Cisco MDS 9000 Series System Management Configuration Guide, Release 9.x

First Published: 2021-08-16

Last Modified: 2023-08-18

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

CHAPTER 5

CHAPTER 7

CHAPTER 16

Preface

This preface describes the audience, organization of, and conventions used in the Cisco MDS 9000 Series Configuration Guides. It also provides information on how to obtain related documentation, and contains the following chapters:

Audience

To use this installation guide, you need to be familiar with electronic circuitry and wiring practices, and preferably be an electronic or electromechanical technician.

Document Conventions

This document uses the following conventions:

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Warnings use the following conventions:

Warning This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device. Statement 1071.

Related Documentation

The documentation set for the Cisco MDS 9000 Series Switches includes the following documents.

Release Notes

http://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/products-release-notes-list.html

Regulatory Compliance and Safety Information

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/hw/regulatory/compliance/RCSI.html

Compatibility Information

http://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/products-device-support-tables-list.html

Installation and Upgrade

http://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/products-installation-guides-list.html

Configuration

http://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/products-installation-and-configuration-guides-list.html

CLI

http://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/products-command-reference-list.html

Troubleshooting and Reference

http://www.cisco.com/c/en/us/support/storage-networking/mds-9000-nx-os-san-os-software/tsd-products-support-troubleshoot-and-alerts.html

To find a document online, use the Cisco MDS NX-OS Documentation Locator at:

http://www.cisco.com/c/en/us/td/docs/storage/san_switches/mds9000/roadmaps/doclocater.html

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business impact you’re looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco Marketplace.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

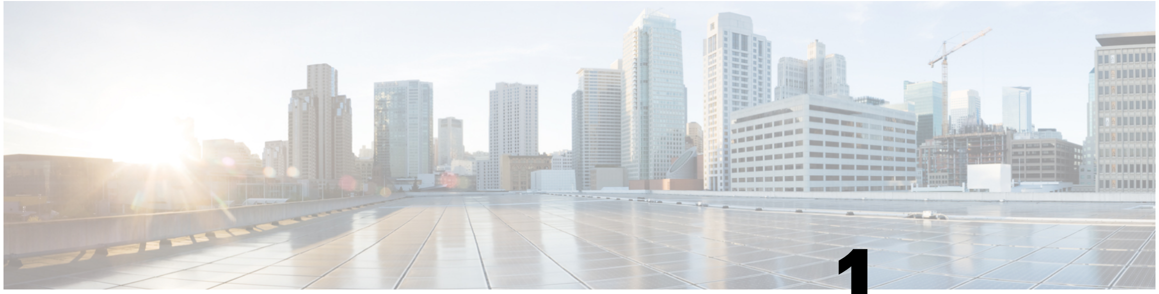

CHAPTER 1

New and Changed Information

Change Summary

Feature Name	Description	Release	Where Documented
Secure Syslog	The Secure Syslog Servers feature allows you to securely log system messages to a configured remote logging server using TLS.	9.2(1)	Configuring System Message Logging, on page 27

CHAPTER 2

System Management Overview

You can use the system management features to monitor and manage a switch using Cisco MDS NX-OS software. These features include Call Home, SNMP, RMON, SPAN, and the Embedded Event Manager (EEM).

Cisco Fabric Services

The Cisco MDS NX-OS software uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database distribution and to promote device flexibility. CFS simplifies SAN provisioning by automatically distributing configuration information to all switches in a fabric.

For information on configuring CFS, see Using the CFS Infrastructure, on page 7.

System Messages

System messages are monitored remotely by accessing the switch through Telnet, SSH, or the console port, or by viewing the logs on a system message logging server. Log messages are not saved across system reboots.

For information about configuring system messages, see Configuring System Message Logging, on page 27.

Call Home

Call Home provides email-based notification of critical system events. A versatile range of message formats are available for optimal compatibility with pager services, standard email, or XML-based automated parsing applications. Common uses of this feature may include direct paging of a network support engineer, email notification to a Network Operations Center, and utilization of Cisco Smart Call Home services for direct case generation with the Technical Assistance Center.

For information about configuring Call Home, see Configuring Call Home, on page 49.

Scheduler

The Cisco MDS command scheduler feature helps you schedule configuration and maintenance jobs in any switch in the Cisco MDS 9000 Family switches.You can use this feature to schedule jobs on a one-time basis or periodically. The Cisco NX-OS command scheduler provides a facility to schedule a job (set of CLI commands) or multiple jobs at a specified time in the future. The jobs can be executed once at a specified time in the future or at periodic intervals.

For information on configuring the Cisco MDS command scheduler feature, see Scheduling Maintenance Jobs, on page 113.

System Processes and Logs

The health of a switch can be monitored by various system processes and logs. The Online Health Management System (system health) is a hardware fault detection and recovery feature. This Health Management System ensures the general health of switching, services, and supervisor modules in any switch in the Cisco MDS 9000 Family.

For information on monitoring the health of the switch, see System Status Monitoring, on page 127.

Embedded Event Manager

Embedded Event Manager (EEM) monitors events that occur on your device and takes action to recover or troubleshoot these events, based on your configuration. EEM consists of three major components:

- Event statements—Events to monitor from another Cisco NX-OS component that may require some action, workaround, or notification.
- Action statements —An action that EEM can take, such as sending an email or disabling an interface, to recover from an event.
- Policies—An event paired with one or more actions to troubleshoot or recover from the event.

For information on configuring EEM, see About Embedded Event Manager, on page 159.

SNMP

Simple Network Management Protocol (SNMP) is an application layer protocol that facilitates the exchange of management information between network devices. In all Cisco MDS 9000 Family switches, three SNMP versions are available: SNMPv1, SNMPv2c, and SNMPv3. The CLI and SNMP use common roles in all switches in the Cisco MDS 9000 Family. You can use SNMP to modify a role that was created using the CLI and vice versa.

Users, passwords, and roles for all CLI and SNMP users are the same. A user configured through the CLI can access the switch using SNMP (for example, the DCNM-SAN or the Device Manager) and vice versa.

For information on configuring SNMP, see Configuring SNMP, on page 235.

RMON

RMON is an Internet Engineering Task Force (IETF) standard monitoring specification that allows various network agents and console systems to exchange network monitoring data. You can use the RMON alarms and events to monitor Cisco MDS 9000 Family switches running the Cisco SAN-OS Release 2.0(1b) or later or Cisco Release NX-OS 4.1(3) or later software.

For information on configuring RMON, see Configuring RMON, on page 179.

Pathtrace

The Pathtrace feature builds on the Traceroute feature to provide information about interfaces, such as ingress and egress interface names and the number of transmitted and received frames and errors, at each hop in the path between 2 devices in a fabric. Pathtrace provides an end-to-end view of the shortest path without the need to connect to individual switches and check the Fabric Shortest Path First (FSPF) topology hop by hop.

For information on using the Pathtrace feature, see Using Pathtrace, on page 217.

Domain Parameters

The Fibre Channel domain (fcdomain) feature performs principal switch selection, domain ID distribution, FC ID allocation, and fabric reconfiguration functions as described in the FC-SW-2 standards. The domains are configured on a per-VSAN basis. If you do not configure a domain ID, the local switch uses a random ID.

For information on configuring the Fibre Channel domain feature, see Configuring Domain Parameters, on page 263.

SPAN

The Switched Port Analyzer (SPAN) feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic through a Fibre Channel interface. Traffic through any Fibre Channel interface can be replicated to a special port called the SPAN destination port (SD port). Any Fibre Channel port in a switch can be configured as an SD port. Once an interface is in SD port mode, it cannot be used for normal data traffic. You can attach a Fibre Channel analyzer to the SD port to monitor SPAN traffic.

For information on SPAN feature, see Monitoring Network Traffic Using SPAN, on page 291.

Fabric Configuration Server

The Fabric Configuration Server (FCS) provides discovery of topology attributes and maintains a repository of configuration information of fabric elements. A management application is usually connected to the FCS on the switch through an N port. In the Cisco MDS 9000 Family switch environment, multiple VSANs constitute a fabric, where one instance of the FCS is present per VSAN.

For information on configuring FCS, see Configuring Fabric Configuration Server, on page 329.

Uniform Timestamps

While debugging issues on a switch, it is critical to arrange logs in chronological order to understand the sequence of events that resulted in the issue. MDS logs use different time formats and it is tedious to merge and sort them to understand the timeline of events. The Uniform Timestamps feature allows onboard syslogs, accounting logs, and logs of various MDS software components to use a uniform timestamp format. This allows quick merging and sorting of multiple logs together to build a complex timeline onboard the switch. Logs can also be exported from the switch (by using the **show tech-support** command and so on) and processed easily by hand, scripts, or data-mining applications.

This feature enables RFC 5424 format timestamps. This format is supported by many other devices and vendors so that logs can also be merged with other products to build an end to end timeline through a fabric. Ensure that clocks on all the devices are synchronized before attempting this.

This feature does not change the format of syslogs that are exported to an external syslog server via the syslog protocol.

For more information, see the **system timestamp format** command in the Cisco MDS 9000 Series Command Reference, Release 8.x.

Configuring Uniform Timestamps

To enable RFC 5424 compliant timestamping in logs, perform the step below:

Procedure

Step 1 Enter global configuration mode:

switch# **configure terminal**

Step 2 Enable RFC 5424 compliant timestamping in logs:

switch# **system timestamp format rfc5424**

CHAPTER 3

Using the CFS Infrastructure

Cisco Fabric Services (CFS) provides a common infrastructure for automatic configuration synchronization in the fabric. It provides the transport function as well as a rich set of common services to the applications. CFS has the ability to discover CFS-capable switches in the fabric and discover application capabilities in all CFS-capable switches.

Information About CFS

The Cisco MDS NX-OS software uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database distribution and to foster device flexibility. It simplifies SAN provisioning by automatically distributing configuration information to all switches in a fabric.

Several Cisco MDS NX-OS applications use the CFS infrastructure to maintain and distribute the contents of a particular application's database.

Many features in the Cisco MDS switches require configuration synchronization in all switches in the fabric. Maintaining configuration synchronization across a fabric is important to maintain fabric consistency. In the absence of a common infrastructure, such synchronization is achieved through manual configuration at each switch in the fabric. This process is tedious and error prone.

Cisco MDS NX-OS Features Using CFS

The following Cisco NX-OS features use the CFS infrastructure:

- N Port Virtualization
- FlexAttach Virtual pWWN
- NTP
- Dynamic Port VSAN Membership
- Distributed Device Alias Services

- IVR topology
- SAN device virtualization
- TACACS+ and RADIUS
- User and administrator roles
- Port security
- iSNS
- Call Home
- Syslog
- fctimer
- SCSI flow services
- Saved startup configurations using the Fabric Startup Configuration Manager (FSCM)
- Allowed domain ID lists
- RSCN timer
- iSLB

CFS Features

CFS has the following features:

- Peer-to-peer protocol with no client-server relationship at the CFS layer.
- Three scopes of distribution.
  - Logical scope—The distribution occurs within the scope of a VSAN.
  - Physical scope—The distribution spans the entire physical topology.
  - Over a selected set of VSANs—Some applications, such as Inter-VSAN Routing (IVR), require configuration distribution over some specific VSANs. These applications can specify to CFS the set of VSANs over which to restrict the distribution.
- Three modes of distribution.
  - Coordinated distributions—Only one distribution is allowed in the fabric at any given time.
  - Uncoordinated distributions—Multiple parallel distributions are allowed in the fabric except when a coordinated distribution is in progress.
  - Unrestricted uncoordinated distributions—Multiple parallel distributions are allowed in the fabric in the presence of an existing coordinated distribution. Unrestricted uncoordinated distributions are allowed to run in parallel with all other types of distributions.
- Supports a merge protocol that facilitates the merge of application configuration during a fabric merge event (when two independent fabrics merge).

Enabling CFS for an Application

All CFS-based applications provide an option to enable or disable the distribution capabilities. Features that existed prior to Cisco SAN-OS Release 2.0(1b) have the distribution capability disabled by default and must have distribution capabilities enabled explicitly.

