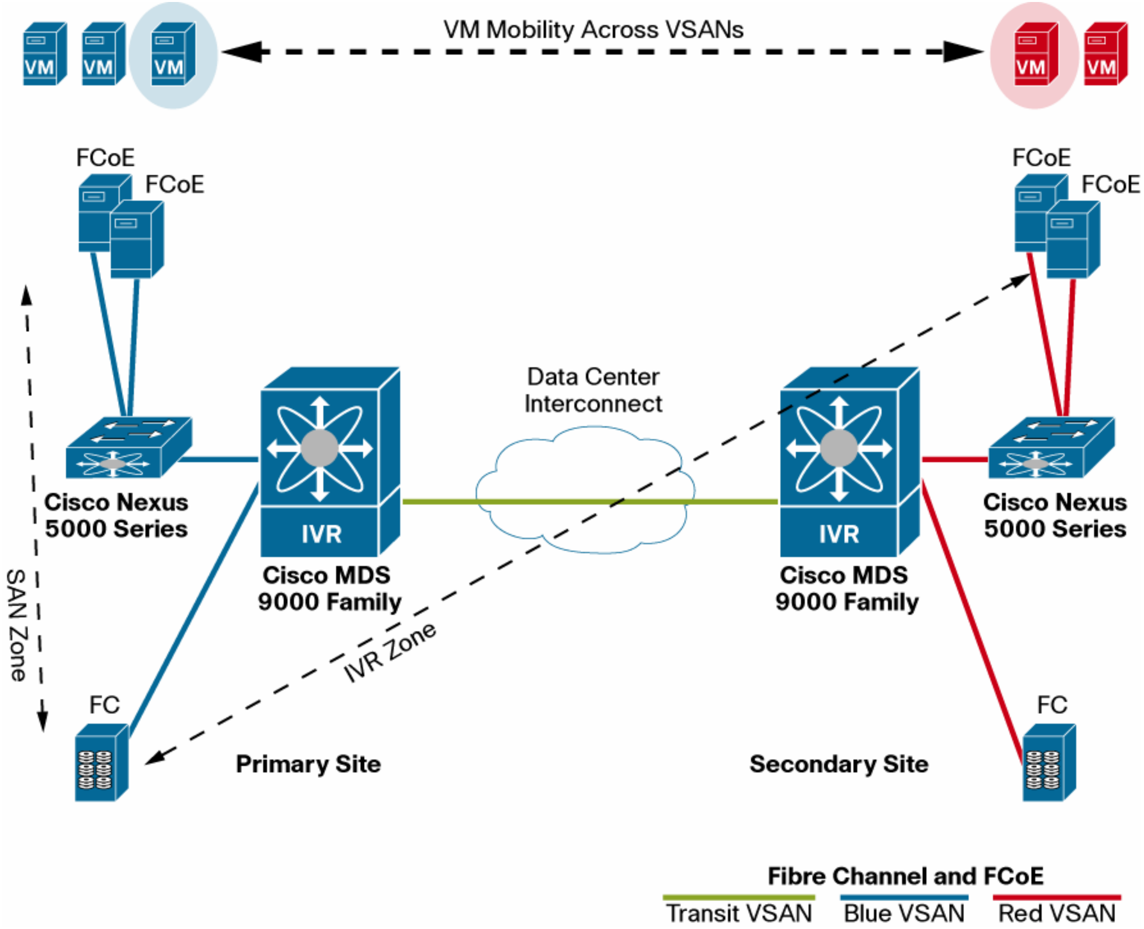

### Network-Level High Availability: PortChannel

One means of delivering high availability at the network level is the aggregation of multiple physical Inter-Switch Links (ISLs) into a single logical interface. This aggregation allows fabric administrators to provide link redundancy, greater aggregated bandwidth, and load balancing. Cisco calls this technology SAN PortChannel (Figure 8). An important advantage of Cisco SAN PortChannel technology is that the bundled physical links can be located on any port on any module in the switch, protecting against both link failures, such as cable breaks and faulty optics, and switching module failures. The Cisco SAN PortChannel solution does not show any performance degradation over long distances and has no specific cabling requirements. The Cisco SAN PortChannel provides load balancing and can deliver predictable and robust performance independent of distance. Load balancing can be:

- Flow based (all frames between source and destination follow the same links for a given flow)
- Exchange based (the first frame in an exchange is assigned to a link, and then subsequent frames in the exchange follow the same link); this method provides finer granularity for load balancing while preserving the order of frames for each exchange

Before configuring a SAN PortChannel, consider the following guidelines:

- Configure the SAN PortChannel using Fibre Channel ports from multiple line cards in the Cisco MDS 9000 Family switch and from both expansion modules in the Cisco Nexus 5000 Series Switch, to provide increased availability (in case one of the modules fails).
- Make sure that all ports in a SAN PortChannel are connected to the same sets of switches. SAN PortChannels require point-to-point connections between the same set of switches.

Cisco NX-OS can detect configuration errors across the Cisco MDS 9000 Family and Cisco Nexus 5000 Series Switches, and it responds by shutting down the links and generating a misconfiguration message.

