

Transparently Migrate a SAN from a Heterogeneous Environment to a Cisco MDS 9000 Family SAN

What You Will Learn

Today's SAN administrators are faced with the need for more storage capacity and speed. They require high-performance and redundant SAN networks that can both meet their current demands and scale for growth in the future. To accommodate these new requirements, SAN administrators often need to migrate or upgrade from their existing storage networks.

Cisco® MDS 9000 Family SAN switches are recognized across the data center industry for their reliability, flexibility, performance, and investment protection. The Cisco MDS 9000 Family portfolio includes the Cisco MDS 9700 and 9500 Series Multilayer Directors, the multiprotocol Cisco MDS 9200 Series Multiservice Switches, and the fixed form-factor Cisco MDS 9100 Multilayer Fabric Switches. These switches provide flexibility, redundancy, high availability, and high performance at the core and edge levels with room for future growth.

Migrating SANs from one vendor to another requires a specific plan that includes design, configuration, and implementation processes along with post migration analysis. This document helps you evaluate appropriate options for SAN conversion from third-party solutions to Cisco SANs using the Cisco MDS 9000 Family.

For additional information about Cisco SAN design principles, see the [Cisco SAN design guides](#).

Scope of the Document

This document provides an overview of SAN technology terms, various features, interoperability requirements, and licensing and other verification checks to consider when migrating to a Cisco MDS 9000 Family SAN. This document also discusses some design parameters and best practices to help guide the migration process.

Migration Concepts

When migrating to a Cisco MDS 9000 Family SAN, you can choose among three migration methods: rip and replace, cap and grow, and interoperate. The choice of migration method is determined by several criteria, including whether you want a single-vendor or mixed-vendor operation, risk-mitigation needs, migration timeline, connectivity requirements and overall fabric capacity during the migration process.

- **Rip and Replace:** As the name suggests, with this approach you simply replace third-party switches with preconfigured Cisco MDS 9000 Family switches.
- **Cap and Grow:** Cisco MDS 9000 Family switches are connected in a parallel topology to an existing third-party SAN. New and existing servers and targets are attached to the Cisco MDS 9000 Family switches. Third-party switches are removed one by one while both SANs are running and serving traffic simultaneously and independently during the migration. This option allows greater scalability and growth in the future.
- **Interoperate:** Cisco MDS 9000 Family switches are connected to third-party switches using interoperate mode. Both vendors work together for a period of time before third-party switches are removed in phases.

Migration Process

Migrating or upgrading a SAN from one vendor to the Cisco MDS 9000 Family product line can be relatively easy if proper guidelines are followed. Differences in technical terms can sometimes be confusing, however, because every vendor has its own technical terms for similar functions. For ease of migration, the migration process is divided into steps, narrowing the change window required, focusing the tasks, and helping mitigate risk and ease deployment. These are the main steps:

- **Prepare:** Analyze the current storage infrastructure, business requirements, and risks. Identify critical servers, storage subsystems, and applications. Prepare a rollback procedure in case rollback is required. Prepare or update the SAN and storage diagram to meet new requirements. Prepare all device configurations (zone conversion, VSAN configuration, etc.) in advance and have them ready during the change window. Depending on migration method used, most of the configuration can be completed ahead of time.
- **Plan and design:** Identify migration options and create a migration strategy. Identify any new additions and future requirements for the SAN fabric at this stage. This step will help the new SAN environment have enough flexibility to meet your needs longer.
- **Operate and optimize:** Perform the actual migration, move cables and chassis, and configure the new setup. After migration is complete, you can implement continuous monitoring and optimization to identify and mitigate risk and tune the infrastructure to accommodate new projects and applications as the need arises.

Prepare

The process of SAN migration starts with preparation. This step helps you define, scale, and meet your migration goals.

- **Inventory your network:** Prepare a list of hosts, targets, and switches and their hardware versions.
- **Verify compatibility:** Verify your inventory with the software and hardware compatibility matrix and switch interoperability matrix.
- **Upgrade components:** You may need to upgrade some components to meet the requirements of the support matrices. Upgrading will reduce the likelihood that incompatible hardware or conflicts with existing software will delay the migration process. You then may need to upgrade the hardware and software on the list you prepared in the previous steps.
- **Assess the SAN:** Before starting the migration, collect current metrics and plan to collect future metrics for proper assessment. This step will help you avoid bottlenecks later in the migration process or in the near future. Statistics such as bandwidth requirements (based on existing and new needs) and projected growth for bandwidth, targets, hosts, etc. can help you gauge the right set of requirements.
- **Validate applications:** To set service-level agreements (SLAs), application validation is essential. You need to consider current and expected future latency associated with growth. In addition, multipath connectivity is required during migration. Cisco Prime™ Data Center Network Manager (DCNM) can be useful for performing this check.

You should also validate any hardware or software upgrade by testing application-level connectivity between the fabrics along with intended initiator-target pairs. When required, this testing should also be conducted on important features and functions for any site-to-site replication, data mobility, etc. However, exhaustive feature and function testing is not always practical because it may require dedicated test ports in the production fabric and storage subsystems, but such tasks, when they can be performed, boost confidence in the migration for the operation team.

The following information about the existing SAN network for each fabric also will help you define an appropriate migration plan:

- Total number of host and server ports
- Total number of storage (disk and tape) ports
- Total bandwidth requirements from the host edge
- Total bandwidth requirements from the storage edge
- Current oversubscription ratio from the host to storage
- Expected oversubscription ratio from the host to storage
- List of third-party fabric licenses in use

Plan and Design

The planning and design phase involves both physical and architectural elements.

Physical Planning

Physical planning includes identification of space, cooling, power, power distribution unit (PDU), cabling, and cable rack requirements. Different chassis from different vendors have their own sets of requirements. More details about specific Cisco MDS 9000 Family chassis can be found in the individual [data sheets](#). Some of the important hardware components that should be considered are discussed here.

Chassis Power, Cooling, and Airflow

Chassis cooling characteristics and proper spacing for airflow are important for efficient operation of the chassis (Table 1). For Cisco MDS 9000 Family switches, the hardware installation guide provides details about the height, width, and depth of the chassis.

Table 1. Airflow Direction for Cisco MDS 9000 Family Switches

Cisco Chassis Name	Airflow Direction
Cisco MDS 9710 Multilayer Director	Front (port) to back
Cisco MDS 9250i Multiservice Fabric Switch	Front (port) to back
Cisco MDS 9513 Multilayer Director	Left to right
Cisco MDS 9506 Multilayer Director	Left to right
Cisco MDS 9222i Multiservice Modular Switch	Left to right
Cisco MDS 9148 Multilayer Fabric Switch	Back to front (port)

For more information, please refer to the [hardware installation guides](#).

Power Planning

For the new SAN switches, verify that you have the correct amount of AC and DC power for proper operation, the correct power cord connectors, and PDUs and uninterruptible power supplies (UPSs) with the appropriate capacity. The Cisco [site preparation checklist](#) includes more information about power requirements in the technical specification sections of the respective [hardware installation guides](#).

Architectural Planning

Architectural planning includes all design-related details, including network topology, cable diagrams, cabling techniques, power-plug connections and positions, cabling mechanisms for different chassis, PDU placement, and air conditioning and air circulation requirements. Architectural planning requires more information and analysis than physical planning, including information about:

- Fibre Channel cable connections to the chassis (how are they stacked across the chassis and blade: vertical or horizontal?)
- Power cable connections to the chassis (some chassis have front-end and some have back-end connections)
- PDU connectors and types
- Space for new hardware (form factors and rack unit size and depth)
- Cable length specifications
- Air space required around the chassis for proper air flow
- Placement of PDUs and power cable connections for the chassis
- Console connections to the chassis
- Front-door placement and space required at the front of the chassis
- Space required for maintenance (for example, for pulling out the line card or replacing the fan tray in the Cisco MDS 9513)

Software Interoperability Planning

For migration, switch interoperability is an important consideration. Switches from different vendors should be able to communicate with each other, and software interoperability plays a major role in helping ensure that they can. Cisco has a variety of guides to address interoperability concerns when interoperability mode is considered. Interoperability requires storage, host firmware, driver compatibility. Software running on SAN switches must also be compatible. Interoperability guides such as the [Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide](#) along with the [Cisco MDS 9000 NX-OS Software Release Notes](#), [Cisco MDS 9000 NX-OS Hardware and Software Compatibility Information and Feature Lists](#), [Cisco Data Center Interoperability Support Matrix](#), and [Cisco MDS NX-OS Release Compatibility Matrix for Storage Service Interface Images](#) can help address any interoperability questions. Although Cisco recommends tested and verified code levels for interoperability, original storage manufacturer (OSM) partners may have different levels of code releases and support matrices. In such cases, please refer to OSM partner's support matrix for the code level.

Interoperability Modes

Multivendor switch interoperability is part of the Fibre Channel standards. INCITS introduced the FC-SW-2 standard, which defines switch-to-switch interconnectivity and operation requirements, including features such as fabric addressing and configuration, Fabric Shortest Path First (FSPF) Protocol, zone merge, and distributed services parameters. Most vendors support (or have supported) standards-based interoperability. However, some vendors support proprietary operating modes to position their product features and functions that differ from the Fibre Channel standards. This support results in an environment in which switches from one vendor may not interact properly with switches from another vendor, making interoperability impossible. Cisco supports interoperability with other vendors to provide customers with more options and flexibility in creating SAN solutions. Cisco offers a comprehensive set of interoperability modes to allow interoperation with third-party switches.