Applications introduced in Cisco SAN-OS Release 2.0(1b) or later, or MDS NX-OS Release 4.1(1) or later have the distribution enabled by default.

The application configuration is not distributed by CFS unless distribution is explicitly enabled for that application.

CFS Protocol

The CFS functionality is independent of the lower layer transport. Currently, in Cisco MDS switches, the CFS protocol layer resides on top of the Fiber Channel 2 (FC2) layer and is peer-to-peer with no client-server relationship. CFS uses the FC2 transport services to send information to other switches. CFS uses a proprietary SW_ILS (0x77434653) protocol for all CFS packets. CFS packets are sent to or from the switch domain controller addresses.

CFS can also use IP to send information to other switches.

Applications that use CFS are completely unaware of the lower layer transport.

CFS Distribution Scopes

Different applications on the Cisco MDS 9000 Family switches need to distribute the configuration at various levels:

- VSAN level (logical scope)

  Applications that operate within the scope of a VSAN have the configuration distribution restricted to the VSAN. An example application is port security where the configuration database is applicable only within a VSAN.

- Physical topology level (physical scope)

  Applications might need to distribute the configuration to the entire physical topology spanning several VSANs. Such applications include NTP and DPVM (WWN-based VSAN), which are independent of VSANs.

- Between selected switches

  Applications might only operate between selected switches in the fabric. An example application is SCSI flow services, which operates between two switches.

CFS Distribution Modes

CFS supports different distribution modes to support different application requirements: coordinated and uncoordinated distributions. Both modes are mutually exclusive. Only one mode is allowed at any given time.

Uncoordinated Distribution

Uncoordinated distributions are used to distribute information that is not expected to conflict with that from a peer. An example is local device registrations such as iSNS. Parallel uncoordinated distributions are allowed for an application.

Coordinated Distribution

Coordinated distributions can have only one application distribution at a given time. CFS uses locks to enforce this. A coordinated distribution is not allowed to start if locks are taken for the application anywhere in the fabric. A coordinated distribution consists of three stages:

1. A fabric lock is acquired.
2. The configuration is distributed and committed.
3. The fabric lock is released.

Coordinated distribution has two variants:

- CFS driven —The stages are executed by CFS in response to an application request without intervention from the application.
- Application driven—The stages are under the complete control of the application.

Coordinated distributions are used to distribute information that can be manipulated and distributed from multiple switches, for example, the port security configuration.

Unrestricted Uncoordinated Distributions

Unrestricted uncoordinated distributions allow multiple parallel distributions in the fabric in the presence of an existing coordinated distribution. Unrestricted uncoordinated distributions are allowed to run in parallel with all other types of distributions.

CFS Connectivity in a Mixed Fabric

CFS is an infrastructure component that also runs on the Cisco Nexus 5000 Series switches and the Cisco MDS 9000 switches. A mixed fabric of different platforms (such as the Cisco Nexus 7000 Series, Cisco Nexus 5000 Series, and Cisco MDS 9000 switches) can interact with each other.

Using CFSoIP and CFSoFC, the respective CFS clients can also talk to their instances running on the other platforms. Within a defined domain and distribution scope, CFS can distribute the client's data and configuration to its peers running on other platforms.

All three platforms support both CFSoIP and CFSoFC. However, the Cisco Nexus 7000 Series and Cisco Nexus 5000 Series switches require an FC or FCoE plugin and corresponding configuration in order for CFSoFC to operate. Both options are available by default on the Cisco MDS 9000 switches.

Note Some applications are not compatible with their instances running on different platforms. Therefore, Cisco recommends that you carefully read the client guidelines for CFS distribution before committing the configuration.

For more information on CFS for the Cisco Nexus 5000 Series and Cisco MDS 9000 switches, see the Cisco Nexus 5000 Series NX-OS System Management Configuration Guide and the Cisco MDS 9000 Family NX-OS System Management Configuration Guide, respectively.

Locking the Fabric

When you configure (first time configuration) a Cisco NX-OS feature (or application) that uses the CFS infrastructure, that feature starts a CFS session and locks the fabric. When a fabric is locked, the Cisco NX-OS software does not allow any configuration changes from a switch to this Cisco NX-OS feature, other than the switch holding the lock, and issues a message to inform the user about the locked status. The configuration changes are held in a pending database by that application.

If you start a CFS session that requires a fabric lock but forget to end the session, an administrator can clear the session. If you lock a fabric at any time, your user name is remembered across restarts and switchovers. If another user (on the same machine) tries to perform configuration tasks, that user's attempts are rejected.

For information on verifying CFS lock status, refer to Verifying CFS Lock Status, on page 24.

Committing Changes

A commit operation saves the pending database for all application peers and releases the lock for all switches.

In general, the commit function does not start a session; only a lock function starts a session. However, an empty commit is allowed if configuration changes are not previously made. In this case, a commit operation results in a session that acquires locks and distributes the current database.

When you commit configuration changes to a feature using the CFS infrastructure, you receive a notification about one of the following responses:

- One or more external switches report a successful status—The application applies the changes locally and releases the fabric lock.
- None of the external switches report a successful state—The application considers this state a failure and does not apply the changes to any switch in the fabric. The fabric lock is not released.

Note Once the **feature commit** is done the running configuration has been modified on all switches participating in the feature's distribution. You can then use the **copy running-config startup-config fabric** command to save the running-config to the startup-config on all the switches in the fabric.

CFS Merge Support

An application keeps the configuration synchronized in a fabric through CFS. Two such fabrics might merge as a result of an ISL coming up between them. These two fabrics could have two different sets of configuration information that need to be reconciled in the event of a merge. CFS provides notification each time an application peer comes online. If a fabric with M application peers merges with another fabric with N application peers and if an application triggers a merge action on every such notification, a link-up event results in M*N merges in the fabric.

CFS supports a protocol that reduces the number of merges required to one by handling the complexity of the merge at the CFS layer. This protocol runs per application per scope. The protocol involves selecting one switch in a fabric as the merge manager for that fabric. The other switches do not play any role in the merge process.

During a merge, the merge manager in the two fabrics exchange their configuration databases with each other. The application on one of them merges the information, decides if the merge is successful, and informs all switches in the combined fabric of the status of the merge.

In case of a successful merge, the merged database is distributed to all switches in the combined fabric and the entire new fabric remains in a consistent state.

CFS Distribution over IP

You can configure CFS to distribute information over IP for networks containing switches that are not reachable over Fibre Channel. CFS distribution over IP supports the following features:

- Physical distribution over an entirely IP network.
- Physical distribution over a hybrid Fibre Channel and IP network with the distribution reaching all switches that are reachable over either Fibre Channel or IP.

Note The switch attempts to distribute information over Fibre Channel first and then over the IP network if the first attempt over Fibre Channel fails. CFS does not send duplicate messages if distribution over both IP and Fibre Channel is enabled.

- Distribution over IP version 4 (IPv4) or IP version 6 (IPv6).

Note CFS cannot distribute over both IPv4 and IPv6 from the same switch.

- Keepalive mechanism to detect network topology changes using a configurable multicast address.
- Compatibility with Cisco MDS SAN-OS Release 2.x.
- Distribution for logical scope applications is not supported because the VSAN implementation is limited to Fibre Channel.

Figure 1: Network Example 1 with Fibre Channel and IP Connections, on page 12 shows a network with both Fibre Channel and IP connections. Node A forwards an event to node B over Fibre Channel. Node B forwards the event node C and node D using unicast IP. Node C forwards the event to node E using Fibre Channel.

Figure 1: Network Example 1 with Fibre Channel and IP Connections

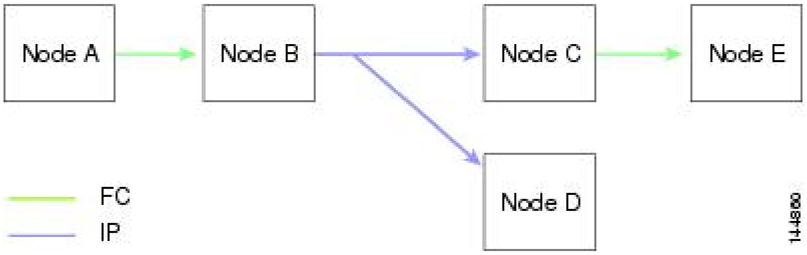

Figure 2: Network Example 2 with Fibre Channel and IP Connections, on page 12 is the same as Figure 1: Network Example 1 with Fibre Channel and IP Connections, on page 12 except that node D and node E are connected using Fibre Channel. All processes is the same in this example because node B has node C and node D the distribution list for IP. Node C does not forward to node D because node D is already in the distribution list from node B.

Figure 2: Network Example 2 with Fibre Channel and IP Connections

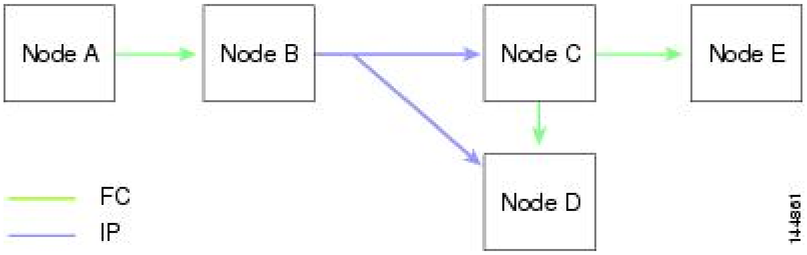

Figure 3: Network Example 3 with Fibre Channel and IP Connections, on page 13 is the same as Figure 2: Network Example 2 with Fibre Channel and IP Connections, on page 12 except that node D and node E are

connected using IP. Both node C and node D forward the event to E because the node E is not in the distribution list from node B.

Figure 3: Network Example 3 with Fibre Channel and IP Connections

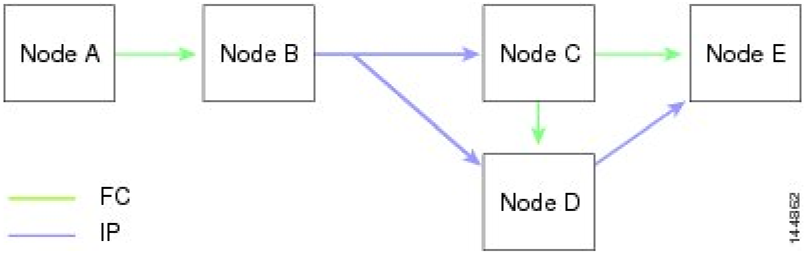

Static IP Peers for CFS

CFS over IP can also be used with static IP peers. In this case, dynamic discovery over IP multicast is disabled and CFS distribution is done only on the peers configured statically.

CFS uses the list of configured IP addresses to communicate with each peer and learn the peer switch WWN. After learning the peer switch WWN, CFS marks the switch as CFS-capable and triggers application-level merging and database distribution.

Multicast forwarding is disabled by default in some devices. For example, the IBM Blade chassis has multicast forwarding disabled, especially on external Ethernet ports, and there is no method to enable it. N port virtualization devices use only IP as the transport medium and do not have ISL connectivity or a Fibre Channel domain. Such devices may benefit from using static IP peers for CFS.

The following MDS 9000 features require static IP peer configuration for CFS over IP distribution:

- N port virtualization devices have IP as the communication channel because NPV switches do not have FC domain. NPV devices use CFS over IP as the transport medium.
- FlexAttach virtual pWWN distribution on CFS region 201 that links only the NPV-enabled switches.

About CFS Regions

A CFS region is a user-defined subset of switches for a given feature or application in its physical distribution scope.When a SAN is spanned across a vast geography, you may need to localize or restrict the distribution of certain profiles among a set of switches based on their physical proximity. Before MDS SAN-OS Release 3.2.(1) the distribution scope of an application within a SAN was spanned across the entire physical fabric without the ability to confine or limit the distribution to a required set of switches in the fabric. CFS regions enables you to overcome this limitation by allowing you to create CFS regions, that is, multiple islands of distribution within the fabric, for a given CFS feature or application. CFS regions are designed to restrict the distribution of a feature's configuration to a specific set or grouping of switches in a fabric.