**Figure 8.** SAN Port Channel and Trunking

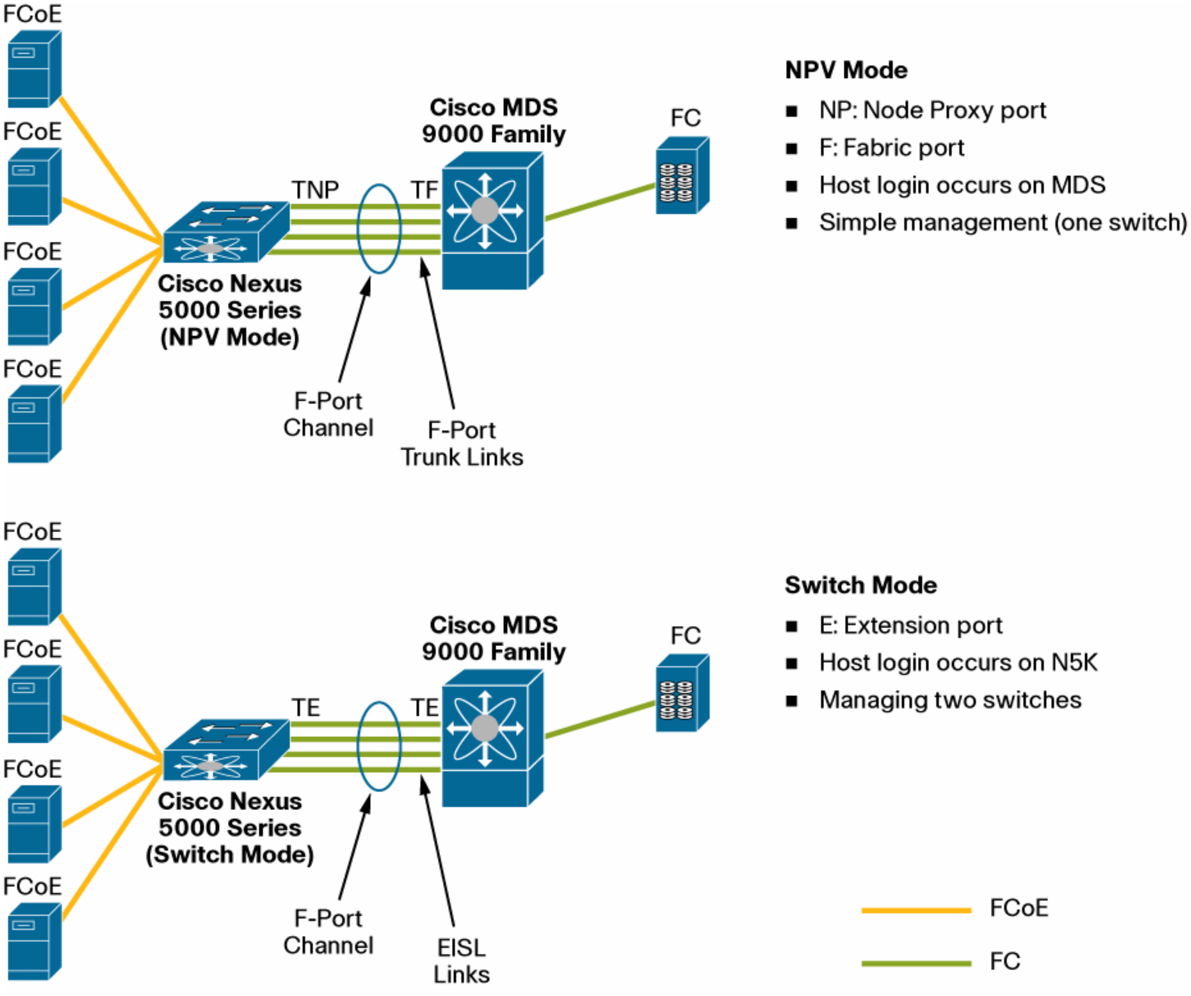

**NPV Mode**

- NP: Node Proxy port
- F: Fabric port
- Host login occurs on MDS
- Simple management (one switch)

**Switch Mode**

- E: Extension port
- Host login occurs on N5K
- Managing two switches

**Common Advantages**

- Trunking: Each link can carry traffic for multiple VSANs (Tx port)
- Aggregation: Dynamic load sharing across the links in a channel (flow-based or exchange-based)
- Use ports from both N5K Expansion modules and from multiple MDS modules (increased availability)
- Easier management (multiple links seen as one link only)

### Transparent Fabric Services: Zoning and Aliases Services

Since the Cisco MDS 9000 Family and the Cisco Nexus 5000 Series are powered by the same Cisco NX-OS unified data center operating system, fabric services such as enhanced zoning and aliases services can be used to enhance interoperability.

In a fabric, all switches must be capable of supporting enhanced zoning to enable the feature, but this consideration is applicable to a Cisco Nexus 5000 Series Switch in switch mode. Interoperability or zoning enforcement is generally not a concern when the Cisco Nexus 5000 Series Switch operates in NPV mode (fabric service functions are completely dependent on the core switch). When a third-party core switch is used, zones may be limited to the World Wide Port Name (WWPN), and other proprietary zoning methods (for example, physical port number) may be eliminated. Enhanced zoning provides considerable advantages, such as reduced zone database size, administration consistency, and protection against undesired merging. Consult the documentation for more information.

### Distributed Device Alias Service

In addition to zone-based aliases, the Cisco MDS 9000 Family and the Cisco Nexus 5000 Series Switches both support distributed device alias services (device aliases) on a fabricwide basis. Device aliases simplify management tasks when the WWPN of a device must be specified to configure applications such as zoning, dynamic port VSAN membership (DPVM), port security, and troubleshooting commands.

When operating in enhanced mode, the features accept a device alias in its native format. In enhanced mode, the device alias is stored in the configuration and distributed in its native device alias format, so applications such as zoning can automatically keep track of the device alias membership changes and enforce them accordingly, providing a single point of change.

Enhanced mode, or native device alias configurations, are not accepted in VSANs in interoperability mode.

### Homogeneous Authentication, Authorization, and Accounting

Use of Cisco NX-OS across the Cisco MDS 9000 Family and the Cisco Nexus 5000 Series conserves the authorization, authentication, and accounting (AAA) infrastructure. User accounts, user groups, and device identities can be consolidated and stored on a standard RADIUS or TACACS+ device, and the syslog can be stored on a common syslog server.

Prior to unified I/O, the LAN and SAN network elements were managed by separate LAN and SAN administrative teams. To retain this organizational structure in a unified network, a way is needed to separate LAN and SAN administrative roles as well as help ensure their independence from one another. This is accomplished by using role-based access control (RBAC): a user account can be associated in the AAA server with one or more roles, for instance, by setting multiple attribute-value pairs in a RADIUS server. User roles are defined by rules that specify the access permissions allowed for each person assigned to that role. Each user role can have multiple rules, and each user can belong to more than one role. The Cisco MDS 9000 Family and the Cisco Nexus 5000 Series Switches provide the following default user roles:

- Network-admin (super user): This role has complete read and write access to the entire device.
- Network-operator: This role has read access plus access to some basic troubleshooting commands such as ping.

Cisco recommends defining the following three roles (more specific roles may be needed):

- **Unified-admin:** This role includes all actions that affect both LAN and SAN operations. This role is specific of the Cisco Nexus 5000 Series.
- **LAN-admin:** This role includes a set of actions that affect LAN operation while denying any actions that could affect SAN operations. This role is specific of the Cisco Nexus 5000 Series.
- **SAN-admin:** This role is associated with a different set of rules when the user logs in on a Cisco MDS 9000 Family switch or a Cisco Nexus 5000 Series Switch. If the user logs in on a Cisco Nexus 5000 Series Switch, the role includes a set of actions that affect SAN operation while denying any actions that could affect LAN operations. If the user logins on a Cisco MDS 9000 Family switch, since SAN administration in a unified environment is not different from administration performed in a separate SAN today, this role may be equivalent to the network-admin role.

### Unified Fabric Access Security

The Fibre Channel Security Protocol (FC-SP) provides switch-to-switch and host-to-switch authentication to overcome security challenges for enterprisewide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DH-CHAP) is an FC-SP protocol that provides authentication between Cisco MDS 9000 Family switches and Cisco Nexus 5000 Series Switches and other devices.