Cisco provides four interoperability modes to support interoperability with different switch vendors: mode 1 (Fibre Channel standards based), mode 2 (Brocade native part ID [PID] = 0), mode 3 (Brocade native PID = 1), and mode 4 (McData native).

Brocade has two modes to support interoperability: native mode for its own switches, and standard mode (or open fabric mode) to support McData switches.

Table 2 summarizes the Cisco interoperability modes and their compatibility with third-party switches.

Table 2. Cisco Interoperability Modes and Compatibility with Third-Party Switches

Cisco Interoperability Mode	Brocade Native Mode	Brocade Interoperability Mode	McData Native Mode	McData Open Fabric Mode
Native	–	–	–	–
Mode 1	–	Yes	–	Yes
Mode 2	Yes (PID 0)	–	–	–
Mode 3	Yes (PID 1)	–	–	–
Mode 4	–	–	Yes	–

–: Not supported or does not work

Yes: Supported configuration

The interoperability mode of Cisco MDS 9000 Family SAN Switches can be enabled on a per-VSAN basis with no requirement to reboot the switch. When you enable the vendor native interoperability mode on a Cisco switch, no additional configuration is required on Brocade or McData switches running in their native modes.

- Default or Cisco MDS native mode:** This is the default mode or behavior for a VSAN that is communicating with a Cisco MDS 9000 Family switch-based SAN. Cisco MDS native mode is fully compatible with Fibre Channel standards. Advanced features such as trunking, PortChannels, and VSANs are not supported on third-party switches or Cisco MDS 9000 Family ports connected to third-party switches.
- Interoperability mode 1:** This is the FC-MI standard interoperability mode. This mode interoperates with Brocade switches that have been configured with Brocade interoperability mode. This mode is VSAN specific. Brocade reduces the capabilities of features such as port zoning, trunking, QuickLoop, Fabric Assist, Secure Fabric OS, and virtual flow control in this mode.
- Interoperability mode 2:** This mode, also known as the interoperability mode for existing Brocade switches, allows transparent integration with Brocade switches running in native mode with the core value of PID = 0.

- **Interoperability mode 3:** This mode was introduced for Brocade switches that contained more than 16 ports. With this interoperability mode, Cisco switches will interoperate with Brocade switches in their native mode and operating with a core value of PID = 1.
- **Interoperability mode 4:** This mode, also known as interoperability mode 4 for existing switches, provides interoperability between Cisco MDS 9000 Family switches and McData switches operating in native mode. This mode supports only domain IDs 1 through 31.
 - Adding Cisco MDS 9000 Family Switches to existing third-party fabrics does not require outages.
 - The Cisco MDS 9000 Family switch interoperability mode affects only the configured VSAN; all other VSANs are unaffected.
 - Cisco MDS 9000 Family switch traffic engineering (TE) ports can simultaneously carry VSANs that are running any or all interoperability modes as well as Cisco MDS 9000 Family switch native mode.
 - No configuration changes are needed to the Brocade switches if they are already set to native interoperability mode.

Table 3 summarizes interoperability modes. For more detailed information, please refer to the [Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide](#).

Table 3. Cisco Interoperability Mode and Feature Limitations with Third-Party Switches

Cisco Interoperability Mode	Description	Brocade or McData Mode	Domain Range	Domain ID and Port Support
Mode 1	Standards-based interoperability	Brocade interoperability mode 1 and McData open fabric mode	97-127	No
Mode 2	Brocade native	Brocade PID = 0	1-239	Yes
Mode 3	Brocade native	Brocade PID = 1	1-239	Yes
Mode 4	McData native	McData native	1-31	Yes

Open Fabric mode

In most cases, interoperability between Cisco and third-party Fibre Channel switches should be considered a temporary solution for the duration of a migration to address ongoing concerns about code-level compatibility, feature compatibility and restrictions, etc.

Licensing

Before migrating from third-party SANs to Cisco SANs, it is important to obtain the correct license set for Cisco MDS 9000 Family switches. Most Cisco MDS 9000 Family software features are included in the Base switch license. However, some features are logically grouped into add-on packages that must be licensed separately. Examples include the Cisco MDS 9000 Enterprise Package, Cisco MDS 9000 Mainframe Package, and Cisco Prime DCNM for SAN Advanced Edition. Table 4 provides a comparison of Brocade and Cisco licenses.

Table 4. License Comparison: Brocade and Cisco

Brocade License	Equivalent Cisco License
10 Gigabit Fibre Channel over IP (FCIP) and Fibre Channel (10-Gbps license)	Standard feature (Base license) on MDS 9250i
Adaptive networking with quality of service (QoS)	Enterprise license
Advanced extension (buffer-to-buffer [B2B] credits)	Enterprise license
Advanced Fibre Connection (FICON) acceleration	Mainframe license
Brocade Extended Fabrics	Standard feature (Base license)
Brocade Fabric Watch	Standard feature (Base license)

Brocade License	Equivalent Cisco License
Brocade Inter-Switch Link (ISL) trunking	Standard feature (Base license)
Fibre Channel over Ethernet (FCoE)	Standard feature (Base license)
FICON Management Server	Mainframe license
High-performance extension over FCIP and Fibre Channel; also known as FC-IP services	Standard license on 9250i and 9222i SAN Extension license on SSN-16 and MSM 18/4 module
Integrated routing	Enterprise license

Migration Tools: Zone Migration Wizard

To migrate a third-party SAN to a Cisco MDS 9000 Family SAN, Cisco provides a Zone Migration Wizard in the Cisco Prime DCNM for SAN module. The wizard collects and automatically converts third-party zones into Cisco MDS 9000 Family zones. This GUI-based tool helps migrate the zoning configuration to Cisco SAN switches.

Step 1. Open the Zone Migration Wizard.

- a. Log into Cisco Prime DCNM using administrator or equivalent credentials.
- b. Click DCNM-SAN at the top-right corner of the screen.
- c. A new window opens and prompts you to log in again. Enter the administrator or equivalent user credentials again.
- d. When you are logged in, the Cisco Prime DCNM for SAN dashboard appears. Choose Zone > Migrate Non-MDS Database. The following screen opens:

Step 1 of 5: Migrate Zones from Switch

Generate MDS Zone config file from McData or Brocade. McData zones are not stored on switch, the file format must be 'WWN=<aliasName>' - one entry per line.
NOTE: Tested on Brocade version 3.1.1 and McData version 04.01.00 only.

Switch:

User:

Password:

VSAN Id:

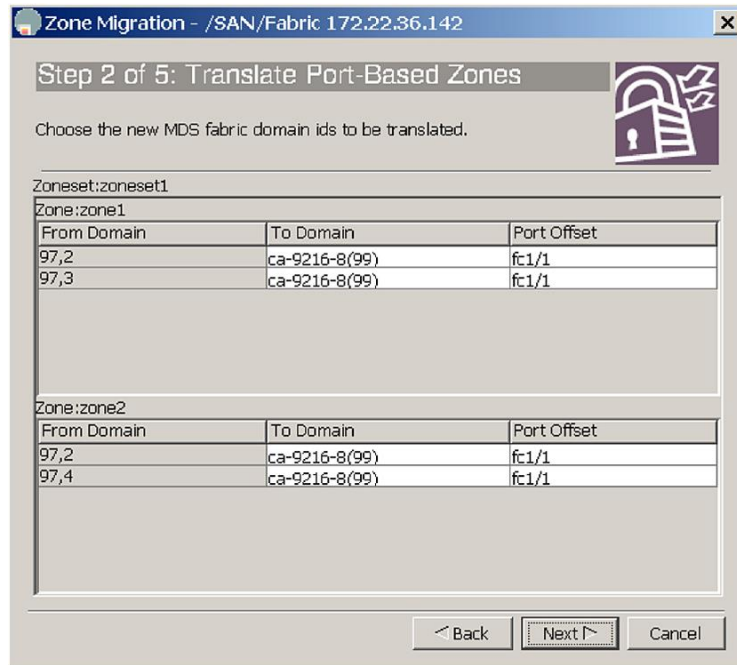
Migrate From: Brocade Effective Defined

Migrate McData Alias?


Alias File: ...