Note You can only configure a CFS region on physical switches in a SAN. You cannot configure a CFS region in a VSAN.

Example CFS Scenario: Call Home is an application that triggers alerts to Network Administrators when a situation arises or something abnormal occurs. When the fabric covers many geographies and with multiple Network Administrators who are each responsible for a subset of switches in the fabric, the Call Home application sends alerts to all Network Administrators regardless of their location. For the Call Home application to send message alerts selectively to Network Administrators, the physical scope of the application has to be fine tuned or narrowed down, which is achieved by implementing CFS regions.

CFS regions are identified by numbers ranging from 0 through 200. Region 0 is reserved as the default region, and contains every switch in the fabric. You can configure regions from 1 through 200. The default region maintains backward compatibility. If there are switches on the same fabric running releases of SAN-OS before Release 3.2(1), only features in Region 0 are supported when those switches are synchronized. Features from other regions are ignored when those switches are synchronized.

If the feature is moved, that is, assigned to a new region, its scope is restricted to that region; it ignores all other regions for distribution or merging purposes. The assignment of the region to a feature has precedence in distribution over its initial physical scope.

You can configure a CFS region to distribute configurations for multiple features. However, on a given switch, you can configure only one CFS region at a time to distribute the configuration for a given feature. Once you assign a feature to a CFS region, its configuration cannot be distributed within another CFS region.

Guidelines and Limitations

All switches in the fabric must be CFS capable. A Cisco MDS 9000 Family switch is CFS capable if it is running Cisco SAN-OS Release 2.0(1b) or later, or MDS NX-OS Release 4.1(1) or later. Switches that are not CFS capable do not receive distributions and result in part of the fabric not receiving the intended distribution.

CFS has the following guidelines and limitations:

- Implicit CFS usage—The first time you issue a CFS task for a CFS-enabled application, the configuration modification process begins and the application locks the fabric.
- Pending database—The pending database is a temporary buffer to hold uncommitted information. The uncommitted changes are not applied immediately to ensure that the database is synchronized with the database in the other switches in the fabric. When you commit the changes, the pending database overwrites the configuration database (also known as the active database or the effective database).
- CFS distribution enabled or disabled on a per-application basis—The default (enable or disable) for CFS distribution state differs between applications. If CFS distribution is disabled for an application, then that application does not distribute any configuration nor does it accept a distribution from other switches in the fabric.
- Explicit CFS commit—Most applications require an explicit commit operation to copy the changes in the temporary buffer to the application database, to distribute the new database to the fabric, and to release the fabric lock. The changes in the temporary buffer are not applied if you do not perform the commit operation.

Default Settings

Table 1: Default CFS Parameters , on page 14 lists the default settings for CFS configurations.

Table 1: Default CFS Parameters

Parameters	Default
CFS distribution on the switch	Enabled.

Parameters	Default
Database changes	Implicitly enabled with the first configuration change.
Application distribution	Differs based on application.
Commit	Explicit configuration is required.
CFS over IP	Disabled.
Static IP Peers for CFS	Disabled.
IPv4 multicast address	239.255.70.83
IPv6 multicast address	ff15:efff:4653

Configuring CFS

This section describes the configuration process.

Disabling CFS Distribution on a Switch

By default, CFS distribution is enabled. Applications can distribute data and configuration information to all CFS-capable switches in the fabric where the applications exist. This is the normal mode of operation.

You can globally disable CFS on a switch, including CFS over IP to isolate the applications using CFS from fabric-wide distributions while maintaining physical connectivity.

Note When CFS is globally disabled on a switch, CFS operations are restricted to the switch and all CFS commands continue to function as if the switch were physically isolated.

To globally disable or enable CFS distribution on a switch, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **no cfs distribute**

Globally disables CFS distribution for all applications on the switch, including CFS over IP.

Step 3 switch(config)# **cfs distribute**

Enables (default) CFS distribution on the switch.

Committing Changes

You can commit changes for a specified feature by entering the **commit** command for that feature.

Discarding Changes

If you discard configuration changes, the application flushes the pending database and releases locks in the fabric. Both the abort and commit functions are only supported from the switch from which the fabric lock is acquired.

You can discard changes for a specified feature by using the **abort** command for that feature.

Saving the Configuration

Configuration changes that have not been applied yet (still in the pending database) are not shown in the running configuration. The configuration changes in the pending database overwrite the configuration in the effective database when you commit the changes.

⚠ **Caution** If you do not commit the changes, they are not saved to the running configuration.

The CISCO-CFS-MIB contains SNMP configuration information for any CFS-related functions. Refer to the *Cisco MDS 9000 Family MIB Quick Reference* for more information on this MIB.

Clearing a Locked Session

You can clear locks held by an application from any switch in the fabric. This option is provided to rescue you from situations where locks are acquired and not released.

The clear the CFS locks, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **dpvm abort**

Terminates the configuration from the switch where the configuration lock was acquired previously. This method clears the CFS locks in the entire fabric.

Clears the CFS locks for the DPVM application in the entire fabric.

Step 3 switch(config)# **clear dpvm session**

Clears the sessions from any switch in the fabric.

Clears the CFS locks for the DPVM application.

Enabling CFS over IP

Enabling or Disabling CFS over IPv4

To enable or disable CFS over IPv4, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **cfs ipv4 distribute**

Globally enables CFS over IPv4 for all applications on the switch.

Step 3 switch(config)# **no cfs ipv4 distribute**

```
This will prevent CFS from distributing over IPv4 network.
Are you sure? (y/n) [n] y
```

Disables (default) CFS over IPv4 on the switch.

Enabling or Disabling CFS Over IPv6

To enable or disable CFS over IPv6, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **cfs ipv6 distribute**

Globally enables CFS over IPv6 for all applications on the switch.

Step 3 switch(config)# **no cfs ipv6 distribute**

Disables (default) CFS over IPv6 on the switch.

Configuring IP Multicast Address for CFS over IP

All CFS over IP enabled switches with similar multicast addresses form one CFS over IP fabric. CFS protocol specific distributions, such as the keepalive mechanism for detecting network topology changes, use the IP multicast address to send and receive information.

Note CFS distributions for application data use directed unicast.

You can configure a CFS over IP multicast address value for either IPv4 or IPv6. The default IPv4 multicast address is 239.255.70.83 and the default IPv6 multicast address is ff15:efff:4653.

Configuring IP Multicast Address for CFS over IPv4

To configure an IP multicast address for CFS over IPv4, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **cfs ipv4 mcast-address 239.255.1.1**

```
Distribution over this IP type will be affected
Change multicast address for CFS-IP ?
```

Configures the IPv4 multicast address for CFS distribution over IPv4. The ranges of valid IPv4 addresses are 239.255.0.0 through 239.255.255.255 and 239.192/16 through 239.251/16.

Step 3 switch(config)# **no cfs ipv4 mcast-address 239.255.1.1**

Example:

```
Distribution over this IP type will be affected
Change multicast address for CFS-IP ?Are you sure? (y/n)  [n] y
```

Reverts to the default IPv4 multicast address for CFS distribution over IPv4. The default IPv4 multicast address for CFS is 239.255.70.83.

Configuring IP Multicast Address for CFS over IPv6

To configure an IP multicast address for CFS over IPv6, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **cfs ipv6 mcast-address ff15::e244:4754**

```
Distribution over this IP type will be affected
Change multicast address for CFS-IP ?
```

```
Are you sure? (y/n)  [n] y
```

Configures the IPv6 multicast address for CFS distribution over IPv6. The range of valid IPv6 addresses is ff15::/16 (ff15::0000:0000 through ff15::ffff:ffff) and ff18::/16 (ff18::0000:0000 through ff18::ffff:ffff).

Step 3 switch(config)# **no cfs ipv6 mcast-address ff15::e244:4754**

```
Distribution over this IP type will be affected
Change multicast address for CFS-IP ?
Are you sure? (y/n)  [n] y
```

Reverts to the default IPv6 multicast address for CFS distribution over IPv6. The default IPv6 multicast address for CFS over IP is ff15::efff:4653.

Configuring Static IP Peers for CFS

To configure a static IP peer address for CFS over IP, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **cfs static-peers**

```
WARNING: This mode will stop dynamic discovery and rely only on the static peers. For this
 mode to be in effect, at least one static peer will need to be configured.
Do you wish to continue? (y/n)  [n] y
```

switch(config-cfs-static)#

Enters CFS static peers configuration mode and disables dynamic discovery of peers using multicast forwarding. For this to take effect, at least one static peer needs to be configured in Step 3.

Step 3 switch(config)# **no cfs static-peers**

```
WARNING: This will remove all existing peers and start dynamic discovery.
Do you wish to continue? (y/n)  [n] y
```

Disables CFS static peer discovery and enables dynamic peer discovery using multicast forwarding on all switches.

Step 4 switch(config-cfs-static)# **ip address 1.2.3.4**

switch(config-cfs-static)#**ip address 1.2.3.5**

switch(config-cfs-static)# **end**

switch#

Adds the IP address to the static peers list and marks the switch as CFS-capable. To display the static IP peers list, use the **show cfs static peers** command.

Step 5 switch(config-cfs-static)# **no ip address 1.2.3.3**

switch(config-cfs-static)# **end**

Removes the IP address from the static peers list and moves the switch to dynamic peer discovery using multicast forwarding.

Step 6 switch# **show cfs static peers**

Displays the IP address, WWN, and the status of CFS static peer request:

- Discovery Inprogress
- Local
- Reachable
- Unreachable
- Local IP not present
- Rediscovery and distribution disabled

Note The IP address and WWN must be configured on the local switch. If CFS does not receive the local switch information, then CFS cannot start any discovery for peer switches.

Configuring CFS Regions

Creating CFS Regions

To create a CFS region, perform this task:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **cfs region 4**

Creates a region, for example, number 4.

Assigning Applications to CFS Regions

To assign an application on a switch to a region, perform this task:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **cfs region 4**

Creates a region, for example, number 4.

Step 3 switch(config-cfs-region)# **ntp**

switch(config-cfs-region)# **callhome**

Adds application(s).

Moving an Application to a Different CFS Region

You can move an application to a different CFS region, for example from Region 1 (originating region) with NTP and Call Home applications to Region 2 (target region).

To move an application, perform this task:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **cfs region 2**

Enters Region 2.

Step 3 switch(config-cfs-region)# **ntp**

switch(config-cfs-region)# **callhome**

Indicates application(s) to be moved into Region 2 that originally belong to Region 1. For example, here, the NTP and Call Home applications are moved to Region 2.

Note If you try adding an application to the same region more than once, you see the error message, “Application already present in the same region.”

Removing an Application from a Region

Removing an application from a region is the same as moving the application back to the default region or to Region 0, that is, bringing the entire fabric into the scope of distribution for the application.

To remove applications from Region 1, perform this task:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **cfs region 1**

Enters Region 1.

Step 3 switch(config-cfs-region)# **no ntp**

switch(config-cfs-region)# **no callhome**

Removes application(s) that belong to Region 1, which you want to move.

Deleting CFS Regions

Deleting a region is nullifying the region definition. All the applications bound by the region are released back to the default region by deleting that region.

To delete a region (for example, a region numbered 4), perform this task:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **no cfs region 4**

Example:

```
WARNING: All applications in the region wiil be  moved to default region.
Are you sure? (y/n)  [n]
```

Deletes the Region 4.

Note After Step 2, you see the warning, “All the applications in the region will be moved to the default region.”

Verifying CFS Configurations

To display the CFS configuration information, perform one of the following tasks:

Command	Purpose
show cfs status	Displays the status of CFS distribution on the switch.

Command	Purpose
show cfs application	Displays the applications that are currently registered with CFS.
show cfs lock	Displays all the locks that are currently acquired by any application.
show cfs status	Verifies the CFS over IP configuration.
show cfs region brief	Displays brief information about the CFS regions.
show cfs region	Displays detailed information about the CFS regions.

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference* .