The Cisco MDS 9000 Family and Cisco Nexus 5000 Series Switches enforce a strong security policy for the FC-SP credentials used in a fabric. All FC-SP secrets must contain numbers and case-sensitive letters only and be between 8 and 64 characters in length. The Cisco MDS 9000 Family and Cisco Nexus 5000 Series Switches can perform FC-SP authentication using a common AAA server (Figure 9). The use of a RADIUS or TACACS+ server is recommended for fabrics with more than five switches. For fabrics with fewer than five switches, FC-SP secrets can be managed locally in the fabric without the use of an external AAA server.

**Figure 9.** SAN Security in the Unified Fabric

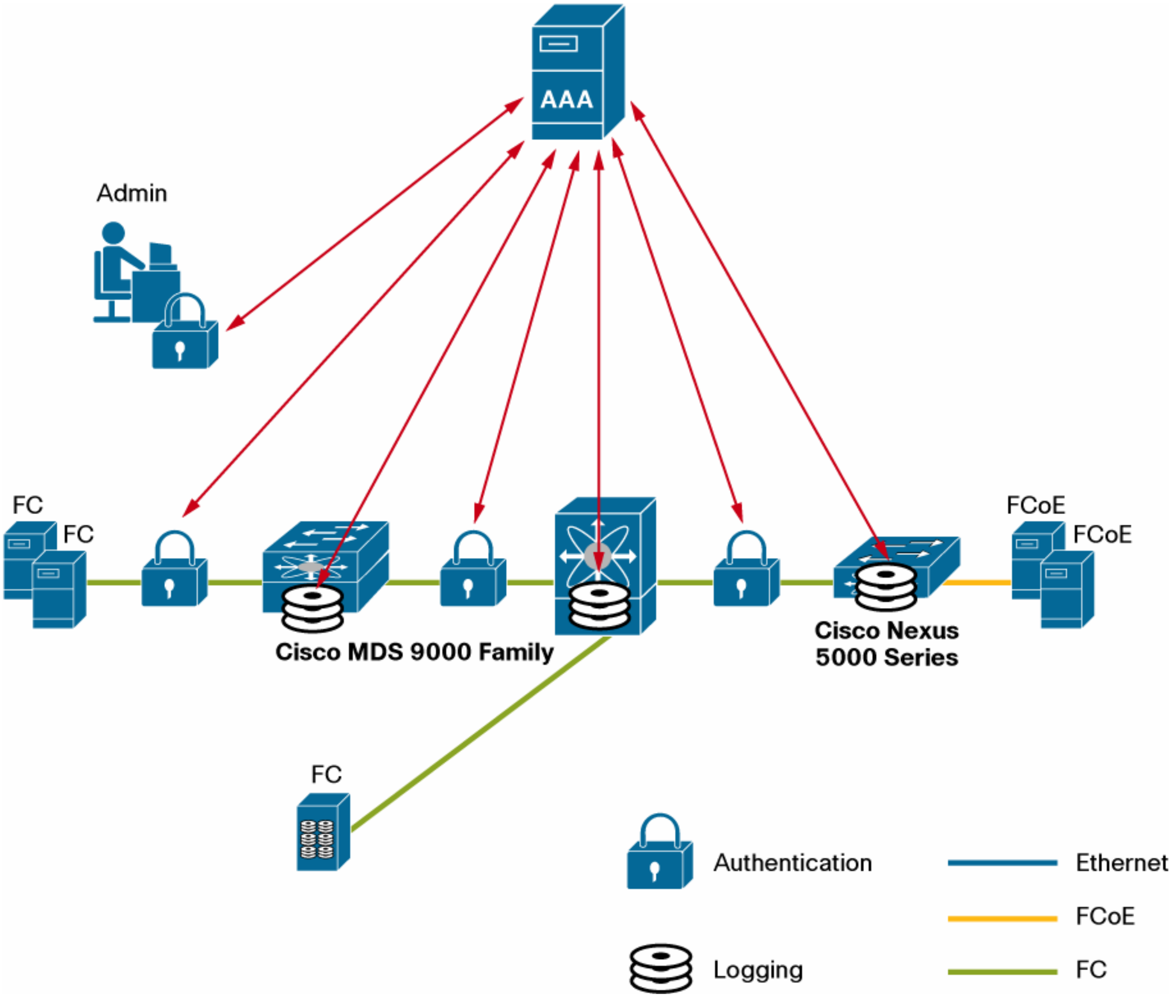

### Configuration Synchronization with Cisco Fabric Services

Some features in the Cisco Nexus 5000 Series and Cisco MDS 9000 Family switches require configuration synchronization with other switches in the network to function correctly, but synchronization through manual configuration at each switch in the network can be a tedious and error-prone process.

Cisco Fabric Services provides a common infrastructure for automatic configuration synchronization in the network. It provides the transport function and a set of services common to the features. Cisco Fabric Services can discover switches in the network that have Cisco Fabric Services capability and can discover feature capabilities in all switches that support Cisco Fabric Services.

The Cisco MDS 9000 Family and the Cisco Nexus 5000 Series both support Cisco Fabric Services message distribution over Fibre Channel, IPv4, and IPv6 networks. If the switch is provisioned with Fibre Channel ports, Cisco Fabric Services over Fibre Channel is enabled by default. Cisco Fabric Services over IP must be enabled explicitly.

### Quality of Service

The Cisco Nexus 5000 Series Switches automatically classify all Fibre Channel and FCoE control and data traffic in the FCoE system class, which provides a selective no-drop service for storage traffic.

The Cisco MDS 9000 Family switches allow a further classification of the storage traffic, with a granularity that can extend up to an individual initiator and target and logical unit number (ITL).

The QoS feature in the Cisco MDS 9000 Family switches prioritizes data traffic in distinct levels of service (low, medium, or high priority) to help ensure that data traffic for latency-sensitive applications receives higher priority than throughput-intensive applications.