Next > Close

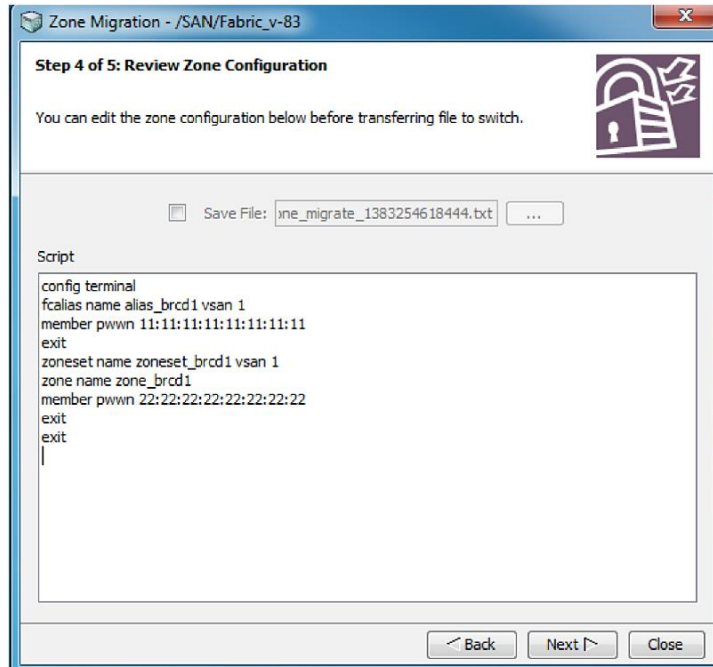
- Step 2. Convert the domain and port-based zone members into Cisco MDS 9000 Family interface-based zone members:



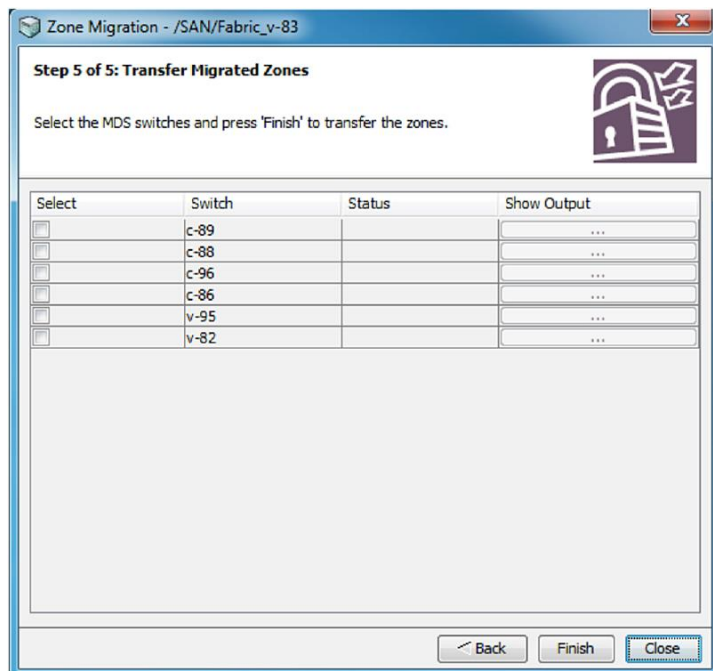
- Step 3. If any zone member types are unknown, you can manually change them:



- Step 4. Script compatible with the Cisco MDS 9000 Family command-line interface (CLI) is produced, which can be either automatically applied to the switch or VSAN or saved to a text file for future use:



- Step 5. Select the Cisco MDS 9000 Family switches to which you want to apply the zoning configuration:



Note: This step is not related to the Cisco Prime DCNM for SAN seed switch.

Migration Method

SAN best practices typically call for two fabrics for redundancy, referred to as Fabric A and Fabric B in this document. Several migration options are available, but the three methods discussed here are preferred by Cisco.

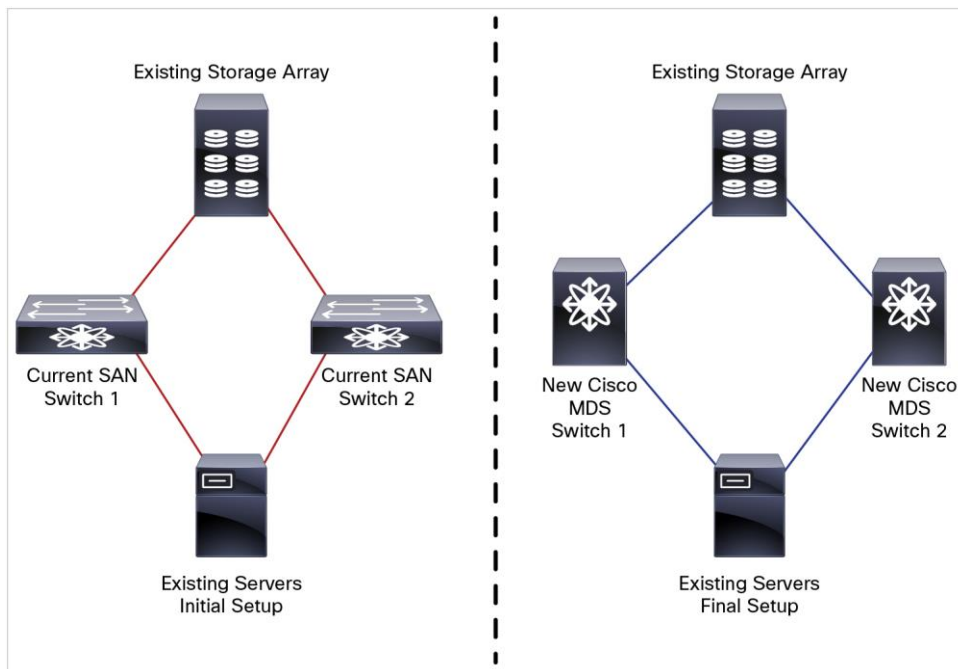
Rip and Replace

The rip-and-replace approach to migration is also called fabric-by-fabric replacement. The advantage of this option is that the migration process takes very little time to complete, and you avoid interoperability challenges by replacing third-party switches with Cisco MDS 9000 Family switches on a per-fabric basis.

With this option, Fabric A hardware will be replaced while Fabric B remains running, providing redundancy and reducing downtime. As a best practice, you should disable the host and target host bus adapters (HBAs) connected to Fabric A prior to the migration to avoid any impact on applications. After the Fabric A hardware is replaced and verified to be up and running, all host and target connections to Fabric A are reenabled. After verification that server-to-target operation is restored over Fabric A, you repeat the same process on Fabric B. The conversion of Fabric B can occur in the same change window or in a subsequent change window at a later date, depending on user requirements.

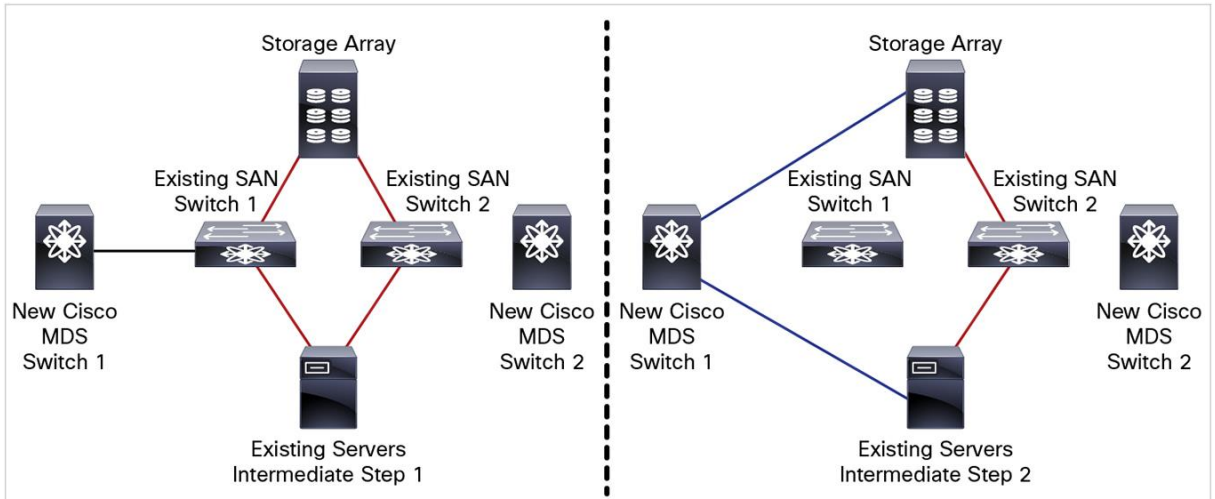
Figure 1 shows the migration process.

Figure 1. Rip-and-Replace Migration Method



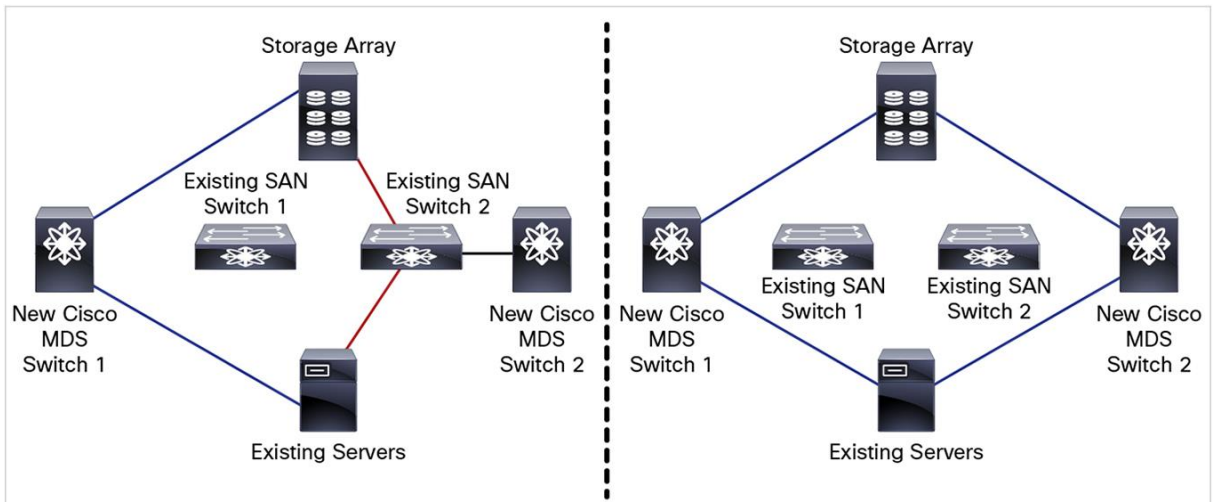
Step 1. Connect new Cisco MDS Switch 1 to existing third-party Switch 1 and preconfigure the new switch by adding zone information, propagating the zone database (using the Cisco Prime DCNM for SAN Zone Migration Wizard), configuring the VSAN, and selecting the appropriate interoperability mode so that the new switch is ready to take over the connections between the existing host and storage ports. You can also use the Cisco Prime DCNM for SAN Zone Migration Wizard to help you migrate the zone configuration. Verify the path to the storage through Switch 2 using Symantec Dynamic Multipathing (DMP), EMC PowerPath, IBM Subsystem Device Driver (SDD), or IBM Multipath IO (MPIO). Now I/O is flowing only through Switch 2 (Figure 2).