Verifying CFS Distribution Status

The **show cfs status** command displays the status of CFS distribution on the switch.

```
switch# show cfs status
Distribution : Enabled
Distribution over IP : Disabled
IPv4 multicast address : 239.255.70.83
IPv6 multicast address : ff15::efff:4653
```

Verifying Application Registration Status

The **show cfs application** command displays the applications that are currently registered with CFS. The first column displays the application name. The second column indicates whether the application is enabled or disabled for distribution (enabled or disabled). The last column indicates the scope of distribution for the application (logical, physical, or both).

Note The **show cfs application** command only displays applications registered with CFS. Conditional services that use CFS do not appear in the output unless these services are running.

```
switch# show cfs application
----------------------------------------------
 Application    Enabled   Scope
----------------------------------------------
 ntp            No        Physical-fc-ip
 fscm           Yes       Physical-fc
 role           No        Physical-fc-ip
 rscn           No        Logical
 radius         No        Physical-fc-ip
 fctimer        No        Physical-fc
 syslogd        No        Physical-fc-ip
 callhome       No        Physical-fc-ip
 fcdomain       No        Logical
 fc-redirect    Yes       Physical-fc
 device-alias   Yes       Physical-fc
Total number of entries = 11
```

The **show cfs application name** command displays the details for a particular application. It displays the enabled/disabled state, timeout as registered with CFS, merge capability (if it has registered with CFS for merge support), and the distribution scope.

```
switch# show cfs application name ntp
 Enabled        : Yes
 Timeout        : 5s
 Merge Capable  : Yes
 Scope          : Physical
 Region     : Default
```

Verifying CFS Lock Status

The **show cfs lock** command displays all the locks that are currently acquired by any application. For each application the command displays the application name and scope of the lock taken. If the application lock is taken in the physical scope, then this command displays the switch WWN, IP address, user name, and user type of the lock holder. If the application is taken in the logical scope, then this command displays the VSAN in which the lock is taken, the domain, IP address, user name, and user type of the lock holder.

```
switch# show cfs lock
Application: ntp
Scope      : Physical
--------------------------------------------------------------------
 Switch WWN               IP Address       User Name    User Type
--------------------------------------------------------------------
 20:00:00:05:30:00:6b:9e  10.76.100.167    admin        CLI/SNMP v3
 Total number of entries = 1
Application: port-security
Scope      : Logical
-----------------------------------------------------------
 VSAN   Domain   IP Address       User Name    User Type
-----------------------------------------------------------
 1      238      10.76.100.167    admin        CLI/SNMP v3
 2      211      10.76.100.167    admin        CLI/SNMP v3
 Total number of entries = 2
```

The **show cfs lock name** command displays the lock details similar for the specified application.

Lock Information for the Specified Application

```
switch# show cfs lock name ntp
Scope      : Physical
--------------------------------------------------------------------
 Switch WWN               IP Address       User Name    User Type
--------------------------------------------------------------------
 20:00:00:05:30:00:6b:9e  10.76.100.167    admin        CLI/SNMP v3
 Total number of entries = 1
```

Verifying the CFS over IP Configuration

To verify the CFS over IP configuration, use the **show cfs status** command.

```
switch# show cfs status
Distribution : Enabled
Distribution over IP : Disabled
IPv4 multicast address : 239.255.70.83
IPv6 multicast address : ff15::efff:4653
```

Verifying IP Multicast Address Configuration for CFS over IP

To verify the IP multicast address configuration for CFS over IP, use the **show cfs status** command.

```
switch# show cfs status
Fabric distribution Enabled
IP distribution Enabled mode ipv4
IPv4 multicast address : 10.1.10.100
IPv6 multicast address : ff13::e244:4754
```

Verifying Static IP Peer Configuration

To verify the IP peer configuration, use the **show cfs status** command.

```
switch# show cfs status
Distribution: Enabled
Distribution over IP: Enabled - mode IPv4 (static)
IPv4 multicast address : 239:255:70:83
IPv6 multicast address : ff15::efff:4563
```

To display the status of static IP peers discovery, use the **show cfs static peers** command.

```
switch# show cfs static peers
-------------------------------------------------------------
IP Address     WWN                         Status
-------------------------------------------------------------
192.0.2.4     00:00:00:00:00:00:00:00      Discovery in progress
192.0.2.5     20:00:00:0d:ec:06:55:b9      Reachable
192.0.2.6     20:00:00:0d:ec:06:55:c0      Local
```

Verifying CFS Regions

To display the CFS regions, perform this task:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **show cfs region brief**

Displays brief information about the CFS regions.

Step 3 switch(config)# **show cfs region**

Displays detailed information about the CFS regions.

Note To successfully form CFS peer you may configure common mcast IP on two different switches which are connected to two different management switches.

Additional References

For additional information related to implementing CFS, see the following section:

MIBs

MIBs	MIBs Link
• CISCO-CFS-CAPABILITY-MIB • CISCO-CFS-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

CHAPTER 4

Configuring System Message Logging

This chapter describes how to configure system message logging on Cisco MDS 9000 Series switches.

Feature History for System Message Logging

Table 2: Feature History for Configuring SAN Analytics

Feature Name	Release	Feature Information
Secure Remote System Message Logging	9.2(2)	Support for CFS distribution of the **secure** option was added.
Secure Remote System Message Logging	9.2(1)	The Secure Remote System Message Logging feature allows you to securely log system messages to a remote logging server using TLS.

Information About System Message Logging

With the system message logging software, you can save messages in a log file or direct the messages to other devices. By default, the switch logs normal but significant system messages to a log file and sends these messages to the system console. This feature provides you with the following capabilities:

- Provides logging information for monitoring and troubleshooting
- Allows you to select the types of captured logging information
- Allows you to select the destination server to forward the captured logging information properly configured system message logging server.

Note When the switch first initializes, the network is not connected until initialization completes. Therefore, messages are not redirected to a system message logging server for a few seconds.

Log messages are not saved across system reboots. However, a maximum of 100 log messages with a severity level of critical and below (levels 0, 1, and 2) are saved in NVRAM.

Table 3: Internal Logging Facilities , on page 28 describes some samples of the facilities supported by the system message logs.

Table 3: Internal Logging Facilities

Facility Keyword	Description	Standard or Cisco MDS Specific
acl	ACL manager	Cisco MDS 9000 Family specific
all	All facilities	Cisco MDS 9000 Family specific
auth	Authorization system	Standard
authpriv	Authorization (private) system	Standard
bootvar	Bootvar	Cisco MDS 9000 Family specific
callhome	Call Home	Cisco MDS 9000 Family specific
cron	Cron or at facility	Standard
daemon	System daemons	Standard
fcc	FCC	Cisco MDS 9000 Family specific
fcdomain	fcdomain	Cisco MDS 9000 Family specific
fcns	Name server	Cisco MDS 9000 Family specific
fcs	FCS	Cisco MDS 9000 Family specific
flogi	FLOGI	Cisco MDS 9000 Family specific
fspf	FSPF	Cisco MDS 9000 Family specific
ftp	File Transfer Protocol	Standard
ipconf	IP configuration	Cisco MDS 9000 Family specific
ipfc	IPFC	Cisco MDS 9000 Family specific
kernel	Kernel	Standard
local0 to local7	Locally defined messages	Standard
lpr	Line printer system	Standard
mail	Mail system	Standard

Facility Keyword	Description	Standard or Cisco MDS Specific
mcast	Multicast	Cisco MDS 9000 Family specific
module	Switching module	Cisco MDS 9000 Family specific
news	USENET news	Standard
ntp	NTP	Cisco MDS 9000 Family specific
platform	Platform manager	Cisco MDS 9000 Family specific
port	Port	Cisco MDS 9000 Family specific
port-channel	PortChannel	Cisco MDS 9000 Family specific
qos	QoS	Cisco MDS 9000 Family specific
rdl	RDL	Cisco MDS 9000 Family specific
rib	RIB	Cisco MDS 9000 Family specific
rscn	RSCN	Cisco MDS 9000 Family specific
securityd	Security	Cisco MDS 9000 Family specific
syslog	Internal system messages	Standard
sysmgr	System manager	Cisco MDS 9000 Family specific
tlport	TL port	Cisco MDS 9000 Family specific
user	User process	Standard
uucp	UNIX-to-UNIX Copy Program	Standard
vhbad	Virtual host base adapter daemon	Cisco MDS 9000 Family specific
vni	Virtual network interface	Cisco MDS 9000 Family specific
vrrp_cfg	VRRP configuration	Cisco MDS 9000 Family specific
vrrp_eng	VRRP engine	Cisco MDS 9000 Family specific
vsan	VSAN system messages	Cisco MDS 9000 Family specific
vshd	vshd	Cisco MDS 9000 Family specific
wwn	WWN manager	Cisco MDS 9000 Family specific
xbar	Xbar system messages	Cisco MDS 9000 Family specific
zone	Zone server	Cisco MDS 9000 Family specific

Table 4: Error Message Severity Levels , on page 30 describes the severity levels supported by the system message logs.

Table 4: Error Message Severity Levels

Level Keyword	Level	Description	System Message Definition
emergencies	0	System unusable	LOG_EMERG
alerts	1	Immediate action needed	LOG_ALERT
critical	2	Critical conditions	LOG_CRIT
errors	3	Error conditions	LOG_ERR
warnings	4	Warning conditions	LOG_WARNING
notifications	5	Normal but significant condition	LOG_NOTICE
informational	6	Informational messages only	LOG_INFO
debugging	7	Debugging messages	LOG_DEBUG

Note Refer to the *Cisco MDS 9000 Family System Messages Reference* for details on the error log message format.

System Message Logging

The System Message Logging feature allows system messages to be logged for later reference. This feature has the following capabilities:

- Provides logging information for monitoring and troubleshooting.
- Allows the user to select the types of captured logging information.
- Allows the user to forward the captured logging information to remote logging servers.

Messages are time stamped to enhance real time debugging and message management.

By default, the switch logs normal but significant system messages to an onboard logfile and the system console as they occur. The onboard logfile is circular and can store up to the last 1200 messages. Messages stored in the onboard logfile can be viewed using the CLI.

System messages may be displayed in real time in a user's session to the switch. This allows real time monitoring of switch events when troubleshooting. The minimum severity of messages to be displayed to sessions is configurable.

System messages may also be logged to remote logging servers. Up to three remote destinations may be configured. These may be a mix of IPv4 and IPv6 addresses. By default, when a remote logging destination is configured, system messages are sent using UDP. From Cisco MDS NX-OS Release 9.2(1), logging over a secure Transport Layer Security (TLS) connection and mutual device authentication is supported. The Cisco MDS device is the TLS client and initiates a connection to the remote logging server. This allows transport encryption for secure logging over an unsecure network. From Cisco MDS NX-OS Release 9.2(2), distribution over Cisco Fabric Services (CFS) of secure syslog server configurations is supported.

Tip To be able to compare system messages from multiple devices ensure that all devices have the correct time. This will allow the sequence of events involving multiple devices to be understood. Device clocks can be synchronised by using NTP.

The system messages to be logged to each destination can be filtered based on the facility and the severity level.

SFP Diagnostics

The error message related to SFP failures is written to the syslog. You can listen to the syslog for events related to SFP failures. The values, low or high alarm, and the warning are checked for the following parameters:

- TX Power
- RX Power
- Temperature
- Voltage
- Current

The SFP notification trap indicates the current status of the alarm and warning monitoring parameters for all the sensors based on the digital diagnostic monitoring information. This notification is generated whenever there is a change in the status of at least one of the monitoring parameters of the sensors on the transceiver in an interface.

The CISCO-INTERFACE-XCVR-MONITOR-MIB contains the SFP notification trap information. Refer to the *Cisco MDS 9000 Family MIB Quick Reference* for more information on this MIB.

Outgoing System Message Logging Server Facilities

All system messages have a logging facility and a level. The logging facility can be thought of as *where* and the level can be thought of as *what*.

The single system message logging daemon (syslogd) sends the information based on the configured **facility** option. If no facility is specified, local7 is the default outgoing facility.

The internal facilities are listed in Table 3: Internal Logging Facilities , on page 28 and the outgoing logging facilities are listed in Table 5: Outgoing Logging Facilities , on page 31.