**Figure 10.** Storage Traffic QoS in the Unified I/O Fabric

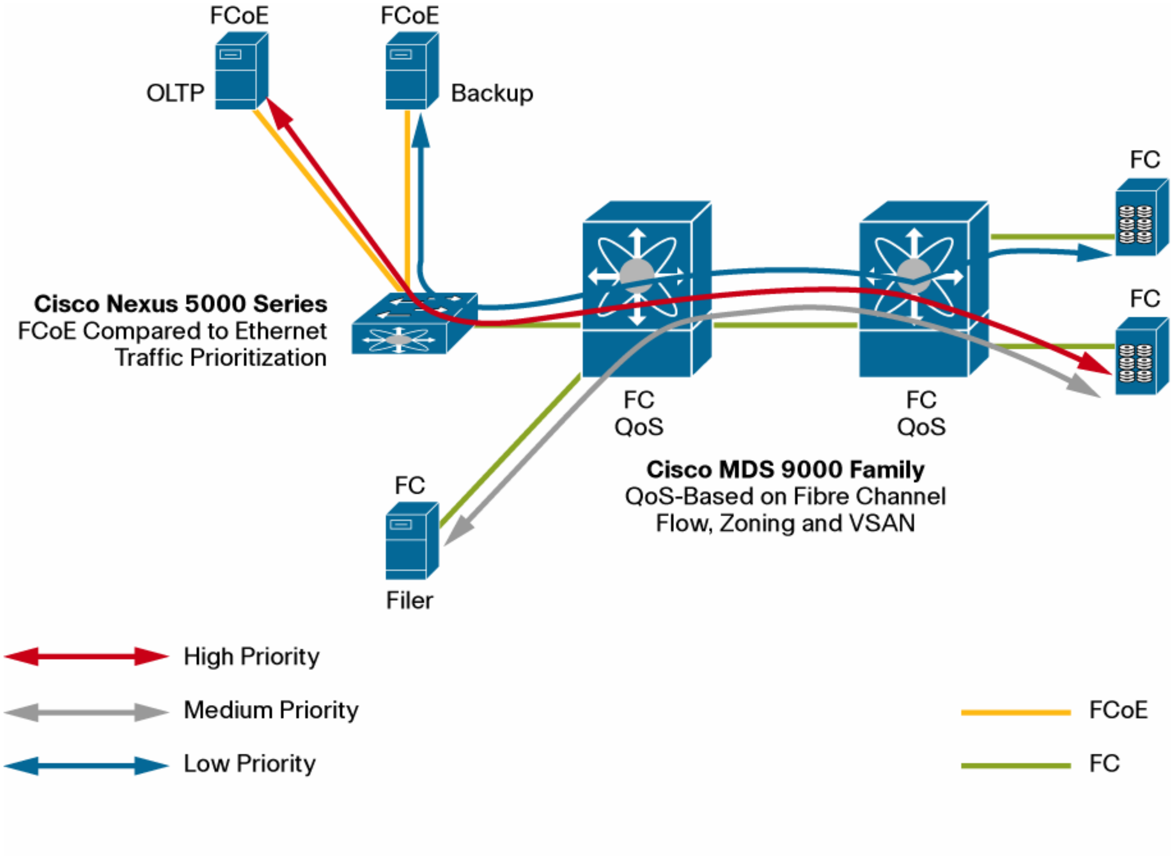

In the typical user scenario shown in Figure 10, online transaction processing (OLTP) traffic is given priority over backup traffic. QoS helps ensure shorter delays and higher bandwidth to high-priority traffic during congestion. If the ISL is congested when the OLTP server sends a frame, the frame is queued in the high-priority queue and is forwarded almost immediately since the high-priority queue is not congested. The scheduler assigns it priority over the backup traffic in the low-priority queue.

### Simplified Operations and Training

The administration of a network composed of devices belonging to the Cisco MDS 9000 Family and Cisco Nexus 5000 Series is greatly simplified, since Cisco NX-OS provides the same command-line interface (CLI) across the Cisco MDS 9000 Family and the Cisco Nexus 5000 Series. The command syntax and usage is the same, wherever applicable, across the two platforms, Cisco Fabric Manager supports the unified fabric, and Cisco Device Manager is invoked through Cisco Fabric Manager and is functionally equivalent for the Cisco MDS 9000 Family and the Cisco Nexus 5000 Series. SAN troubleshooting tools also are the same:

- FCtrace traces the route followed by data traffic and computing interswitch (hop-to-hop) latency. The trace frame is routed normally through the network until it reaches the F-port connected to the end node with the given WWPN or Fibre Channel ID (FC-ID); then the frame is looped back to the originator. If the destination cannot be reached, FCtrace displays path information up to the point of failure.
- FCping verifies the reachability of a node by checking its end-to-end connectivity, based on the node FC-ID, the destination WWPN, or the device alias information.
- Ethanalyzer is a Cisco NX-OS protocol analyzer based on the Wireshark (formerly Ethereal) open source code. Ethanalyzer is a command-line version of Wireshark that captures and decodes packets. Ethanalyzer facilitates troubleshooting by analyzing the control-plane traffic.
- Switched Port Analyzer (SPAN), sometimes called port mirroring or port monitoring, creates a copy of network traffic for analysis by a network analyzer. The network analyzer can be a Cisco SwitchProbe, a Fibre Channel analyzer, or a remote monitoring (RMON) probe.

### Multiprotocol Options: Fibre Channel, iSCSI, and FCoE

A common scenario is a data center that hosts a nonhomogeneous set of servers. In this case, the SAN architect needs advanced and flexible options to provide connectivity to the consolidated enterprise storage, to lower the overall cost of the solution while maintaining the same features and functions.

#### Small Computer System Interface over IP

The Small Computer System Interface over IP (iSCSI) protocol is a popular choice to provide enterprise storage access to low-end servers and, because it can fully support virtual machine mobility, to deploy entry-level server virtualization solutions.

The Cisco MDS 9000 Family offers a fully integrated Fibre Channel–to–iSCSI gateway solution that allows iSCSI clients to access, over a 1-Gbps classical Ethernet connection, the consolidated Fibre Channel storage. The Cisco MDS 9000 Family iSCSI implementation provides all the capabilities available to a Fibre Channel initiator (including VSAN, advanced security, and zoning) to the iSCSI initiator, simplifying migration and enabling hybrid deployments.

The capability to have a combination of Cisco MDS 9000 Family switches, Cisco Nexus 5000 Series Switches, and Cisco Nexus 2000 Series Fabric Extenders supporting Fibre Channel, FCoE, and iSCSI offers greater consolidation while providing a unified management interface (Figure 11).