Figure 2. Rip-and-Replace Step 1 and Step 2



Step 2. Move the cables connecting the host and storage from third-party Switch 1 to new Cisco MDS Switch 1. Verify that the host and the storage ports can log into new Cisco MDS Switch 1 and that the ports are in the correct VSAN and part of the correct zone and zone set. Also verify application connectivity through both paths (Figure 3).

Figure 3. Rip-and-Replace Step 3 and Step 4



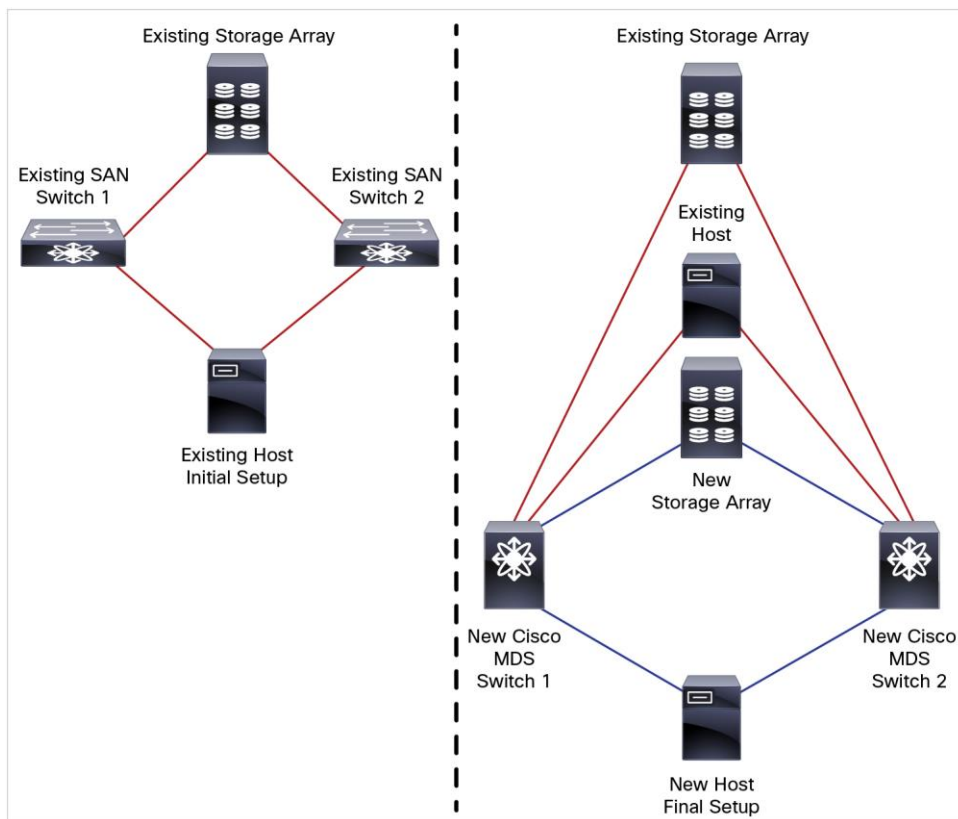
Step 3. Resume the I/O path on the original primary fabric using new Cisco MDS Switch 1. All the connection flow should now go through new Cisco MDS Switch 1 and SAN Switch 2 simultaneously.

Step 4. Repeat the instructions in step 2 thru step 3 to migrate the path from third-party Switch 2 to new Cisco MDS Switch 2.

Cap and Grow

With the cap-and-grow method, a complete new storage environment setup is created using new Cisco MDS 9000 Family switches. After the new infrastructure is built, you can start adding to the Cisco MDS 9000 Family fabric and start moving existing storage and host ports from the existing SAN as well. This method is appropriate when you are deploying a new application environment or upgrading a large number of storage devices. You could start with one director-class switch and end up swapping the whole SAN infrastructure fabric during the scheduled change window. For example, with Fabric A and Fabric B in production, start with completely new Fabric C and Fabric D. Start deploying new storage and host ports to new Fabric C and Fabric D. After these are operational, start migrating existing host and storage ports from Fabric A and Fabric B to new Fabric C and Fabric D, respectively. This approach provides continuity for the business along with room to expand the infrastructure. It also reduces interoperability risks and migration complexity and provides the flexibility to roll back the changes in the event of any emergency situation. In this type of migration, the Cisco MDS 9000 Family switches are introduced in parallel with the existing SAN infrastructure (Figure 4).

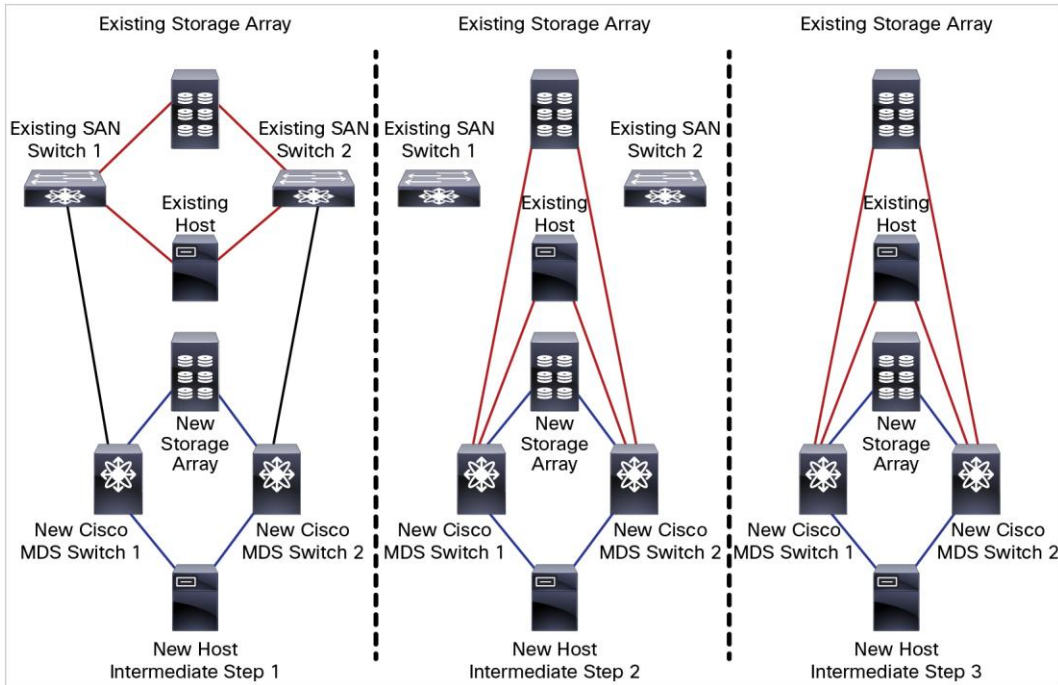
Figure 4. Cap-and-Grow Migration Method



- Step 1. Connect the new Cisco MDS 9000 Family switches to the new storage devices and new hosts. Then perform zone migration, VSAN migration, and migration of other configurations from the third-party switches to the Cisco MDS 9000 Family switches.
- Step 2. After configuration migration is complete, move the cables from the existing hosts and storage to the new Cisco MDS 9000 Family switches (Figure 5).

Step 3. Existing SAN switches can now be decommissioned (Intermediate Step 3 from Figure 5).

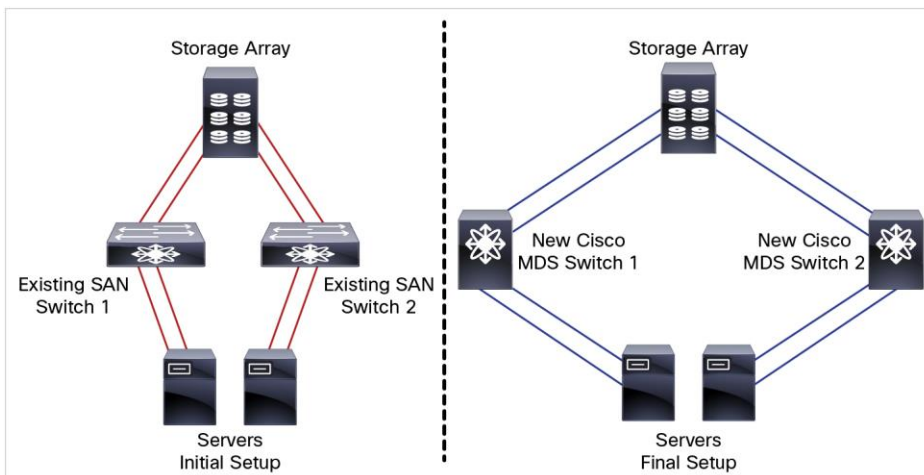
Figure 5. Cap-and-Grow Migration Method Intermediate Steps 2 and 3



Interoperate

With the interoperate approach, production traffic experiences little downtime, but more time is needed to complete the entire migration process. The timeline can stretch from a few weeks to a few months, depending on the size of the SAN infrastructure. With this method, the Cisco hardware is integrated into the third-party switch SAN environment. Then, slowly, the storage traffic is transferred to the Cisco MDS 9000 Family switches one switch, one application, and one blade chassis at a time. Cisco MDS 9000 Family SAN switches can interoperate with Brocade SAN switches and can offer the same scalability and capabilities. Figure 6 provides an overview of the interoperate migration process.