Table 5: Outgoing Logging Facilities

Facility Keyword	Description	Standard or Cisco MDS Specific
auth	Authorization system	Standard
authpriv	Authorization (private) system	Standard
cron	Cron or at facility	Standard
daemon	System daemons	Standard
ftp	File Transfer Protocol	Standard

Facility Keyword	Description	Standard or Cisco MDS Specific
kernel	Kernel	Standard
local0 to local7	Locally defined messages	Standard (local7 is the default)
lpr	Line printer system	Standard
mail	Mail system	Standard
news	USENET news	Standard
syslog	Internal system messages	Standard
user	User process	Standard
uucp	UNIX-to-UNIX Copy Program	Standard

System Message Logging Configuration Distribution

You can enable fabric distribution for all Cisco MDS switches in the fabric. When you perform system message logging configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric.

You automatically acquire a fabric-wide lock when you issue the first configuration command after you enabled distribution in a switch. The system message logging server uses the effective and pending database model to store or commit the commands based on your configuration. When you commit the configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. After making the configuration changes, you can choose to discard the changes by terminating the changes instead of committing them. In either case, the lock is released. See Using the CFS Infrastructure, on page 7 for more information on the CFS application.

Fabric Lock Override

If you have performed a system message logging task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.

Tip The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

Guidelines and Limitations for System Message Logging

- CFS distribution must be enabled for synchronized distribution of secure syslog configuration across a fabric.
- In Cisco MDS NX-OS Release 9.2(1), you can either configure the secure option for remote system logging servers or CFS distribution of the system logging configuration. You cannot configure both. If

you try configuring a secure remote destination when CFS distribution for logging is enabled, you will be prompted with a message to disable CFS distribution for logging before configuring a secure remote destination, and vice versa.

- CA certificates must be installed for a TLS connection to be used and mutual authentication of secure remote logging server connections. Hence, a warning message is displayed after each secure syslog configuration command. For information on configuring CA certificates, see the "Configuring Certificate Authorities and Digital Certificates" chapter in the Cisco MDS 9000 Series Security Configuration Guide, Release 9.x.
- Any system messages that are logged before any remote syslog servers are reachable (such as supervisor active or online messages) will not be sent to the syslog server.

When merging two fabrics with CFS that have different system message logging configurations, follow these guidelines:

- Be aware that the merged configuration is a union of the existing and received configuration for each switch in the fabric.
- Verify that the merged configuration will only have a maximum of three unique system message logging servers.

⚠

Caution If the merged configuration contains more that three servers, the merge will fail.

For detailed concepts on CFS merge, see CFS Merge Support, on page 11.

Default Settings

Table 6: Default System Message Log Settings , on page 33 lists the default settings for system message logging.

Table 6: Default System Message Log Settings

Parameters	Default
System message logging to the console	Enabled for messages at the critical severity level.
System message logging to sessions	Disabled.
Onboard logging file size	4194304 bytes.
Onboard logging file name	messages
Remote server facility	local7
Remote logging destinations	Not configured.
Unsecure remote server destination port	UDP 514
Secure remote server destination port	TCP 6514

Parameters	Default
CA certificates	Not installed.

Configuring System Message Logging

System logging messages are sent to the console based on the default (or configured) logging facility and severity values.

Task Flow for Configuring System Message Logging

Follow these steps to configure system message logging:

Procedure

Step 1 Enable or disable message logging.

Step 2 Configure console severity level.

Step 3 Configure monitor severity level.

Step 4 Configure module log severity level.

Step 5 Configure facility severity levels.

Step 6 Configure the onboard log file.

Step 7 Configure system message logging servers.

Step 8 Configure system message logging distribution.

Enabling or Disabling Message Logging

You can disable logging to the console or enable logging to a specific Telnet or SSH session.

- When you disable or enable logging to a console session, that state is applied to all future console sessions. If you exit and log in again to a new session, the state is preserved.
- When you enable or disable logging to a Telnet or SSH session, that state is applied only to that session. If you exit and log in again to a new session, the state is not preserved.

To enable or disable the logging state for a Telnet or SSH session, follow these steps:

Procedure

Step 1 switch# **terminal monitor**

Enables logging for a Telnet or SSH session.

Note Logging to the console session is enabled by default.

Step 2 switch# **terminal no monitor**

Disables logging for a Telnet or SSH session.

Note A Telnet or SSH session is disabled by default.

Configuring Console Severity Level

When logging is enabled for a console session (default), you can configure the severity levels of messages that appear on the console. The default severity for console logging is 2 (critical).

Note The current critical (default) logging level is maintained if the console baud speed is 9600 baud (default). All attempts to change the console logging level generates an error message. To increase the logging level (above critical), you must change the console baud speed to 38400 baud.

To configure the severity level for the console session, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **logging console 3**

Configures console logging at level 3 (error). Logging messages with a severity level of 3 or above are displayed on the console.

Step 3 switch(config)# **no logging console**

Reverts console logging to the factory set default severity level of 2 (critical). Logging messages with a severity level of 2 or above are displayed on the console.

Configuring Monitor Severity Level

When logging is enabled for a monitor session (default), you can configure the severity levels of messages that appear on the monitor. The default severity for monitor logging is 5 (notifications).

To configure the severity level for a monitor session, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **logging monitor 3**

Configures monitor logging at level 3 (error). Logging messages with a severity level of 3 or above are displayed on the monitor.

Step 3 switch(config)# **no logging monitor**

Reverts monitor logging to the factory set default severity level of 5 (notifications). Logging messages with a severity level of 5 or above are displayed on the console.

Configuring Module Logging

By default, logging is enabled at level 7 for all modules. You can enable or disable logging for each module at a specified level.

To enable or disable the logging for modules and configure the severity level, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **logging module 1**

Configures module logging at level 1 (alerts) for all modules.

Step 3 switch(config)# **logging module**

Configures module logging for all modules in the switch at the default level 5 (notifications).

Step 4 switch(config)# **no logging module**

Disables module logging.

Configuring Facility Severity Levels

To configure the severity level for a logging facility (see Table 3: Internal Logging Facilities , on page 28), follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **logging level kernel 4**

Configures Telnet or SSH logging for the kernel facility at level 4 (warning). As a result, logging messages with a severity level of 4 or above are displayed.

Step 3 switch(config)# **no logging level kernel 4**

Reverts to the default severity level 6 (informational) for the Telnet or SSH logging for the kernel facility.

Note Use the **show logging info** command to display the default logging levels for the facilities listed in Table 3: Internal Logging Facilities , on page 28.

Configuring the Onboard Log File

By default, the switch logs normal but significant system messages to a log file and sends these messages to the system console. Log messages are not saved across system reboots. The logging messages that are generated may be saved to a log file. You can configure the name of this file and restrict its size as required. The default log file name is messages.

The file name can have up to 80 characters and the file size ranges from 4096 bytes to 4194304 bytes.

To send log messages to a file, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **logging logfile messages 3**

Configures logging of information for errors or events above with a severity level 3 or above to the default log file named messages.

Step 3 switch(config)# **logging logfile ManagerLog 3**

Configures logging of information for errors or events with a severity level 3 or above to a file named ManagerLog using the default size of 10,485,760 bytes.

Step 4 switch(config)# **logging logfile ManagerLog 3 size 3000000**

Configures logging information for errors or events with a severity level 3 or above to a file named ManagerLog. By configuring a size, you are restricting the file size to 3,000,000 bytes.

Step 5 switch(config)# **no logging logfile**

Disables logging messages to the logfile.

You can rename the log file using the **logging logfile** command.

The location of the log file cannot be changed. You can use the **show logging logfile** and clear logging logfile commands to view and delete the contents of this file. You can use the **dir log:** command to view logging file statistics. You can use the **delete log:** command to remove the log file.

You can copy the logfile to a different location using the **copy log:** command using additional copy syntax.

Configuring System Message Logging to Remote Logging Destinations

To configure system message logging to a remote logging destination, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **logging server** *name* [*severity-level*] [**port** *number*] [**secure** [**trustpoint client-identity** *name*]] [**facility** *facility-name*]

Configures system message logging to a remote destination at the specified hostname, IPv4, or IPv6 address. Specify the minimum severity of forwarded messages with the *severity-level* parameter. Use the **port** option to override the default destination port number. Use the **secure** option to use TCP, use the secure destination port, and encrypt the connection to the remote logging server using TLS. For TLS mutual authentication to succeed, identity certificates signed by trusted CAs must be installed using **crypto** commands. By default, certificates from all trust points are sequentially tried until authentication succeeds. Optionally, the certificates used for authentication may be restricted to a single trust point by specifying the **trustpoint client-identity** option. Use the **facility** option to specify a different logging category.

Step 3 switch(config)# **syslog priority** 1 **msg** "test message"

(Optional) Logs a test message to all system message logging destinations. This may be used to verify that logging to remote destinations is working.

Step 4 switch(config)# **no logging server** *name*

Removes the specified server as a destination for system message logs.

Configuring the Origin ID for System Messages

To specify the hostname, IP address, or a text string in the system messages that are sent to remote syslog servers, follow these steps:

Procedure

Step 1 switch# **configure**

Enters configuration mode.

Step 2 switch(config)# **logging origin-id** {**hostname** | **ip** *address* | **string** *word*}

Specifies the hostname, IP address, or a text string in the system messages that are sent to remote syslog servers.

Configuring System Message Logging Servers

You can configure a maximum of three system message logging servers. To send log messages to a UNIX system message logging server, you must configure the system message logging daemon on a UNIX server. Log in as a privileged user, and follow these steps:

Procedure

Step 1 Add the following line to the /etc/syslog.conf file.

local1.debug /var/log/ myfile .log

Note Be sure to add five tab characters between **local1.debug** and **/var/log/***myfile***.log**. Refer to entries in the /etc/syslog.conf file for further examples.

The switch sends messages according to the specified facility types and severity levels. The **local1** keyword specifies the UNIX logging facility used. The messages from the switch are generated by user processes. The **debug** keyword specifies the severity level of the condition being logged. You can set UNIX systems to receive all messages from the switch.

Step 2 Create the log file by entering these commands at the UNIX shell prompt:

$ touch /var/log/ myfile .log

$ chmod 666 /var/log/ myfile .log

Step 3 Make sure the system message logging daemon reads the new changes by entering this command:

$ kill -HUP ~cat /etc/syslog.pid~

Configuring System Message Logging Distribution

To enable fabric distribution for system message logging server configurations, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **logging distribute**

Enables the system message logging server configuration to be distributed to all switches in the fabric, acquires a lock, and stores all future configuration changes in the pending database.

Step 3 switch(config)# **no logging distribute**

Disables (default) system message logging server configuration distribution to all switches in the fabric.

Committing Changes

To commit the system message logging server configuration changes, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **logging commit**

Distributes the configuration changes to all switches in the fabric, releases the lock, and overwrites the effective database with the changes made to the pending database.

Discarding Changes

To discard the system message logging server configuration changes, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **logging abort**

Discards the system message logging server configuration changes in the pending database and releases the fabric lock.

Fabric Lock Override

To use administrative privileges and release a locked system message logging session, use the **clear logging session** command.

```
switch# clear logging session
```

Displaying System Message Logging Information

To display the system message logging information, perform one of the following tasks:

Command	Purpose
show logging	Displays current system message logging.
show logging nvram	Displays NVRM log contents.

Command	Purpose
show logging logfile	Displays the log file.
show logging level	Displays logging facility.
show logging info	Displays logging information.
show logging last 2	Displays last few lines of a log file.
show logging module	Displays switching module logging status.
show logging monitor	Displays monitor logging status.
show logging server	Displays server information.

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference* .

Use the **show logging** command to display the current system message logging configuration . See Examples Current System Message Logging, on page 41 to Remote Logging Server Information, on page 46.

Note When using the **show logging** command, output is displayed only when the configured logging levels for the switch are different from the default levels.