**Figure 11.** Transparent Integration of Fibre Channel, FCoE, and iSCSI

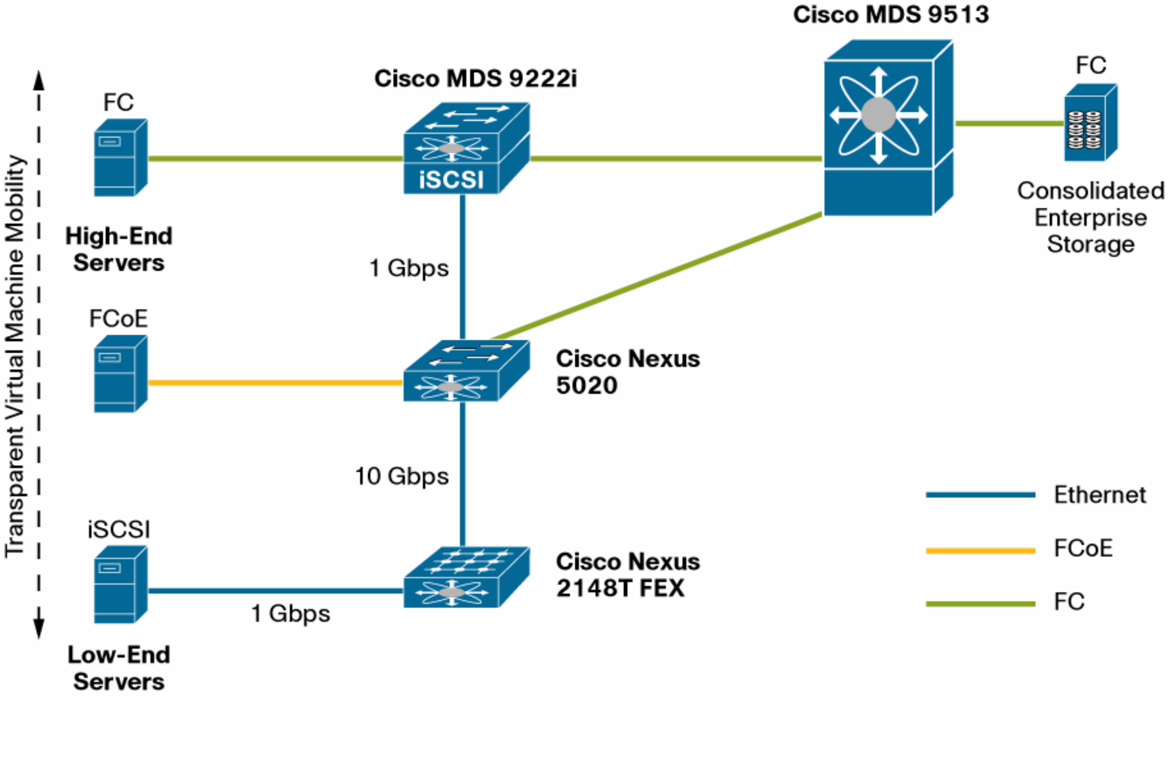

## SAN Extension and Data Replication

Most disaster recovery strategies require a mirror copy of primary storage data in the remote secondary data center. Several products on the market offer synchronous or asynchronous replication: for example, EMC CLARiiON MirrorView, RecoverPoint, Symmetrix Remote Data Facility (SRDF), and NetApp SnapMirror. The interconnection between locations can be achieved with Cisco SAN extension technologies such as native Fibre Channel transport and Fibre Channel over IP (FCIP). The Cisco MDS 9000 I/O Accelerator (IOA) provides a fabric-based, highly integrated, topology-independent acceleration service to improve remote I/O performance and bandwidth utilization over metropolitan-area networks (MANs) and WANs.

### Metropolitan Area: Native Fibre Channel Transport

An organization with direct access to dark fiber has several options for extending SAN connectivity:

- Fibre Channel can be transmitted directly over the dark fiber infrastructure.
- A coarse wavelength-division multiplexing (CWDM) connection allows up to eight wavelengths to share the same fiber, increasing the available bandwidth, and can host mixed SAN and Ethernet traffic.
- Dense wavelength-division multiplexing (DWDM) uses different optical frequencies to allow more than 32 channels of communication to share a dark fiber infrastructure (more than 320-Gbps aggregated bandwidth).

The approaches based on optical technologies rely on native flow control techniques to achieve a lossless and performance-optimized interconnection.

The Cisco Nexus 5000 Series software currently supports cable lengths up to 300 meters. However, if a Cisco Nexus 5000 Series Switch is connected to a CNA, then the maximum distance may be lower: for instance, 50 meters for a Menlo-based first -generation CNA. Because of this limitation, the current FCoE implementation is unsuitable for deployment as a SAN extension solution.

The Fibre Channel interfaces of the Cisco Nexus 5000 Series Switches allow up to 64 buffers, enabling connectivity up to 32 kilometers at 4 Gbps; for longer distances, Cisco Nexus 5000 Series Switches can be connected to a core Cisco MDS 9000 Family switch to take advantage of the larger number of buffers.

### Long-Distance SAN Extension: FCIP

For longer distances and additional flexibility, FCIP communications travel across TCP/IP networks, supporting almost unlimited distance between the source and target of a replication configuration. With any long-distance communication, the farther apart the endpoints in the network, the greater the potential for network latency to affect performance. Nonetheless, FCIP has been successfully deployed for distances exceeding 10,000 miles.

### Cisco MDS 9000 Family Solutions for SAN Extension

The Cisco MDS 9000 Family of fabric switches and directors provides native Fibre Channel interfaces with a large number of buffer credits for long-distance connection and integrates multiprotocol support for FCIP, performing local fabric switching and SAN extension services on a single platform (Figure 12).