Figure 6. Interoperate Migration Method

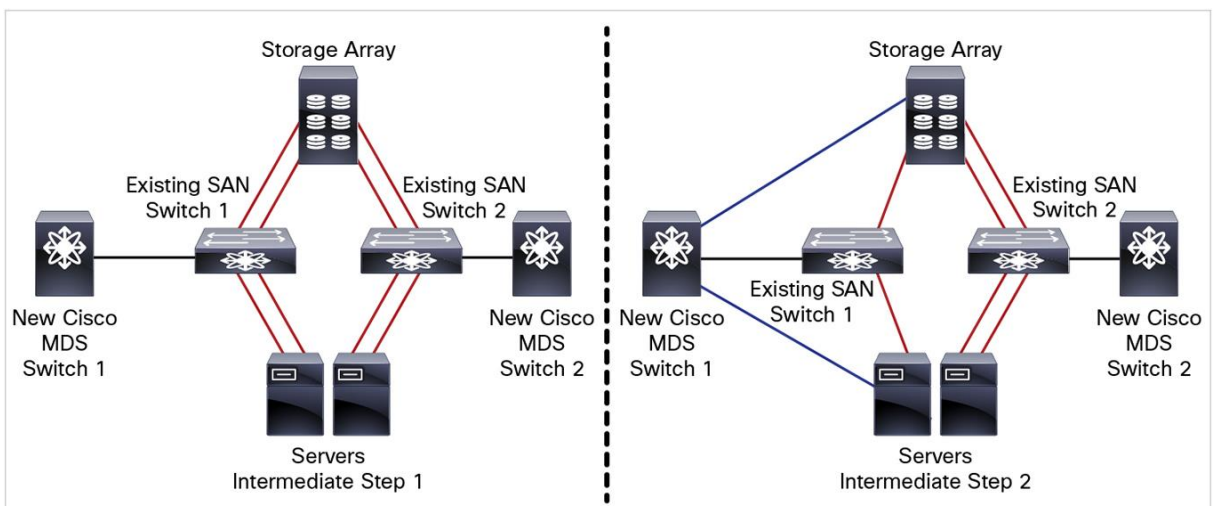


Basic Interoperate Migration Process

The following steps show the basic process for an interoperate migration.

- Step 1. Connect Cisco MDS Switch 1 and Cisco MDS Switch 2 to Brocade Switch 1 and Brocade Switch 2, respectively. Configure the appropriate interoperate mode on the Cisco MDS 9000 Family switches. After the interoperate mode is configured, start moving VSAN, zoning database, and other configurations from SAN Switch 1 to Cisco MDS Switch 1. Perform similar actions on SAN Switch 2 and Cisco MDS Switch 1.
- Step 2. After the configuration migration is complete, connect new Cisco MDS Switch 1 to the existing storage array and servers and start moving the production traffic through it. At this point, both the Brocade and Cisco SAN switches are passing production traffic in parallel (Figure 7).


Figure 7. Starting Interoperate Migration



- Step 3. After the Cisco MDS 9000 Family Switch operations are verified, move the remaining connections from SAN Switch 1 to Cisco MDS Switch 1.

- Step 4. At this point, SAN Switch 1 is ready to retire. Take it offline (Figure 8).

Figure 8. Retiring SAN Switch 1




Step 5. Repeat steps 2, 3, and 4 to replace SAN Switch 2 with Cisco MDS Switch 2 and then take SAN Switch 2 offline.

Interoperate Migration in a Core-Edge Network

A slightly more complex design using the interoperate migration method is required for a core-edge network in which core director-class switches are connected to the storage array on one side, and edge switches are connected on the other side (Figure 9).


Figure 9. Interoperate Migration Method for a Core-Edge Network



Step 1. Connect core Cisco MDS Switch 1 and core Cisco MDS Switch 2 to Brocade core SAN Switch 1 and Brocade core SAN Switch 2, respectively. Configure the appropriate interoperate mode on the Cisco MDS 9000 Family switches. After the interoperate mode is configured, start moving VSANs, zoning database, and other configurations from core SAN Switch 1 to core Cisco MDS Switch 1. Perform similar actions on core SAN Switch 2 and core Cisco MDS Switch 1.

Step 2. After the configuration migration is complete, connect core Cisco MDS Switch 1 to the existing storage array and servers and start moving the production traffic through it. At this point, both Brocade and Cisco SAN switches are passing production traffic in parallel (Figure 10).


Figure 10. Starting Interoperate Migration for a Core-Edge Network



Step 3. After the Cisco MDS 9000 Family switch operations are verified, move the remaining connections from core SAN Switch 1 to core Cisco MDS Switch 1.

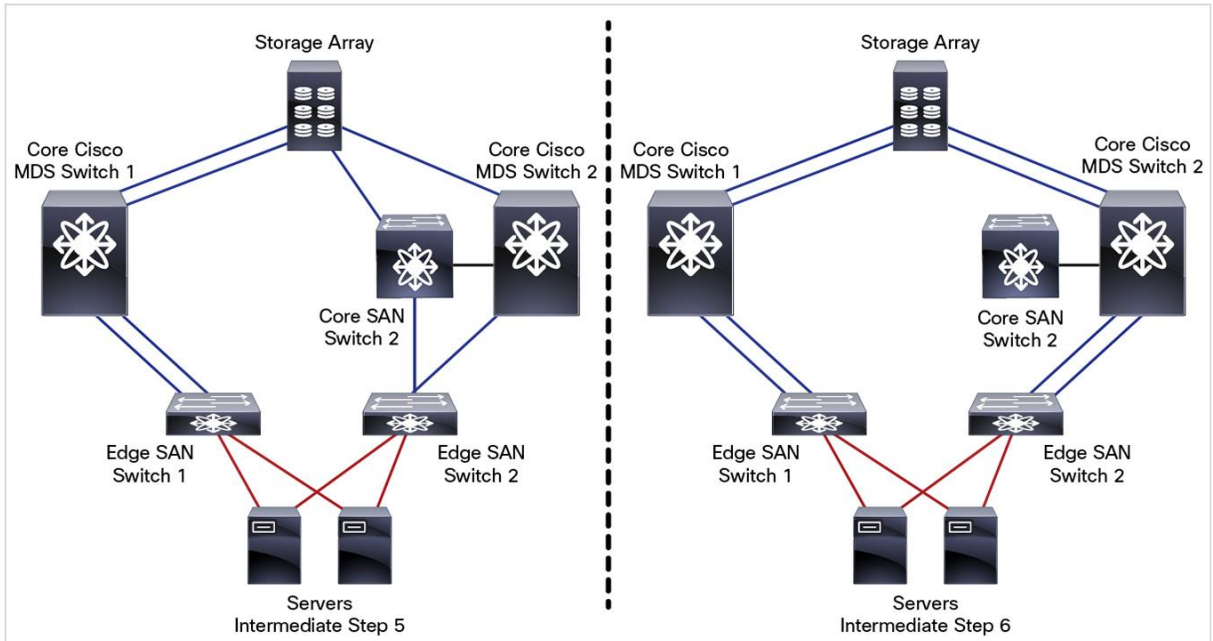
Step 4. At this point, core SAN Switch 1 is ready to retire. Take it offline (Figure 11).

Figure 11. Taking the SAN Switch Offline in a Core-Edge Network



Step 5. Repeat steps 2, 3, and 4 to replace core SAN Switch 2 with core Cisco MDS Switch 2 and then take core SAN Switch 2 offline (Figure 12).

Figure 12. Replacing SAN Switch 2 in a Core-Edge Network



The interoperate migration method has a number of advantages and disadvantages:

- Advantages
 - Cisco MDS 9000 Family switches are immediately integrated into the existing SAN.
 - Both VSANs compatible with the Cisco MDS 9000 Family and VSANs that are not compatible can be used concurrently in the same fabric with third-party switches.
 - Use Inter-VSAN Routing (IVR) can be used on Cisco MDS 9000 Family switches to reduce fabric merge risks.
 - No downtime is required to connect the Cisco MDS 9000 Family switches to the fabric. You can migrate at your own pace.
- Disadvantages
 - Multivendor switches interacting in the same fabric could result in unknown issues.
 - Firmware upgrades and downgrades on the third-party switches may be required to allow interoperability.
 - One storage port can be used by multiple hosts and applications, so the movement of one storage port may affect multiple storage ports, applications, and hosts (a situation known as a spider web).
 - If you are connecting Cisco Unified Computing System™ (Cisco UCS®) servers to Brocade switches with no VSAN, The Cisco UCS SAN profile will need to be updated with the Cisco MDS 9000 Family default VSAN. This update is nondisruptive on the Cisco UCS side.
 - Changing or modifying the interoperate mode will interrupt VSAN traffic where IVR or interoperate mode is configured. A small VSAN for interoperate mode containing only the ISLs in conjunction with IVR is preferred. After the migration is complete, remove the interoperate ISL and delete the interoperate VSAN.