Current System Message Logging

The following example displays the current system message logging settings and contents of the onboard log file:

```
switch# show logging

Logging console:                enabled (Severity: critical)
Logging monitor:                enabled (Severity: debugging)
Logging linecard:               enabled (Severity: debugging)
Logging server:                 enabled
{172.20.102.34}
        server severity:        debugging
        server facility:        local7
{10.77.202.88}
        server severity:        debugging
        server facility:        local7
{10.77.202.149}
        server severity:        debugging
        server facility:        local7
Logging logfile:                enabled
        Name - messages: Severity - debugging Size - 4194304
Facility        Default Severity        Current Session Severity
--------        ----------------        ------------------------
kern                    6                       6
user                    3                       3
mail                    3                       3
daemon                  7                       7
auth                    0                       7
syslog                  3                       3
```

```
lpr                     3                       3
news                    3                       3
uucp                    3                       3
cron                    3                       3
authpriv                3                       7
ftp                     3                       3
local0                  3                       3
local1                  3                       3
local2                  3                       3
local3                  3                       3
local4                  3                       3
local5                  3                       3
local6                  3                       3
local7                  3                       3
vsan                    2                       2
fspf                    3                       3
fcdomain                2                       2
module                  5                       5
sysmgr                  3                       3
zone                    2                       2
vni                     2                       2
ipconf                  2                       2
ipfc                    2                       2
xbar                    3                       3
fcns                    2                       2
fcs                     2                       2
acl                     2                       2
tlport                  2                       2
port                    5                       5
flogi                   2                       2
port_channel            5                       5
wwn                     3                       3
fcc                     2                       2
qos                     3                       3
vrrp_cfg                2                       2
ntp                     2                       2
platform                5                       5
vrrp_eng                2                       2
callhome                2                       2
mcast                   2                       2
rdl                     2                       2
rscn                    2                       2
bootvar                 5                       2
securityd               2                       2
vhbad                   2                       2
rib                     2                       2
vshd                    5                       5
0(emergencies)          1(alerts)       2(critical)
3(errors)               4(warnings)     5(notifications)
6(information)          7(debugging)
Feb 14 09:50:57 switchname %TTYD-6-TTYD_MISC: TTYD TTYD started
Feb 14 09:50:58 switchname %DAEMON-6-SYSTEM_MSG: precision = 8 usec
...
```

Use the **show logging nvram** command to view the log messages saved in NVRAM. Only log messages with a severity level of critical and below (levels 0, 1, and 2) are saved in NVRAM.

NVRM Log Contents

The following example displays the NVRM log contents:

```
switch# show logging nvram

Jul 16 20:36:46 switchname %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2209, ret_val = -105)
Jul 16 20:36:46 switchname %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2199, ret_val = -105)
Jul 16 20:36:46 switchname %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2213, ret_val = -105)
Jul 16 20:36:46 switchname %KERN-2-SYSTEM_MSG: unable to alloc and fill in a
new mtsbuf (pid=2213, ret_val = -105)
...
```

Log File

The following example displays the onboard log file:

```
switch# show logging logfile

Jul 16 21:06:50 %DAEMON-3-SYSTEM_MSG: Un-parsable frequency in /mnt/pss/ntp.drift
Jul 16 21:06:56 %DAEMON-3-SYSTEM_MSG: snmpd:snmp_open_debug_cfg: no snmp_saved_dbg_uri ;
Jul 16 21:06:58 switchname %PORT-5-IF_UP: Interface mgmt0 is up
Jul 16 21:06:58 switchname %MODULE-5-ACTIVE_SUP_OK: Supervisor 5 is active
...
```

Console Logging Status

The following example displays the console logging status:

```
switch# show logging console

Logging console:                enabled (Severity: notifications)
```

Logging Facility

The following example displays the logging level of each switch facility:

```
switch# show logging level

Facility        Default Severity        Current Session Severity
--------        ----------------        ------------------------
kern                    6                       6
user                    3                       3
mail                    3                       3
daemon                  7                       7
auth                    0                       7
syslog                  3                       3
lpr                     3                       3
news                    3                       3
uucp                    3                       3
cron                    3                       3
authpriv                3                       7
ftp                     3                       3
local0                  3                       3
local1                  3                       3
local2                  3                       3
```

```
local3                  3                       3
local4                  3                       3
local5                  3                       3
local6                  3                       3
local7                  3                       3
vsan                    2                       2
fspf                    3                       3
fcdomain                2                       2
module                  5                       5
sysmgr                  3                       3
zone                    2                       2
vni                     2                       2
ipconf                  2                       2
ipfc                    2                       2
xbar                    3                       3
fcns                    2                       2
fcs                     2                       2
acl                     2                       2
tlport                  2                       2
port                    5                       5
flogi                   2                       2
port_channel            5                       5
wwn                     3                       3
fcc                     2                       2
qos                     3                       3
vrrp_cfg                2                       2
ntp                     2                       2
platform                5                       5
vrrp_eng                2                       2
callhome                2                       2
mcast                   2                       2
rdl                     2                       2
rscn                    2                       2
bootvar                 5                       2
securityd               2                       2
vhbad                   2                       2
rib                     2                       2
vshd                    5                       5
0(emergencies)          1(alerts)       2(critical)
3(errors)               4(warnings)     5(notifications)
6(information)          7(debugging)
```

Logging Information

The following example displays the current system message logging settings:

```
switch# show logging info

Logging console:                enabled (Severity: critical)
Logging monitor:                enabled (Severity: debugging)
Logging linecard:               enabled (Severity: debugging)
Logging server:                 enabled
{192.168.1.34}
        server severity:        debugging
        server facility:        local7
{192.168.1.88}
        server severity:        debugging
        server facility:        local7
{192.168.1.149}
        server severity:        debugging
        server facility:        local7
```

```
Logging logfile:                enabled
        Name - messages: Severity - debugging Size - 4194304
Facility        Default Severity        Current Session Severity
--------        ----------------        ------------------------
kern                    6                       6
user                    3                       3
mail                    3                       3
daemon                  7                       7
auth                    0                       7
syslog                  3                       3
lpr                     3                       3
news                    3                       3
uucp                    3                       3
cron                    3                       3
authpriv                3                       7
ftp                     3                       3
local0                  3                       3
local1                  3                       3
local2                  3                       3
local3                  3                       3
local4                  3                       3
local5                  3                       3
local6                  3                       3
local7                  3                       3
vsan                    2                       2
fspf                    3                       3
fcdomain                2                       2
module                  5                       5
sysmgr                  3                       3
zone                    2                       2
vni                     2                       2
ipconf                  2                       2
ipfc                    2                       2
xbar                    3                       3
fcns                    2                       2
fcs                     2                       2
acl                     2                       2
tlport                  2                       2
port                    5                       5
flogi                   2                       2
port_channel            5                       5
wwn                     3                       3
fcc                     2                       2
qos                     3                       3
vrrp_cfg                2                       2
ntp                     2                       2
platform                5                       5
vrrp_eng                2                       2
callhome                2                       2
mcast                   2                       2
rdl                     2                       2
rscn                    2                       2
bootvar                 5                       2
securityd               2                       2
vhbad                   2                       2
rib                     2                       2
vshd                    5                       5
0(emergencies)          1(alerts)       2(critical)
3(errors)               4(warnings)     5(notifications)
6(information)          7(debugging)
```

Last Few Lines of a Log File

The following example displays the last few lines of a log file:

```
switch# show logging last 2

Nov 8 16:48:04 switchname %LOG_VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from pts/1
 (171.71.58.56)
Nov 8 17:44:09 switchname %LOG_VSHD-5-VSHD_SYSLOG_CONFIG_I: Configuring console from pts/0
 (171.71.58.72)
```

Switching Module Logging Status

The following example displays switching module logging status:

```
switch# show logging module


Logging linecard:                enabled (Severity: debugging)
```

Monitor Logging Status

The following example displays the monitor logging status:

```
switch# show logging monitor

Logging monitor:                 enabled (Severity: information)
```

Remote Logging Server Information

The following example displays the configured remote logging server information:

```
switch# show logging server
Logging server:                  enabled
{192.168.113.1}
        server severity:         notifications
        server facility:         local7
        server VRF:              default
        server port:             55552
        server transport:        secure
{192.168.106.50}
        server severity:         notifications
        server facility:         local7
        server VRF:              default
        server port:             55551
        server transport:        secure
{192.168.229.220}
        server severity:         notifications
        server facility:         local7
        server VRF:              default
        server port:             55552
```

Additional References

For additional information related to implementing system message logging, see the following section:

MIBs

MIBs	MIBs Link
• CISCO-SYSLOG-EXT-MIB • CISCO-SYSLOG-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

CHAPTER 5

Configuring Call Home

Call Home provides email-based notification of critical system events. A versatile range of message formats are available for optimal compatibility with pager services, standard email, or XML-based automated parsing applications.

Note Cisco Autonotify is upgraded to a new capability called Smart Call Home. Smart Call Home has significant functionality improvement over Autonotify and is available across the Cisco product range. For detailed information on Smart Call Home, see the Smart Call Home page at this location: http://www.cisco.com/go/smartcall/ .

This chapter includes the following sections:

Information About Call Home

The Call Home feature provides message throttling capabilities. Periodic inventory messages, port syslog messages, and RMON alert messages are added to the list of deliverable Call Home messages. If required you can also use the Cisco Fabric Services application to distribute the Call Home configuration to all other switches in the fabric.

The Call Home service provides email-based notification of critical system events. A versatile range of message formats are available for optimal compatibility with pager services, standard email, or XML-based automated parsing applications.

Common features may include the following:

- Paging the network support engineer
- Emailing the Network Operations Center
- Raising a direct case with the Technical Assistance Center

The Call Home functionality is available directly through the Cisco MDS 9000 Series switches and the Cisco Nexus 5000 Series switches. It provides multiple Call Home messages, each with separate potential destinations. You can define your own destination profiles in addition to predefined profiles; you can configure up to 50 email addresses for each destination profile. Flexible message delivery and format options make it easy to integrate specific support requirements.

The Call Home feature offers the following advantages:

- Fixed set of predefined alerts for trigger events on the switch.
- Automatic execution and attachment of relevant command output.

Call Home Features

The Call Home functionality is available directly through the Cisco MDS 9000 Series switches and the Cisco Nexus 5000 Series switches. It provides multiple Call Home profiles (also referred to as *Call Home destination profiles*), each with separate potential destinations. You can define your own destination profiles in addition to predefined profiles.

The Call Home function can even leverage support from Cisco Systems or another support partner. Flexible message delivery and format options make it easy to integrate specific support requirements.

The Call Home feature offers the following advantages:

- Fixed set of predefined alerts and trigger events on the switch.
- Automatic execution and attachment of relevant command output.
- Multiple message format options:
    - Short Text—Suitable for pagers or printed reports.
    - Plain Text—Full formatted message information suitable for human reading.
    - XML—Matching readable format using Extensible Markup Language (XML) and document type definitions (DTDs) named Messaging Markup Language (MML). The MML DTD is published on the Cisco.com website at http://www.cisco.com/. The XML format enables communication with the Cisco Systems Technical Assistance Center.
- Multiple concurrent message destinations. You can configure up to 50 email destination addresses for each destination profile.
- Multiple message categories including system, environment, switching module hardware, supervisor module, hardware, inventory, syslog, RMON, and test.
- Secure messages transport directly from your device or through an HTTP proxy server or a downloadable transport gateway (TG). You can use a TG aggregation point to support multiple devices, or in cases where security requires that your devices not be connected directly to the Internet.

Note: Beginning from the Cisco MDS Release 7.3(0)D1(1), all the alerts are classified under the type, Environment and sub type, Minor.

- SUP_FAILURE, POWER_SUPPY_FAILURE, LINECARD_FAILURE alerts are classified under the type, Environment and sub type, Major.

About Smart Call Home

Smart Call Home is a component of Cisco SMARTnet Service that offers proactive diagnostics, real-time alerts, and personalized web-based reports on select Cisco devices.

Smart Call Home provides fast resolution of system problems by analyzing Call Home messages sent from your devices and providing a direct notification path to Cisco customer support.