**Figure 12.** Flexible SAN Extension Solutions

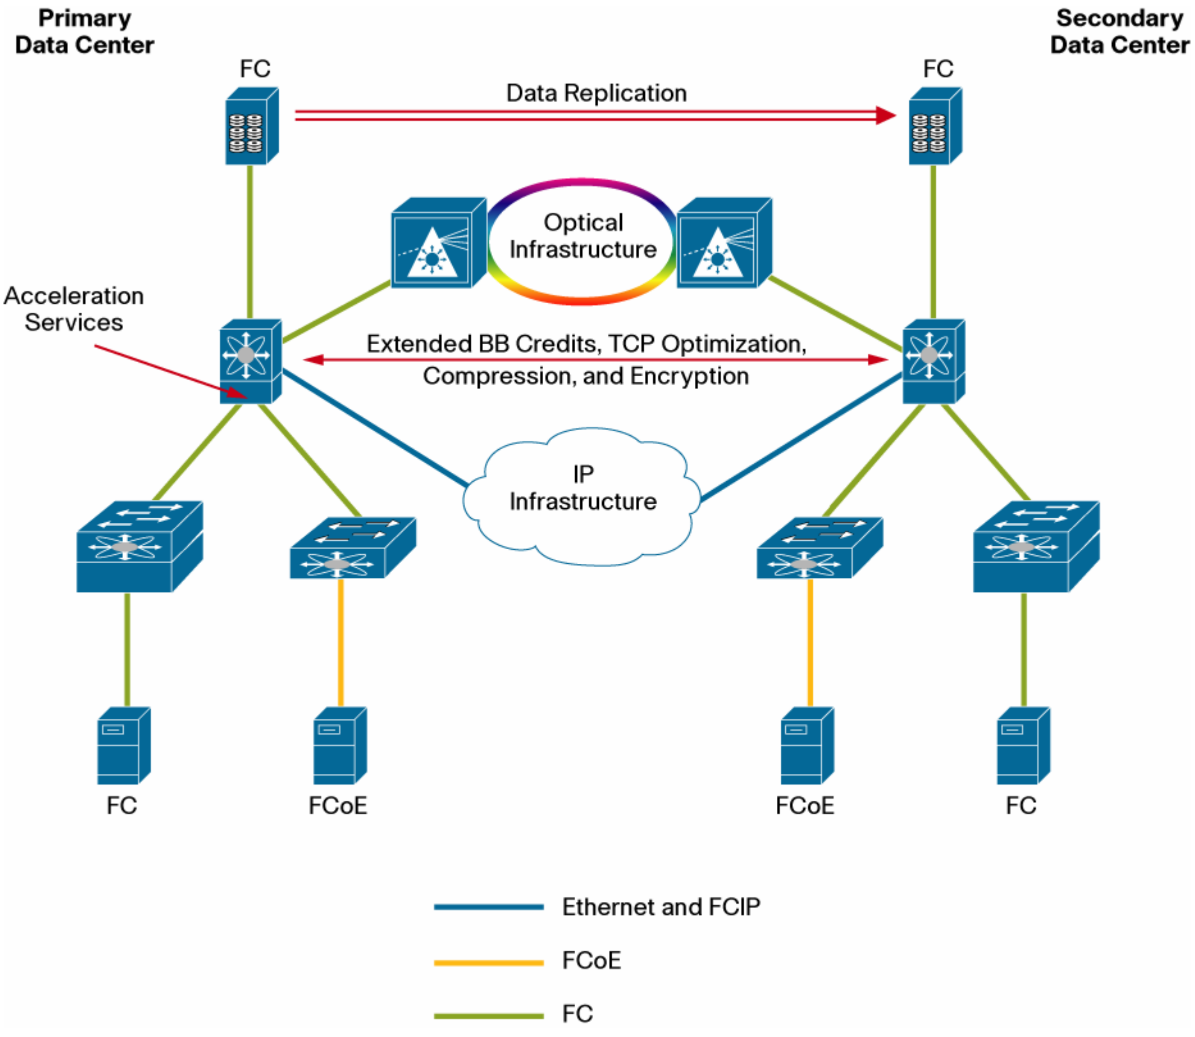

Support for network QoS and VSANs allows storage replication traffic to be segregated effectively and independently managed over a shared network infrastructure (also see Figure 8 and Figure 9). The availability of the Cisco MDS 9000 Family SAN extension links is improved by aggregating links in PortChannels. IOA services significantly improve replication performance and overall network throughput when traffic travels over extended distances. FCIP compression optimizes bandwidth, allowing data replication across low- and intermediate-bandwidth long-distance networks. Cisco MDS 9000 Family hardware helps ensure data confidentiality and integrity, providing native Fibre Channel encryption and native IP Security (IPsec) support for FCIP. Cisco NX-OS includes tools for monitoring and optimizing the performance of a FCIP connection.

## Conclusion

The Cisco MDS 9000 Family and the Cisco Nexus 5000 Series share the same unified data center operating system, Cisco NX-OS, based on industry-tested Cisco IOS® Software and Cisco MDS 9000 SAN-OS Software. Cisco NX-OS is a data center–class operating system built to meet the demands of the virtualized data center, including features such as a modularity, flexible architecture, and continuous system availability.

The operating system homogeneity across the Cisco MDS 9000 Family and the Cisco Nexus 5000 Series brings the advanced capabilities of the Cisco MDS 9000 Family SAN to servers connected over FCoE. Strong synergies are found in the areas of SAN virtualization (VSAN), network-level availability (SAN PortChannels), zoning and fabric services, security, management, and fabricwide configuration support.

The Cisco MDS 9200 Series Multilayer Switches and MDS 9500 Series Multilayer Directors provide the unified fabric with multiple storage services and multiprotocol capabilities, readily available to the Fibre Channel initiators as well as to the FCoE initiators. SAN extension over IP, I/O acceleration services, data replication across heterogeneous storage arrays, and storage media encryption are examples of the advanced features that benefit the FCoE initiators connected through a Cisco Nexus 5000 Series Switch to the Cisco MDS 9000 Family core.

Management and operation is simplified by GUI-based Cisco Fabric Manager and Device Manager support for both the Cisco MDS 9000 Family and the Cisco Nexus 5000 Series.

In addition, the Fibre Channel troubleshooting tools available in Cisco NX-OS and in Cisco Fabric Manager and Device Manager provide end-to-end diagnostics.

The Cisco MDS 9000 Family is the ideal complement to the Cisco Nexus 5000 Series to create a unified fabric capable of providing a homogeneous service level to the well-established Fibre Channel initiators and to the new generation of FCoE initiators.

## For More Information

Please refer to the following links for specific technical documentation:

- Cisco MDS 9000 Family configuration guides: http://www.cisco.com/en/US/products/ps5989/products_installation_and_configuration_guides_list.html
- Cisco Nexus 5000 Series configuration guides: http://www.cisco.com/en/US/products/ps9670/products_installation_and_configuration_guides_list.html

## Appendix A: FCoE Step-by-Step Configuration

This section describes how to configure a server connecting to the Cisco Nexus 5000 Series accessing Fibre Channel storage on a Cisco MDS 9000 Family switch through FCoE. The topology in Figure 13 will be used in both switching and NPV modes.

**Figure 13.** Topology

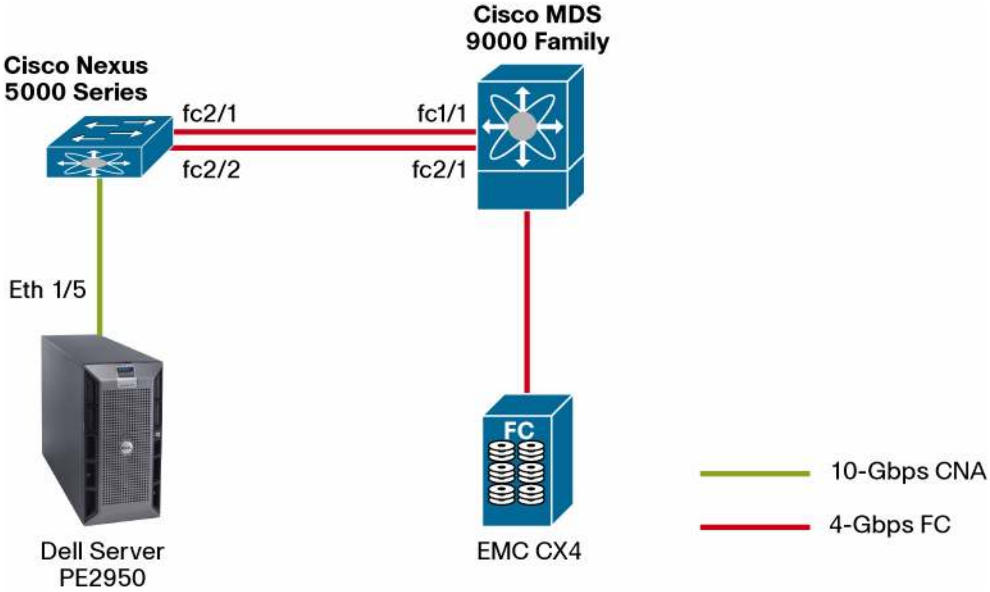

Please note that the goal of this section is to provide guidance on how to connect the Cisco Nexus 5000 Series to the Cisco MDS 9000 Family, so standard Fibre Channel features (for example, zoning, IVR, and AAA) are not documented. For more information about these features, please consult the Cisco MDS 9000 Family configuration guide (referenced in the “For More Information” section).