- Brocade may not support some of the features of port zoning, trunking, QuickLoop, Fabric Assist, etc. in certain interoperate modes.
- As a best practice, this mode should be used temporarily during migration only.

Operate and Optimize: A Continuous Cycle

After the migration process is complete, perform the following actions to verify that the migration was successful:

- Run a Cisco Prime DCNM for SAN health report to verify that all hosts and storage devices have redundant paths.
- Check application performance levels checks and check servers for path redundancy to verify that defined and expected SLAs are being met.
- Back up new SAN configurations so that they are available in the event of a failure.
- Back up switch configurations regularly to protect against unexpected outages. You can run a script at a scheduled time to back up configurations to a Secure FTP (SFTP) server, or you can use Cisco Prime DCNM to back up configurations in the Cisco Prime DCNM database.
- Start retiring the Brocade SAN infrastructure.
- Perform multipath host verification by running a Cisco Prime DCNM host path redundancy check on all hosts after migration is complete.

After the migration is complete, you need to keep the network optimized and run it with optimal efficiency. Cisco Prime DCNM for SAN has features that can help optimize the network.

Cisco Prime DCNM topology discovery is an inherent capability of Cisco Prime DCNM for SAN to accurately depict the current topology and the device state of the connected fabric. This topology discovery also maps end storage and host devices and older switches discovered in the fabric - a very handy capability during migration.

Table 5 lists some of the Cisco Prime DCNM for SAN software features available to help you optimize and operate a new fabric.

Table 5. Cisco Prime DCNM for SAN Features

Feature	Description
Template configuration	Using the Cisco Prime DCNM web client, you can monitor Cisco MDS 9000 Family and Cisco Nexus® Family switch events, performance, and inventory, and perform minor administrative tasks.
Common LAN and SAN discovery	Discover LAN and SAN Cisco devices from a single interface.
Host dashboard	Get information about SAN and virtual hosts.
Summary dashboard	Get information about data center switches, selected SAN and LAN switches, or a group of LAN and SAN switches to see the current status, licensing details, host topology and events, and storage device topology and events.
SAN discovery and topology mapping	Cisco Prime DCNM for SAN provides extensive SAN discovery, topology mapping, and information viewing capabilities. It collects information about the fabric topology through Simple Network Management Protocol (SNMP) queries to the switches connected to it. Cisco Prime DCNM for SAN re-creates a fabric topology, presents it in a customizable map, and provides inventory and configuration information with multiple viewing options.
Inventory management	The Information pane in Cisco Prime DCNM for SAN shows inventory, configuration, and status information for all switches, links, and hosts in the fabric. Inventory management information includes vendor names and models and software and firmware versions. Select a fabric or VSAN from the Logical Domains pane and then select the Summary tab in the Information pane to get a count of the number of VSANS, switches, hosts, and storage elements in the fabric.

Feature	Description
SAN health analysis	<p>The SAN Health Advisor tool is a utility for monitoring performance and collecting statistics. You can perform the following tasks with this tool:</p> <ul style="list-style-type: none"> • Run Performance Monitor to collect I/O statistics • Collect fabric inventory (switches and other devices) • Create a graphical layout of the fabric topology • Create reports of error conditions and statistical data

Limitations, Precautions, and Verifications

While performing the migration, be sure to note the following:

- In most cases, Brocade switches require the entire switch to be taken offline after the domain ID is changed. For example, after changing domain-related configurations, you need to use **switchdisable**, which takes the switch offline.
- In most cases, enabling interoperate mode on Brocade switches running Brocade Fabric OS (FOS) Release 6.0 requires **switchdisable**, which also requires that you take switches offline.
- If you enable interoperate mode, Brocade switches will turn off some features, such as port zoning, trunking, Fabric Assist, Secure Fabric OS, and virtual flow control.
- The Brocade Virtual Fabrics feature requires an external router, but with Cisco VSAN and IVR, there is no such need.
- You should select a principal switch to assign a domain ID to all the switches in the fabric to avoid any duplicates. During the fabric merge process, if a duplicate domain ID exists, the principal switch assigns a new domain ID to one of the duplicate switches in the fabric. The recommended approach is to use a core switch as the principal switch.
- If you are migrating from an AIX or HP-UX network, preserve the existing domain ID on the Cisco MDS 9000 Family VSAN; otherwise, the process is disruptive because these hosts write the Fibre Channel ID (FCID) on the disk. Therefore, in some scenarios, a rip-and-replace migration may be the best solution.
- Modification of the domain ID may be disruptive for all the devices on that switch, but not to the remaining switches in the fabric. To avoid such disruption, the domain ID must be statically configured with the same value as the run-time domain ID.
- All Fibre Channel timers should be set to the default values before you start the actual migration to avoid any outages and conflicts later.
- If possible, always use enhanced device aliases to limit any changes to the device alias to port World Wide Name (pWWN) mapping.
- During zone-set propagation, you must address all conflicts manually.
- It is always advisable to use a seed switch - the same switch every time - preferably a core switch, during zoning configuration.
- Zoning changes cannot be activated from a Brocade switch. As a workaround, use Cisco MDS 9000 Family switches to activate zoning changes. Brocade switches cannot see IVR-enabled devices if Network Address Translation (NAT) is enabled.

The Cisco MDS 9000 Family Advantage

In the world of storage networking, Cisco MDS 9710 and Cisco MDS 9500 Series Multilayer Directors and Cisco MDS 9200 Series Multiservice Switches bring simplicity, flexibility, agility, and performance, providing the high availability and redundancy needed to access the right data at the right time from the right place, independent of the protocols being used. Cisco MDS 9000 Family switches have Fibre Channel, FCoE, and Gigabit Ethernet interfaces to support multiple protocols in the same switch. The Fibre Channel port supports Fibre Channel and FICON protocols; FCoE interfaces run FCoE traffic and Gigabit Ethernet for FCIP and Small Computer System Interface over IP (iSCSI) traffic. Here is a quick overview of the Cisco MDS 9000 Family product portfolio and some of its features.

Cisco MDS 9700 Series Multilayer Directors

Cisco MDS 9700 Series Multilayer Directors deliver superior performance, a fault-tolerant design, and multiprotocol flexibility support with nonstop operation. This platform provides 24 terabits per second (Tbps) of chassis wide throughput using 384 line-rate 2/4/8/10/16-Gbps Fibre Channel ports and 10-Gbps FCoE ports.

Cisco MDS 9500 Series Multilayer Directors

Cisco MDS 9500 Series Multilayer Directors support multiprotocol 1/2/4/8/10-Gbps Fibre Channel and 10-Gbps FCoE with 8.4 Tbps of system bandwidth. It also supports intelligent network services such as VSAN technology, access control lists (ACLs), intelligent frame processing, and fabricwide QoS. This platform also supports smart storage services, including Cisco I/O Accelerator (IOA) and Data Mobility Manager (DMM).

Cisco MDS 9250i Multiservice Fabric Switch

The Cisco MDS 9250i Multiservice Fabric Switch provides superior flexibility for SAN connectivity by delivering multiprotocol convergence and distributed fabric services along with 50 fixed ports in a compact form factor. It has 40 line-rate 16-Gbps Fibre Channel ports, eight 10-Gbps FCoE-capable Ethernet ports, and two 10-Gbps IP storage (FCIP) ports. It supports remote SAN extension with high-performance FCIP for remote replication and other disaster-recovery services along with intelligent fabric services, such as Cisco IOA and DMM.

Cisco MDS 9148 Multiservice Fabric Switch

The Cisco MDS 9148 Multiservice Fabric Switch is a fixed-form switch with 48 Fibre Channel port configurations running at line-rate 8 Gbps. The Cisco MDS 9148 provides full feature compatibility with Cisco MDS 9700 and 9500 Series Multilayer Directors and the Cisco MDS 9200 Series Multiservice Switches for transparent, end-to-end service delivery in large data center core-edge deployments.

Conclusion

Migration from a Brocade SAN to a Cisco SAN requires planning and risk analysis. The process can be relatively easy, however, with proper planning and if proper procedures are defined. Cisco MDS 9000 Family SAN switches offer many features and design functions that facilitate SAN migration between various vendors: for instance, IVR and interoperate mode. Cisco's interoperate mode with IVR helps you migrate the SAN and reduces interoperability failure domains during the migration process. The Cisco Prime DCNM for SAN GUI management tool can easily migrate Brocade SAN switch configurations. Cisco Prime DCNM for SAN has built-in tools to help with this kind of migration and to help merge the fabrics according to defined rules. Further, Cisco has always supported interoperate mode to easily integrate with competitors' products such as Brocade and McData SAN switches. Cisco can provide additional resources such as Cisco Advanced Services through your Cisco account team for more detailed analysis, evaluation, and implementation.