Smart Call Home offers the following features:

- Continuous device health monitoring and real-time diagnostics alerts.
- Analysis of Call Home messages from your device and where appropriate, automatic service request generation, routed to the appropriate TAC team, including detailed diagnostic information to speed problem resolution.
- Web-based access to Call Home messages and recommendations, inventory and configuration information for all Call Home devices. Provides access to associated Field Notices, Security Advisories and End-of-Life Information

.

Table 7: Benefits of Smart Call Home Compared to Autonotify , on page 51 lists the benefits of Smart Call Home.

Table 7: Benefits of Smart Call Home Compared to Autonotify

Feature	Smart Call Home	Autonotify
Low touch registration	The registration process is considerably streamlined. Customers no longer need to know their device serial number or contract information. They can register devices without manual intervention from Cisco by sending a message from those devices. The procedures are outlined at www.cisco.com/go/smartcall.	Requires the customer to request Cisco to add each specific serial number to the database.
Recommendations	Smart Call Home provides recommendations for known issues including those for which SRs are raised and for which SRs are not appropriate but for which customers might want to still take action on.	Autonotify raises SRs for a set of failure scenarios but no recommendations are provided for these.
Device report	Device report includes full inventory and configuration details. Once available, the information in these reports will be mapped to field notices, PSIRTs, EoX notices, configuration best practices and bugs.	No.

Feature	Smart Call Home	Autonotify
History report	The history report is available to look up any message and its contents, including show commands, message processing, analysis results, recommendations and service request numbers for all messages sent over the past three months.	A basic version is available that does not include contents of message.
Network summary report	A report that provides a summary of the make-up of devices and modules in the customer network (for those devices registered with Smart Call home).	No.
Cisco device support	Device Support will be extended across the Cisco product range. See the supported products table at www.cisco.com/go/smartcall.	Deprecated in favor of Smart Call Home in October 2008.

Obtaining Smart Call Home

If you have a service contract directly with Cisco Systems, you can receive automatic case generation from the Technical Assistance Center by registering with the Smart Call Home service.

You need the following items to register:

- The SMARTnet contract number for your switch.
- Your email address
- Your Cisco.com ID

For detailed information on Smart Call Home, including quick start configuration and registration steps, see the Smart Call Home page at this location:

http://www.cisco.com/go/smartcall/

Call Home Destination Profiles

A destination profile contains the required delivery information for an alert notification. Destination profiles are typically configured by the network administrator.

Using alert groups you can select the set of Call Home alerts to be received by a destination profile (predefined or user defined). Alert groups are predefined subsets of Call Home alerts supported in all switches in the Cisco MDS 9000 Series and the Cisco Nexus 5000 Series. Different types of Call Home alerts are grouped into different alert groups depending on their type. You can associate one or more alert groups to each profile as required by your network.

Call Home Alert Groups

An alert group is a predefined subset of Call Home alerts supported in all switches in the Cisco MDS 9000 Series and Cisco Nexus 5000 Series. Alert groups allow you to select the set of Call Home alerts to be received by a destination profile (predefined or user-defined). A Call Home alert is sent to email destinations in a destination profile only if that Call Home alert belongs to one of the alert groups associated with that destination profile.

Using the predefined Call Home alert groups you can generate notification messages when certain events occur on the switch. You can customize predefined alert groups to execute additional **show** commands when specific events occur and to notify you of output other than from the predefined **show** commands.

Customized Alert Group Messages

An alert group is a predefined subset of Call Home alerts supported in all switches in the Cisco MDS 9000 Series and Cisco Nexus 5000 Series switches. Alert groups allow you to select the set of Call Home alerts to be received by a destination profile (predefined or user-defined). The predefined Call Home alert groups generate notification messages when certain events occur on the switch. You can customize predefined alert groups to execute additional show commands when specific events occur.

The output from these additional **show** commands is included in the notification message along with the output of the predefined **show** commands.

Call Home Message Level Feature

The Call Home message level feature allows you to filter messages based on their level of urgency. Each destination profile (predefined and user-defined) is associated with a Call Home message level threshold. Any message with a value lower than the urgency threshold is not sent. Call Home severity levels are not the same as system message logging severity levels.

Syslog-Based Alerts

You can configure the switch to send certain syslog messages as Call Home messages. The messages are sent based on the mapping between the destination profile and the alert group mapping, and on the severity level of the generated syslog message.

To receive a syslog-based Call Home alert, you must associate a destination profile with the syslog alert groups (currently there is only one syslog alert group—syslog-group-port) and configure the appropriate message level.

The syslog-group-port alert group selects syslog messages for the port facility. The Call Home application maps the syslog severity level to the corresponding Call Home severity level (see Table 8: Event Triggers , on page 56). For example, if you select level 5 for the Call Home message level, syslog messages at levels 0, 1, and 2 are included in the Call Home log.

Whenever a syslog message is generated, the Call Home application sends a Call Home message depending on the mapping between the destination profile and the alert group mapping and based on the severity level of the generated syslog message. To receive a syslog-based Call Home alert, you must associate a destination profile with the syslog alert groups (currently there is only one syslog alert group—syslog-group-port) and configure the appropriate message level (see Table 8: Event Triggers , on page 56).

Note Call Home does not change the syslog message level in the message text. The syslog message texts in the Call Home log appear as they are described in the *Cisco MDS 9000 Series System Messages* Reference.

RMON-Based Alerts

You can configure the switch to send Call Home notifications corresponding to RMON alert triggers. All RMON-based Call Home messages have their message level set to NOTIFY (2). The RMON alert group is defined for all RMON-based Call Home alerts. To receive an RMON-based Call Home alert, you must associate a destination profile with the RMON alert group.

General EMail Options Using HTTPS Support

The HTTPS support for Call Home provides a transport method called HTTP. HTTPS support is used for a secure communication, and HTTP is used for nonsecure communication. You can configure an HTTP URL for the Call Home destination profile as a destination. The URL link can be from a secure server or nonsecure server. For a destination profile configured with the HTTP URL, the Call Home message is posted to the HTTP URL link.

Note The Call Home HTTP configuration can be distributed over CFS on the switches running NX-OS Release 4.2(1) and later. The Call Home HTTP configuration cannot be distributed to switches that support the nondistributable HTTP configuration. Switches running lower versions than NX-OS Release 4.2(1) and later will ignore the HTTP configuration.

Multiple SMTP Server Support

Cisco MDS NX-OS and Cisco NX-OS 5000 Series switches support multiple SMTP servers for Call Home. Each SMTP server has a priority configured between 1 and 100, with 1 being the highest priority and 100 being the lowest. If the priority is not specified, a default value of 50 is used.

You can configure up to five SMTP servers for Call Home. The servers are contacted based on their priority. The highest priority server is contacted first. If the message fails to be sent, the next server in the list is contacted until the limit is exhausted. If two servers have equal priority, the one that was configured earlier is contacted.

If a high-priority SMTP server fails, the other servers will be contacted. A time delay may occur while sending a message. The delay is minimal if the attempt to send the message through the first SMTP server is successful. The delay may increase depending on the number of unsuccessful attempts with different SMTP servers.

Note The new configuration process is not related to the old configuration. However, if the SMTP servers are configured using both the old and new schemes, the older configuration is of the highest priority.

Multiple SMTP servers can be configured on any MDS 9000 Series switch, Cisco Nexus 5000 Series switches, and Cisco Nexus 7000 Series switches running Release 5.0(1a) or later.

The new configuration will only be distributed to switches that have multiple SMTP servers. The older switches in the fabric will ignore the new configuration received over CFS.

In a mixed fabric that has CFS enabled, the switches running NX-OS Release 5.0 can configure new functionalities and distribute the new configuration to other switches with Release 5.0 in the fabric over CFS. However, if an existing switch running NX-OS Release 4.x upgrades to Release 5.0, the new configurations will not be distributed to that switch as a CFS merge is not triggered on an upgrade. There are two options to upgrade:

- Apply new configuration only when all the switches in the fabric support them (Recommended option)
- Do an empty commit from an existing NX-OS Release 5.0 switch which has the new configuration

Periodic Inventory Notification

You can configure the switch to periodically send a message with an inventory of all software services currently enabled and running on the switch along with hardware inventory information. The inventory is modified each time the switch is restarted nondisruptively.

Duplicate Message Throttle

You can configure a throttling mechanism to limit the number of Call Home messages received for the same event. If the same message is sent multiple times from the switch within a short period of time, you may be swamped with a large number of duplicate messages.

Call Home Configuration Distribution

You can enable fabric distribution for all Cisco MDS 9000 Series switches and Cisco Nexus 5000 Series switches in the fabric. When you perform Call Home configurations, and distribution is enabled, that configuration is distributed to all the switches in the fabric. However, the switch priority and the Syscontact names are not distributed.

You automatically acquire a fabric-wide lock when you enter the first configuration command operation after you enable distribution in a switch. The Call Home application uses the effective and pending database model to store or commit the configuration changes. When you commit the configuration changes, the effective database is overwritten by the configuration changes in the pending database and all the switches in the fabric receive the same configuration. After making the configuration changes, you can choose to discard the changes by terminating the changes instead of committing them. In either case, the lock is released. See Using the CFS Infrastructure, on page 7 for more information on the CFS application.

Note The switch priority and the Syscontact name are not distributed.

Fabric Lock Override

If you have performed a Call Home task and have forgotten to release the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your changes to the pending database are discarded and the fabric lock is released.

Tip The changes are only available in the volatile directory and are subject to being discarded if the switch is restarted.

Clearing Call Home Name Server Database

When the Call Home name server database is full, a new entry cannot be added. The device is not allowed to come online. To clear the name server database, increase the database size or perform a cleanup by removing unused devices. A total of 20,000 name server entries are supported.

EMC Email Home Delayed Traps

DCNM-SAN can be configured to generate EMC Email Home XML email messages. In SAN-OS Release 3.x or earlier, DCNM-SAN listens to interface traps and generates EMC Email Home email messages. Link traps are generated when an interface goes to down from up or vice versa. For example, if there is a scheduled server reboot, the link goes down and DCNM-SAN generates an email notification.

Cisco NX-OS Release 4.1(3) provides the ability to generate a delayed trap so that the number of generated email messages is reduced. This method filters server reboots and avoids generating unnecessary EMC Email Home email messages. In NX-OS Release 4.1(3), users have the ability to select the current existing feature or this new delayed trap feature.

Event Triggers

This section discusses Call Home trigger events. Trigger events are divided into categories, with each category assigned CLI commands to execute when the event occurs. The command output is included in the transmitted message. Table 8: Event Triggers , on page 56 lists the trigger events.