The following three scenarios can be used when connecting the Cisco Nexus 5000 Series to the Cisco MDS 9000 Family. For all three, the processes for turning on FCoE and configuring the FCoE interfaces are the same:

1. Cisco Nexus 5000 Series in switch mode: PortChannel to Cisco MDS 9000 Family
2. Cisco Nexus 5000 Series in NPV mode (previous to Cisco NX-OS 4.2(1)N1(1)): individual NP uplinks to Cisco MDS 9000 Family
3. Cisco Nexus 5000 Series in NPV mode (Cisco NX-OS 4.2(1)N1(1) or later): NP PortChannel to Cisco MDS 9000 Family

Step 1. (Common to the Three Scenarios): Turn on FCoE

By default on the Cisco Nexus 5000 Series Switch, FCoE is not turned on. To enable the FCoE feature, enter the following:

```
EBC-N5K-1# configure terminal
EBC-N5K-1(config)# feature fcoe
EBC-N5K-1(config)# copy running-config startup-config
```

Cisco NX-OS 4.2(1)N1(1) and later do not require a switch reboot at this point. Earlier releases need to perform a switch reboot:

```
EBC-N5K-1(config)# reload
```

Run the following command to verify that FCoE is enabled:

```
EBC-N5K-1# show feature
Feature Name          Instance  State
--------------------  --------  --------
cimserver             1         disabled
fabric-binding        1         disabled
```

```
fc-port-security      1         disabled
fcoe                  1         enabled
```

### Cisco Nexus 5000 Series in Switch Mode

A Cisco Nexus 5000 Series Switch in switch mode can connect to a Cisco MDS 9000 Family switch through a PortChannel; using trunking, each EISL can carry multiple VSANs.

Step 2.   (Cisco Nexus 5000 Series in Switch Mode): Build a PortChannel to the Cisco MDS 9000 Family Switch

Configure the Cisco Nexus 5000 Series Switch:

```
EBC-N5K-1#configure terminal
EBC-N5K-1(config)# interface san-port-channel 1
EBC-N5K-1(config-if)# interface fc2/1-2
EBC-N5K-1(config-if)# channel-group 1
fc2/1 fc2/2 added to san-port-channel 1 and disabled please do the same operation
on the switch at the other end of the channel, then do "no shutdown" at both ends
to bring them up
EBC-N5K-1(config-if)# no shut
EBC-N5K-1(config-if)# interface san-port-channel 1
EBC-N5K-1(config-if)# no shut
```

Configure the Cisco MDS 9500 Series switch:

```
MDS-9506-1# configure terminal
MDS-9506-1(config)# interface port-channel 1
MDS-9506-1(config-if)# interface fc1/1, fc2/1
MDS-9506-1(config-if)# channel-group 1 force
fc1/1 fc2/1 added to port-channel 1 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both end to bring them up
MDS-9506-1(config-if)# no shut
MDS-9506-1(config-if)# interface port-channel 1
MDS-9506-1(config-if)# no shut
```

To verify that the PortChannel on the Cisco Nexus 5000 Series Switch is running properly, enter the following command:

```
EBC-N5K-1# show interface san-port-channel 1
san-port-channel 1 is trunking
Hardware is Fibre Channel
Port WWN is 24:01:00:0d:ec:a3:0f:c0
Admin port mode is auto, trunk mode is on
snmp link state traps are enabled
Port mode is TE
Port vsan is 1
Speed is 8 Gbps
Trunk vsans (admin allowed and active) (1,100)
Trunk vsans (up) (1,100)
Trunk vsans (isolated) ()
Trunk vsans (initializing) ()
5 minutes input rate 248 bits/sec, 31 bytes/sec, 0 frames/sec
5 minutes output rate 184 bits/sec, 23 bytes/sec, 0 frames/sec
2697 frames input, 208940 bytes
0 discards, 0 errors
```

```
0 CRC, 0 unknown class
0 too long, 0 too short
2703 frames output, 162308 bytes
0 discards, 0 errors
4 input OLS, 4 LRR, 2 NOS, 0 loop inits
6 output OLS, 8 LRR, 0 NOS, 0 loop inits
Member[1] : fc2/1
Member[2] : fc2/2
```

### Step 3. (Common to the Three Scenarios): Create the FCoE Interface

To create the FCoE interface, which is the virtual Fibre Channel (vFC) interface, some basic steps are needed:

1. Create a VSAN.
2. Create an FCoE VLAN.
3. Create the vFC, making sure the vFC is in the correct VSAN.
4. Configure Ethernet interface for FCoE traffic.

Completing these tasks will help ensure that connection of a FCoE CNA to the Cisco Nexus 5000 Series Switch is successful.

### Step 4. Create a VSAN

In the example in this topology, the storage on the Cisco MDS 9000 Family switch resides on VSAN 100, so to help ensure that the vFC interface that is created on the Cisco Nexus 5000 Series Switch can communicate with those storage devices, you need to create this VSAN. Enter the following to complete this task:

```
EBC-N5K-1# configure terminal
EBC-N5K-1(config)# vsan database
EBC-N5K-1(config-vsan-db)# vsan 100
EBC-N5K-1(config-vsan-db)# show vsan
```

### Step 5. Create an FCoE VLAN

As a best practice, you should create a separate VLAN for FCoE traffic, to separate FCoE traffic from other, standard Ethernet traffic. The following demonstrates how to create this FCoE VLAN:

```
EBC-N5K-1# configure terminal
EBC-N5K-1(config)# vlan 100
EBC-N5K-1(config-vlan)# fcoe vsan 100
EBC-N5K-1(config-vlan)# show vlan fcoe
```

### Step 6. Create the vFC Interface

To create the virtual Fibre Channel interface, connect the CNA on Ethernet interface eth1/5. Follow these steps to create the vFC and have it reside in VSAN 100. As a best practice in creating the vFC number, use the number of the physical interface. For example, ethernet 1/5 would have vFC 5. This approach is not a requirement but is a recommended best practice.

```
EBC-N5K-1# configure terminal
EBC-N5K-1(config)# interface vfc 5
EBC-N5K-1(config-if)# bind interface ethernet 1/5
EBC-N5K-1(config-if)# no shut
EBC-N5K-1(config-if)# vsan database
this will get to the VSAN database
```

```
EBC-N5K-1(config-vsan-db)# vsan 100 interface vfc5
EBC-N5K-1(config-vsan-db)# show vsan membership
vsan 1 interfaces:
fc2/1 fc2/2 fc2/3 fc2/4 san-port-channel 1
vsan 100 interfaces:
vfc5 ...
```