For More Information

- **Cisco MDS 9700 Series:**
 - [Cisco MDS 9700 Series Multilayer Directors](#)
 - [Cisco MDS 9710 Multilayer Director data sheet](#)
 - [Cisco MDS 9700 Series Supervisor-1 Module data sheet](#)
 - [Cisco MDS 9700 48-Port 16-Gbps Fibre Channel Switching Module data sheet](#)
 - [Compare models](#): Learn about the similarities and differences among the models in this product series
 - [Data sheets and product literature](#)
 - [At-a-glance documents](#)
 - [Data sheets](#)
 - [Presentations](#)
 - [White papers](#)
- **Cisco MDS 9500 Series:**
 - [Cisco MDS 9513 Multilayer Director](#)
 - [Cisco MDS 9513 Multilayer Director data sheet](#)
 - [Cisco MDS 9000 Family 8-Gbps Advanced Fibre Channel Switching Modules](#)
 - [Cisco MDS 9000 Family 8-Gbps Fibre Channel Switching Modules](#)
- **Cisco MDS 9148 Series:**
 - [Cisco MDS 9148 Multilayer Fabric Switch data sheet](#)
 - [Cisco MDS 9148 Series Fabric Switches data sheet](#)
 - [Cisco MDS 9000 NX-OS Software Release 6.2](#)
 - [Cisco MDS 9000 Family investment protection](#)
 - [Hardware and software compatibility guide with other vendors](#)
- **Cisco MDS 9000 Family hardware installation guides:**
 - [Cisco MDS 9710 Director Hardware Installation Guide](#)
 - [Cisco MDS 9500 Series Hardware Installation Guide](#)
 - [Cisco MDS 9200 Series Hardware Installation Guide](#)
 - [Cisco MDS 9148 Multilayer Fabric Switch Quick Start Guide](#)
 - [Cisco MDS 9100 Series Hardware Installation Guide](#)
 - [Cisco MDS 9500 Series Supervisor-2A Module Tech Note](#)
 - [Interoperability matrix for Cisco Nexus and Cisco MDS 9000 Family products](#)

Appendix: Technology Concepts

Virtual Fabric, LSAN, and VSAN

Brocade Virtual Fabrics: The Brocade Virtual Fabrics feature augments the proven security and fault isolation features of Brocade FOS, enabling organizations to create logical groups of separately managed devices, ports, and switches within a physical SAN. Virtual fabrics and fabric zoning have a complementary relationship. Physical ports or WWNs are assigned to virtual fabrics, and then zones are configured within the virtual fabric. Virtual fabrics may change, for example, when ports are needed or management boundaries change. When the Brocade Virtual Fabrics feature is activated, the capabilities of some features, such as administrative domains and port mirroring, are reduced. Brocade Virtual Fabrics are restrictive in their capabilities compared to Cisco VSANs, which offer greater flexibility and scalability. Virtual fabrics partition the physical infrastructure. The Brocade Virtual Fabrics feature is available on 8-Gbps products that support it, such as the Brocade DCX Backbone, Brocade DCX-4S Backbone, and Brocade 5300 and Brocade 5100 switches.

Brocade administrative domains: An administrative domain is a logical grouping of fabric elements that define the switches, ports, and devices that you can view and modify. Administrative domains partition the administration of a fabric into logical groups and allocate these groups to different user accounts so that these accounts are restricted to manage only the administrative domains assigned to them. You can configure up to 256 administrative domains in a fabric (254 user defined and 2 system defined), numbered from 0 through 255. Each administrative domain is designated by a name and a number.

Brocade logical SAN (LSAN): A Brocade LSAN consists of zones in two or more edge or backbone fabrics that contain the same devices. LSANs essentially provide selective device connectivity between fabrics without forcing you to merge those fabrics. Fibre Channel routers provide multiple mechanisms to manage interfabric device connectivity through extensions to existing switch management interfaces.

Cisco Virtual SAN (VSAN): Cisco pioneered logical fabric separation with the introduction of VSANs in the first Cisco MDS 9000 Family products, introduced in 2002. A Cisco VSAN is a logical fabric in single or multiple switches built on a physical infrastructure to form a single fabric. Every VSAN has its own services, security, and other parameters, providing isolation of any problems within that VSAN boundary only, though the VSANs share the same physical switch and hardware. VSANs can also share frame tagging for shared ISLs. VSANs also support FICON. Multiple VSANs can be defined on a single switch. To separate the VSANs, you must assign each a unique domain ID. A single VSAN can span 239 physical switches, and you can create up to 256 VSANs in a single switch.

Multiprotocol SANs can use Fibre Channel with FCoE across Cisco Nexus switching platforms (Cisco Nexus 7000 and 5000 Series Switches) to span the platforms easily. Up to 256 VSANs can be configured in a switch. Of these, one is the default VSAN (VSAN 1), and another is an isolated VSAN (VSAN 4094). User-specified VSAN IDs range from 2 to 4093.

Brocade Virtual Fabrics and LSAN configurations can be migrated to Cisco VSAN configurations to provide greater scalability, performance, and interoperability. Cisco VSANs are supported across the entire Cisco MDS 9000 Family and Cisco Nexus 7000 and 5000 Series Switches.

Table 6 provides additional information about Brocade LSANs and Virtual Fabrics and Cisco VSANs.

Table 6. Feature Interoperability of Cisco VSANs and Brocade Virtual Fabrics

Feature	Cisco VSANs	Brocade Virtual Fabrics
Virtual Fabric support	Cisco MDS 9000 Family	Brocade DCX , 5300, 5100, DCX 8510, and 6500 platforms only
Maximum number of Brocade Virtual Fabrics and Cisco VSANs per switch	All platforms: 256	Brocade DCX and DCX 8510-8
Frame tagging for shared ISLs	Yes	Yes, with qualifications [*]
FICON support	Yes	Yes, with qualifications [*]
Isolation of Virtual Fabrics	Yes	No
Default Virtual Fabrics	Yes	No
Feature limitations (after enabling Virtual Fabrics and VSAN support)	No	Yes ^{**}
Routing between Virtual Fabrics	Yes	Yes, with qualifications ^{***}

Notes:

Some VSANs are reserved for specific purpose. Cisco VSANs support default Brocade Virtual Fabrics and isolation of Brocade Virtual Fabrics as well.

^{*} Not supported with FICON, Virtual Fabrics routing, McData interoperate, Inter-Chassis Link (ICL) ports, Fibre Channel router edge switch, or Gigabit Ethernet FCIP ports, and can be used only between base switches.

^{**} The following features have limited or no support when the Virtual Fabrics feature is enabled: administrative domain (not supported), encryption (supported only in the default logical switch), port mirroring (not supported), traffic isolation zoning (not supported).

^{***} Requires the use of external ports, Small Form-Factor Pluggables (SFPs), and cables between Virtual Fabrics and the base switch. Also requires the use of line card ports (4 per connection) to route between virtual fabrics (8-Gbps of bandwidth); if more bandwidth is required, more ports must be used (4 ports for every 8-Gbps of bandwidth required).

Note: For a single virtual fabric migration from Brocade, it is easy to migrate to the default Cisco VSAN (VSAN 1). The default VSAN requires only a simple port-to-port mapping between the two fabrics, though the use of VSAN 1 for production traffic is not a best practice. If the existing fabric has multiple Brocade Virtual Fabrics, you will have to create multiple Cisco VSANs to match the different virtual fabric groups.

Inter-VSAN Routing and Virtual Fabric Routing

Cisco defines IVR to control and allow VSAN traffic within its boundaries and to set its own security and traffic policies. This approach enables easy management of the VSAN without disruption of other VSAN traffic. Devices in different VSANs communicate through a super-set zone called an IVR zone set. Only devices in the IVR zone set can see across VSAN boundaries. IVR offers an extension of the VSAN technology to provide cross-VSAN connectivity without the need to merge the routed virtual fabrics. This approach avoids propagation of irrelevant or potentially disruptive fabric events beyond the boundaries of a given VSAN. Using IVR, you can extend connectivity across VSAN boundaries and share a common storage resource among multiple VSANs, without the risk of destabilizing the fabric. IVR supports routing between all VSAN interoperate modes. IVR switches will modify the Fibre Channel headers for all communication between the end devices, including the VSAN number and source and destination FCIDs. Cisco IVR can be easily managed with less overhead. IVR is used mainly in situations in which problems arise with interoperability.

Device Aliases

Device aliases are the user-friendly names given to pWWNs. These aliases use one-to-one mappings to pWWNs and were developed to easily identify devices within the switch. They are used for purposes such as zoning and QoS. There are two types of device aliases: standard and enhanced. With standard aliases, the information is passed to the switch, which substitutes the WWN for the device alias and then passes it to the application or service being used. With enhanced mode, applications accept the device alias name in its native format, rather than expanding the alias to a pWWN. Because applications such as zone servers IVR and Dynamic Port VSAN Membership (DPVM) automatically track and enforce device alias membership changes, you have a single point of change.

Fibre Channel Aliases

Fibre Channel aliases are used to associate one or more pWWNs with a user-friendly name. They are also VSAN specific; hence, if a device is moved from one VSAN to another, a new Fibre Channel alias is needed in the new VSAN. Fibre Channel aliases are propagated through zone-set activation (assuming that the zone-set distribution is set to the full zone set). Fibre Channel aliases are propagated as part of the full database only, if propagation of the full database is allowed in that specific mode.

Table 7 summarizes the differences between Fibre Channel aliases and device aliases.

Table 7. Fibre Channel Alias and Device Alias Comparison

Fibre Channel Alias	Device Alias
Used for zoning purposes only	Multifunction (port security, IVR, zoning, etc.)
Can contain multiple pWWNs	Can have only one pWWN
Configured per VSAN	Not VSAN specific
Used mainly in multivendor environments	Used mainly if the fabric is Cisco MDS 9000 Family only
Propagated through full zone-set distribution	Propagated through Cisco Fabric Service

The primary uses of device aliases and Fibre Channel aliases are summarized here:

- IVR zoning is easier to perform in Cisco Prime DCNM using device aliases.
- Fibre Channel aliases can use only zones and zone sets. Device alias can be used with any services that use Cisco Fabric Service.
- Fibre Channel aliases interoperate with some third-party Fibre Channel switches.
- In Fibre Channel aliases, the full zone set is distributed, so they are available on all switches in the fabric.
- Device aliases are not VSAN specific. After a device alias is created, it applies to that pWWN regardless of the VSAN, whereas with a Fibre Channel alias, a different alias needs to be defined for each VSAN.
- Device aliases are automatically distributed to other Cisco switches attached to the fabric.
- Troubleshooting is easier using device aliases. After a device alias is assigned to a pWWN, any time that the pWWN is displayed, the device alias is also displayed. For example, CLI commands such as **show flogi database** and **show fcns database** will display the pWWN along with the associated device alias.