Step 7. Configure the CNA Ethernet Interface

One final step is needed to be able to bring up the vFC on the Cisco Nexus 5000 Series Switch. You need to make sure that the FCoE VLAN (in this case, VLAN 100) can traverse the Ethernet interface (in this example, eth1/5). Follow these steps to complete this task:

```
EBC-N5K-1#configure terminal
EBC-N5K-1(config)#interface ethernet 1/5
EBC-N5K-1(config-if)# switchport mode trunk
EBC-N5K-1(config-if)# switchport trunk allowed vlan 1,100
```

**Note:** The above command is not needed but if you like to prune the allowed vlans, make sure the FCoE vlan is on he allowed list.

```
EBC-N5K-1(config-if)# spanning-tree port type edge trunk
EBC-N5K-1(config-if)# show interface vfc5
vfc5 is up ....
```

### Cisco Nexus 5000 Series in NPV Mode: Prior to Cisco NX-OS 4.2(1)N1(1)

When you configure the Cisco Nexus 5000 Series Switch in NPV mode, the switch will be reloaded. When you switch between the two modes (switch and NPV), the switch will always need to be reloaded. Since the Cisco Nexus 5000 Series Switch will reload when you enable NPV, it is a best practice to back up the configuration to bootflash memory or some offsite location. To enable NPV on the Cisco Nexus 5000 Series, follow these steps:

```
EBC-N5K-1# copy running-config bootflash:backup-cfg.txt
EBC-N5K-1# configure terminal
EBC-N5K-1(config)# npv enable
```

Verify that boot variables are set and the changes are saved. Changing to npv mode erases the current configuration and reboots the switch in npv mode. Do you want to continue? (y/n): y

**Note:**

- If you have not enabled FCoE yet, do so by performing Step 1: Turn on FCoE.
- The Cisco MDS 9500 Series switch needs to have NPIV enabled, and the interfaces to which the Cisco Nexus 5000 Series Switch is connected must be set to auto or F-port. Those ports also must be on the same VSAN for the connection to work.

Step 1. (NPV Mode with Older Code Only): Enable NP Uplinks

When the Cisco Nexus 5000 Series Switch comes back online, by default all the Fibre Channel interfaces will be set to NP mode with the default VSAN set to VSAN 1. When configuring the Cisco Nexus 5000 Series in NPV mode, you always should verify that you have at least one NP-port running in the correct VSAN in which the CNAs will be residing. Follow these steps to help ensure that the NP uplinks are functioning:

Configure the Cisco Nexus 5000 Series Switch:

```
EBC-N5K-1# configure terminal
```

```
EBC-N5K-1(config)# vsan database
EBC-N5K-1(config-vsan-db)# vsan 100
EBC-N5K-1(config-vsan-db)# vsan 100 interface fc2/1-2
EBC-N5K-1(config-vsan-db)# interface fc2/1-2
EBC-N5K-1(config-if)# no shut
```

**Note:** To verify that the NP links are up and in the correct VSAN, run the following command on the Nexus 5000.

```
EBC-N5K-1# show interface brief
```

Configure the Cisco MDS 9500 Series switch:

```
MDS-9506-1# configure terminal
MDS-9506-1(config)#vsan database
MDS-9506-1(config-vsan-db)# vsan 100 interface fc1/1, fc2/1
MDS-9506-1(config-vsan-db)# interface fc1/1, fc2/1
MDS-9506-1(config-if)# no shut
MDS-9506-1(config-if)# show vsan membership
vsan 1 interfaces:
fc1/2 fc1/3 fc1/5 fc1/6 ...
vsan 100 interfaces:
fc1/1 fc1/4 fc1/7 fc1/8
fc1/9 fc1/10 fc1/11 fc1/12
fc2/1 fc2/2 fc2/3 fc2/4
```

After setting up the VSAN, create the FCoE on the Cisco Nexus 5000 Series Switch in NPV mode is the exact same way as in switch mode. Follow the procedure described in Step 3: Create the FCoE Interface.

### Cisco Nexus 5000 Series in NPV Mode: Cisco NX-OS 4.2(1)N1(1) or Later

Starting from Cisco NX-OS 4.2(1)N1(1), the Cisco Nexus 5000 Series supports trunking and Port Channels to the Cisco MDS 9000 Family switches.

**Note:**

- If you have not enabled FCoE yet, do so by performing Step 1: Turn on FCoE.
- If you have not enabled NPV mode yet, do so by performing the procedure described in the preceding section.
- The Cisco MDS 9500 Series needs to have NPIV enabled, and the interfaces to which the Cisco Nexus 5000 Series Switch is connected must be set to auto or F-port.

Step 1. (NPV Mode with Newer Code Only)

These are the basic configuration steps for building a trunking PortChannel from a Cisco Nexus 5000 Series Switch in NPV mode to a Cisco MDS 9000 Family switch with NPIV enabled.

**Note:**

- By default, the SAN PortChannel is enabled on all configured VSANs and will come up on the matching VSANs on both sides like regular TNP port.
- By default, trunking is enabled. You can enter switchport trunk mode off on the PortChannel context on both ends if you want to disable trunking.

Configure the Cisco Nexus 5000 Series Switch:

```
SwitchA(config-if)#configure terminal
SwitchA(config-if)# interface san-port-channel 1
SwitchA(config-if)# interface fc2/1-2
SwitchA(config-if)# channel-group 1
fc1/1 fc2/1 added to port-channel 1 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring it up
SwitchA(config-if)# no shut
SwitchA(config-if)# interface  san-port-channel 1
SwitchA(config-if)# channel mode active
```

Configure the Cisco MDS 9000 Family switch:

```
MDS-SAN-A(config)# configure terminal
MDS-SAN-A(config)# feature fport-channel-trunk
MDS-SAN-A(config)# interface fc1/1, fc2/1
MDS-SAN-A(config-if)# channel-group 1
fc1/1 fc2/1 added to port-channel 1 and disabled
please do the same operation on the switch at the other end of the port-channel,
then do "no shutdown" at both ends to bring it up
MDS-SAN-A(config-if)# no shut
MDS-SAN-A(config-if)# interface port-channel 1
MDS-SAN-A(config-if)# channel mode active
```

You create the FCoE interface on the Cisco Nexus 5000 Series in NPV mode in the same way as in switch mode. Follow the procedure described in Step 3: Create FCoE Interface.

To show output on both switches, use the commands shown in section "Cisco Nexus 5000 Series in Switch Mode" (for the Cisco MDS 9000 Family, enter show interface port-channel 1; for the Cisco Nexus 5000 Series, enter show interface san-port-channel 1).

**Americas Headquarters**
Cisco Systems, Inc.
San Jose, CA

**Asia Pacific Headquarters**
Cisco Systems (USA) Pte. Ltd.
Singapore

**Europe Headquarters**
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Printed in USA C11-610619-00 07/10