Persistent Fibre Channel IDs

Cisco MDS 9000 Family switches cache assigned FCIDs for each pWWN in volatile memory by default. In the event of any software or hardware failure, these assignments can be wiped out. The use of persistent FCIDs changes this behavior so that the assigned FCIDs and FCID-pWWN mappings are stored in nonvolatile memory. Some traditional operating systems such as HP-UX and AIX use the FCID of the SAN device mapped to the SCSI target number of the storage device to determine the logical unit number and OS storage mapping. Changing the FCID requires the server administrator to remap each LUN on each server. Persistent FCIDs can map the FCID of the storage device as the SCSI target number, so that these devices get the same FCID every time they perform a fabric login (FLOGI) to the switch. You may want to enable this feature less as a security precaution but more to achieve flexibility and availability in the event of migration. The FCID persistence feature is enabled by default on all Cisco MDS 9000 Family switches.

Domain IDs

The domain ID is part of the FCID. Every VSAN has its own unique domain ID on every interconnected switch. When the domain ID is changed, the switch itself will need to reregister with the principal switch in the fabric to verify the uniqueness of the domain ID. As a result, all devices attached to the switch will need to log into the switch again, which could be disruptive. Hence, use of a nonoverlapping static domain ID is preferred, to avoid any disruption from fabric events during migration.

Timers

Timers are extremely important for many purposes. For a Fibre Channel environment, timers can determine the time that packets are allowed to be considered in transit, and they can define various error-detection conditions, etc. The default values for these timers usually don't need to be changed, but when merging fabrics from different vendors, you must be sure that they are set identically in both fabrics. All timers should be the same across all switches because these values are exchanged by E-ports when an ISL is established. They should be left at the default settings on all Brocade switches to make sure that the transition is smooth. All Cisco switches have the same timer settings unless they have been modified manually. Timers also are important parameters for interoperate mode migration. Some valuable timer parameters are the resource allocation time-out value (R_A_TOV), error detect time-out value (E_D_TOV), distributed services time-out value (D_S_TOV), fabric stability time-out value (F_S_TOV), and receiver transmitter time-out value (R_T_TOV).

Fabric Shortest Path First and Dynamic Load Sharing

Brocade Dynamic Load Sharing (DLS) is an exchange-based routing. Cisco uses Fabric Shortest Path First (FSPF) to dynamically compute routes through a fabric by establishing the shortest and quickest path between any two switches. It supports multipath routing based on the link-state protocol and domain ID. Cisco MDS 9000 Family switches use the default src-id, dst-id, and ox-id values (or src-id and dst-id values, if these are configured) to load balance across multiple ISLs, whereas Brocade switches use their default src-id and dst-id values.

Inter-Switch Link and Inter-Chassis Link

Cisco ISLs can be configured between any Cisco MDS 9000 Family switches and line cards. Brocade ICLs use the same algorithm as Cisco Extended ISLs (EISLs), but the links can be used only between like-generation Brocade DCX switches and not with any other models or brands. The Brocade ICLs also need to go through the same application-specific integrated circuit (ASIC) in the back-end of Brocade CR-16 module, which means that the ports used by ICLs have to come from the same ASIC in the back end.

PortChannels and Trunking

Brocade uses the term “trunking,” and Cisco uses the term “PortChannel.” A PortChannel is an aggregation of Fibre Channel and FCIP links into a single logical link to provide a fault-tolerant, high-speed single link. A PortChannel can include all Fibre Channel ports or ISLs between two chassis. Cisco PortChannel technology is supported between different line cards, different ASICs, and different port groups. Cisco MDS 9000 Family switches support a maximum of 16 ISLs per PortChannel. Trunking and PortChannels are not supported between switches from two different vendors. Brocade supports a maximum of 8 ISLs, which can be combined into a single logical ISL.

However, some vendors can continue to use trunking and PortChannels between their own switches while in interoperability mode. This feature can be disabled on a per-port or per-switch basis, and continue to work as expected only if it is allowed by the interoperability mode of the vendor.

VSAN Trunking

VSAN trunking is the trunking of multiple VSANs using a single ISL or group of ISLs and becomes an EISL using VSAN header information. This feature enables a common group of ISLs to be used as a pool for connectivity between switches for multiple fabrics. It uses industry-standard virtual fabric tagging (VFT) extended headers to provide traffic segregation across common trunked ISLs. The primary benefit of VSAN trunking is that it consolidates and reduces the number of distinct ISLs required between switches. For organizations that have multiple fabrics between data centers, VSAN trunking enables a common pool of ISLs to be used, reducing the number of individual ISLs. This approach typically results in substantial cost savings through reduction in the number of dense wavelength-division multiplexing (DWDM) transponders or dark fiber pairs, allowing separate logical VSAN fabrics between sites through VSAN pruning. All VSANs do not need to go through a trunked ISL. Furthermore, individual fabrics often have very different load profiles, and grouping them together can result in higher overall throughput. VSAN trunking also allows a more controlled environment in which priority can be given to specific traffic or devices, and QoS policy can be applied to provide guaranteed bandwidth allocation for specific devices or VSANs.

Zoning

Zones help you define security and provide control over communications between multiple storage devices and user groups. Zones can be created by the administrator to increase security to help prevent data loss through corruption or spoofing. Zoning is enforced by looking at the source and destination ID fields. A zone consists of multiple zone members that can access each other. A device can belong to multiple zones, and zone size can vary. By default, all members are in the default zone unless they are part of some other active zone. Zones can be part of multiple zone sets. The default zone policy for the Cisco MDS 9000 Family switch zone set denies communication between devices. The default zone behavior of permit or deny (all nodes are isolated when not explicitly placed in a zone) may change. The default zone parameter is restricted to the switch on which it is configured, and it is not propagated to other switches. Deny is the recommended setting to help secure the environment.

Zone Set

Group of zones combined together create zone set. A single zone can be part of multiple zone sets. There are two types of zones: active and local. Active zone sets define the zone rules to enforce zoning security. This type of zone set cannot be modified and is distributed to all switches in the VSAN. There can only be one active zone set. A local zone set contains the complete the zone-set database for that switch. This zone set can then be activated to become the active zone set. A VSAN can have multiple local zone sets.

Zone Membership

Zoning can be enforced in two ways: through hard zoning and soft zoning. Hard zoning is enforced with the hardware of each switch for each frame. As soon as a frame reaches the switch, the source and destination IDs are compared with ACL combinations to allow or deny the frame. Hard zoning can be applied to all forms of zoning. Hard zoning is also more secure than soft zoning because it is applied to every frame to help prevent unauthorized access.

Soft zoning is applied only for the duration of interaction between the name server and the end device. If an end device knows the FCID of a device outside its zone, it can access it easily.

A switch can be preconfigured with a set of zones, with zone membership based on the port to which a device is connected (hard zoning). If other proprietary zoning methods (physical port numbers) are eliminated, zones may be limited to the pWWN. Not all vendors can support the same number of zones. Determine the lowest common denominator with Brocade and limit the fabric to the values in Table 8.

Table 8. Zone Types in Interoperability Mode

Zone Type	Cisco MDS 9000 Family Compatible Interoperability Modes
pWWN	All
FCID	Noninteroperability mode only
Fabric pWWN	Noninteroperability mode only
Fibre Channel alias	All
Domain and port	Traditional switch interoperability modes 2, 3, and 4
Symbolic node name	Noninteroperability mode only
Interface and switch WWN	Noninteroperability mode only

Zone-Set Database and Its Propagation

A zone-set database and active zone set are two separate entities. A zone-set database is a local database on each switch that contains all the zone sets, zones, and zone member information, whereas each VSAN in the fabric has a single active zone-set entity derived from the zone-set database of the local switch. This active zone set is distributed to all the switches in the fabric upon activation and remains consistent across all the switches in the fabric, whereas the zone-set database is a local entity and does not need to be homogeneous in the fabric. The zone-set database is not identical on all the switches, which could lead to problems. Multiple switches can be used to configure zoning information at different times, but upon activation only the local switch zone-set database is enforced by the fabric.

This behavior could be disruptive if proper attention is not paid to the zoning methodology, and for that reason some switch vendors recommend use of a seed switch for all zoning configuration. Use of a seed switch can definitely alleviate this problem; however, Cisco MDS 9000 Family switches also provide two commands:

- The **EXEC level zoneset distribute** command distributes the zone-set database of that switch to the entire fabric upon activation.
- The **config level zoneset distribute** command distributes the zone-set database of that switch upon zone-set activation.

Use of the **config level zoneset distribute** command on all switches in the fabric is highly recommended. After this command is activated, all the switches in the fabric will have a consistent zone-set database in the active zone-set entity.

Note: The Cisco Prime DCNM for SAN GUI tool always uses the principal switch as the seed switch for all zoning configuration.




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.

 Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)