Table 8: Event Triggers

Event	Alert Group	Event Name	Description	Call Home Message Level
Call Home	System and CISCO_TAC	SW_CRASH	A software process has crashed with a stateless restart, indicating an interruption of a service.	5
Call Home	System and CISCO_TAC	CRASH_PROC	A software process has crashed with a stateless restart, indicating an interruption of a service.	5
Call Home	System and CISCO_TAC	SW_SYSTEM_INCONSISTENT	Inconsistency detected in software or file system.	5
Call Home	Environmental and CISCO_TAC	TEMPERATURE_ALARM	Thermal sensor indicates temperature reached operating threshold.	6
	Environmental and CISCO_TAC	POWER_SUPPLY_FAILURE	Power supply failed.	6
	Environmental and CISCO_TAC	FAN_FAILURE	Cooling fan has failed.	5

Event	Alert Group	Event Name	Description	Call Home Message Level
Call Home	Line Card Hardware and CISCO_TAC	LINECARD_FAILURE	Line card hardware operation failed.	7
	Line Card Hardware and CISCO_TAC	POWER_UP_DIAGNOSTICS_FAILURE	Line card hardware failed power-up diagnostics.	7
Call Home	Line Card Hardware and CISCO_TAC	PORT_FAILURE	Hardware failure of interface port(s).	6
Call Home	Line Card Hardware, Supervisor Hardware, and CISCO_TAC	BOOTFLASH_FAILURE	Failure of boot compact flash card.	6
Call Home	Supervisor Hardware and CISCO_TAC	NVRAM_FAILURE	Hardware failure of NVRAM on supervisor hardware.	6
Call Home	Supervisor Hardware and CISCO_TAC	FREEDISK_FAILURE	Free disk space is below a threshold on supervisor hardware.	6
Call Home	Supervisor Hardware and CISCO_TAC	SUP_FAILURE	Supervisor hardware operation failed. **Note** When the active supervisor is removed, a switch over occurs. A call home notification for this event will not be sent.	7
		POWER_UP_DIAGNOSTICS_FAILURE	Supervisor hardware failed power-up diagnostics.	7
Call Home	Supervisor Hardware and CISCO_TAC	INBAND_FAILURE	Failure of in-band communications path.	7
Call Home	Supervisor Hardware and CISCO_TAC	EOBC_FAILURE	Ethernet out-of-band channel communications failure.	6
Call Home	Supervisor Hardware and CISCO_TAC	MGMT_PORT_FAILURE	Hardware failure of management Ethernet port.	5
	License	LICENSE_VIOLATION	Feature in use is not licensed, and are turned off after grace period expiration.	6

Event	Alert Group	Event Name	Description	Call Home Message Level
Inventory	Inventory and CISCO_TAC	COLD_BOOT	Switch is powered up and reset to a cold boot sequence.	2
		HARDWARE_INSERTION	New piece of hardware inserted into the chassis.	2
		HARDWARE_REMOVAL	Hardware removed from the chassis.	2
Test	Test and CISCO_TAC	TEST	User generated test.	2
Port syslog	Syslog-group-port	SYSLOG_ALERT	Syslog messages corresponding to the port facility.	2
RMON	RMON	RMON_ALERT	RMON alert trigger messages.	2

Call Home Message Levels

Table 9: Event Categories and Executed Commands

Event Category	Description	Executed Commands
System show module show version show tech-support platform show tech-support sysmgr show hardware show sprom all	Events generated by failure of a software system that is critical to unit operation.	**show tech-supportshow system redundancy status**
Environmental show module show version show environment show logging logfile | tail -n 200	Events related to power, fan, and environment sensing elements such as temperature alarms.	**show moduleshow environment**

Event Category	Description	Executed Commands
Line Card Hardware show module show version show tech-support platform show tech-support sysmgr show hardware show sprom all	Events related to standard or intelligent line card hardware.	**show tech-support**
Supervisor Hardware show module show version show tech-support platform show tech-support sysmgr show hardware show sprom all	Events related to supervisor modules.	**show tech-support**
Inventory show module show version show hardware show inventory show system uptime show sprom all show license usage	Inventory status is provided whenever a unit is cold booted, or when FRUs are inserted or removed. This is considered a noncritical event, and the information is used for status and entitlement.	**show version**
Test show module show version	User generated test message.	**show version**

Call Home messages (sent for syslog alert groups) have the syslog severity level mapped to the Call Home message level (see the Syslog-Based Alerts, on page 53).

This section discusses the severity levels for a Call Home message when using one or more switches in the Cisco MDS 9000 Series and the Cisco Nexus 5000 Series. Call Home message levels are preassigned per event type.

Severity levels range from 0 to 9, with 9 having the highest urgency. Each syslog level has keywords and a corresponding syslog level as listed in Table 10: Severity and Syslog Level Mapping , on page 60.

Note Call Home does not change the syslog message level in the message text. The syslog message texts in the Call Home log appear as they are described in the *Cisco MDS 9000 Series System Messages* Reference.

Note Call Home severity levels are not the same as system message logging severity levels (see the *Cisco MDS 9000 Series System Messages* Reference).

Table 10: Severity and Syslog Level Mapping

Call Home Level	Keyword Used	Syslog Level	Description
Catastrophic (9)	**Catastrophic**	N/A	Network wide catastrophic failure.
Disaster (8)	**Disaster**	N/A	Significant network impact.
Fatal (7)	**Fatal**	Emergency (0)	System is unusable.
Critical (6)	**Critical**	Alert (1)	Critical conditions, immediate attention needed.
Major (5)	**Major**	Critical (2)	Major conditions.
Minor (4)	**Minor**	Error (3)	Minor conditions.
Warning (3)	**Warning**	Warning (4)	Warning conditions.
Notify (2)	**Notification**	Notice (5)	Basic notification and informational messages. Possibly independently insignificant.
Normal (1)	**Normal**	Information (6)	Normal event signifying return to normal state.
Debug (0)	**Debugging**	Debug (7)	Debugging messages.

Message Contents

The following contact information can be configured on the switch:

- Name of the contact person
- Phone number of the contact person
- Email address of the contact person
- Mailing address to which replacement parts must be shipped, if required
- Site ID of the network where the site is deployed
- Contract ID to identify the service contract of the customer with the service provider

Table 11: Short Text Messages , on page 61 describes the short text formatting option for all message types.

Table 11: Short Text Messages

Data Item	Description
Device identification	Configured device name
Date/time stamp	Time stamp of the triggering event
Error isolation message	Plain English description of triggering event
Alarm urgency level	Error level such as that applied to system message

Table 12: Reactive Event Message Format , on page 61, Table 13: Inventory Event Message Format , on page 64, and Table 14: User-Generated Test Message Format , on page 66 display the information contained in plain text and XML messages.

Table 12: Reactive Event Message Format

Data Item (Plain text and XML)	Description(Plain text and XML)	XML Tag (XML only)
Time stamp	Date and time stamp of event in ISO time notation: *YYYY-MM-DD*T*HH:MM:SS*. **Note** The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time - ch:EventTime
Message name	Name of message. Specific event names are listed in the Event Triggers, on page 56.	/mml/header/name
Message type	Specifically “Call Home.”	/mml/header/type - ch:Type
Message group	Specifically “reactive.”	/mml/header/group
Severity level	Severity level of message (see Table 10: Severity and Syslog Level Mapping , on page 60).	/mml/header/level - aml-block:Severity
Source ID	Product type for routing.	/mml/header/source - ch:Series

Data Item (Plain text and XML)	Description(Plain text and XML)	XML Tag (XML only)
Device ID	Unique device identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format is *type@Sid@seria* l, where: • *type* is the product model number from backplane SEEPROM. • **@** is a separator character. • *Sid* is "C," identifying the serial ID as a chassis serial number. • *serial* is the number identified by the Sid field. Example: DS-C9509@C@12345678	/mml/ header/deviceId
Customer ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/header/customerID - ch:CustomerId
Contract ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/header/contractId - ch:ContractId>
Site ID	Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/header/siterId - ch:SiteId
Server ID	If the message is generated from the fabric switch, it is the unique device identifier (UDI) of the switch. Format is *type@Sid@serial* , where: • *type* is the product model number from backplane SEEPROM. • **@** is a separator character. • *Sid* is "C," identifying the serial ID as a chassis serial number. • *serial* is the number identified by the Sid field. Example: DS-C9509@C@12345678	/mml/header/serverId - -blank-
Message description	Short text describing the error.	/mml/body/msgDesc - ch:MessageDescription
Device name	Node that experienced the event. This is the host name of the device.	/mml/body/sysName - ch:SystemInfo/Name
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact - ch:SystemInfo/Contact

Data Item (Plain text and XML)	Description(Plain text and XML)	XML Tag (XML only)
Contact email	Email address of person identified as contact for this unit.	/mml/body/sysContacte-mail - ch:SystemInfo/Contact email
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhoneNumber - ch:SystemInfo/ContactPhoneNumber
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress - ch:SystemInfo/StreetAddress
Model name	Model name of the switch. This is the specific model as part of a product Series name.	/mml/body/chassis/name - rme:Chassis/Model
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo - rme:Chassis/SerialNumber
Chassis part number	Top assembly number of the chassis.	/mml/body/fru/partNo - rme:chassis/Card/PartNumber
Chassis hardware version	Hardware version of chassis.	/mml/body/chassis/hwVersion - rme:Chassis/HardwareVersion
Supervisor module software version	Top level software version.	/mml/body/fru/swVersion - rme:chassis/Card/SoftwareIdentity
Affected FRU name	Name of the affected FRU generating the event message.	/mml/body/fru/name - rme:chassis/Card/Model
Affected FRU serial number	Serial number of affected FRU.	/mml/body/fru/serialNo - rme:chassis/Card/SerialNumber
Affected FRU part number	Part number of affected FRU.	/mml/body/fru/partNo - rme:chassis/Card/PartNumber
FRU slot	Slot number of FRU generating the event message.	/mml/body/fru/slot - rme:chassis/Card/LocationWithinContainer
FRU hardware version	Hardware version of affected FRU.	/mml/body/fru/hwVersion - rme:chassis/Card/SoftwareIdentity
FRU software version	Software version(s) running on affected FRU.	/mml/body/fru/swVersion - rme:chassis/Card/SoftwareIdentity
Command output name	The exact name of the issued command.	/mml/attachments/attachment/name - aml-block:Attachment/Name
Attachment type	Specifically command output.	/mml/attachments/attachment/type - aml-block:Attachment type
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment/mime - aml-block:Attachment/Data encoding

Data Item (Plain text and XML)	Description(Plain text and XML)	XML Tag (XML only)
Command output text	Output of command automatically executed Table 9: Event Categories and Executed Commands , on page 58).	/mml/attachments/attachment/atdata - aml-block:Attachment/Data

Table 13: Inventory Event Message Format

Data Item(Plain text and XML)	Description(Plain text and XML)	XML Tag(XML only)
Time stamp	Date and time stamp of event in ISO time notation: *YYYY-MM-DDTHH:MM:SS*. **Note** The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time - ch:EventTime
Message name	Name of message. Specifically "Inventory Update" Specific event names are listed in the Event Triggers, on page 56.	/mml/header/name
Message type	Specifically "Inventory Update."	/mml/header/type - ch-inv:Type
Message group	Specifically "proactive."	/mml/header/group
Severity level	Severity level of inventory event is level 2 (see Table 10: Severity and Syslog Level Mapping , on page 60).	/mml/header/level - aml-block:Severity
Source ID	Product type for routing at Cisco. Specifically "MDS 9000."	/mml/header/source - ch-inv:Series
Device ID	Unique Device Identifier (UDI) for end device generating message. This field should empty if the message is non-specific to a fabric switch. Format is *type@Sid@serial*, where: • *type* is the product model number from backplane SEEPROM. • *@* is a separator character. • *Sid* is "C," identifying the serial ID as a chassis serial number. • *serial* is the number identified by the Sid field. Example: DS-C9509@C@12345678	/mml/ header /deviceId
Customer ID	Optional user-configurable field used for contact info or other ID by any support service.	/mml/header/customerID - ch-inv:CustomerId

Data Item(Plain text and XML)	Description(Plain text and XML)	XML Tag(XML only)
Contract ID	Optional user-configurable field used for contact info or other ID by any support service.	/mml/header/contractId - ch-inv:ContractId>
Site ID	Optional user-configurable field, can be used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/header/siterId - ch-inv:SiteId
Server ID	If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch. Format is *type@Sid@seria* l, where: • *type* is the product model number from backplane SEEPROM. • **@** is a separator character. • *Sid* is "C," identifying the serial ID as a chassis serial number. • *serial* is the number identified by the Sid field. Example: DS-C9509@C@12345678	/mml/header/serverId - -blank-
Message description	Short text describing the error.	/mml/body/msgDesc - ch-inv:MessageDescription
Device name	Node that experienced the event.	/mml/body/sysName - ch-inv:SystemInfo/Name
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact - ch-inv:SystemInfo/Contact
Contact email	Email address of person identified as contact for this unit.	/mml/body/sysContacte-mail - ch-inv:SystemInfo/Contact email
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhoneNumber - ch-inv:SystemInfo/ContactPhoneNumber
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress - ch-inv:SystemInfo/StreetAddress
Model name	Model name of the unit. This is the specific model as part of a product Series name.	/mml/body/chassis/name - rme:Chassis/Model
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo - rme:Chassis/SerialNumber
Chassis part number	Top assembly number of the chassis.	/mml/body/fru/partNo - rme:chassis/Card/PartNumber
Chassis hardware version	Hardware version of chassis.	/mml/body/fru/hwVersion - rme:chassis/Card/SoftwareIdentity

Data Item(Plain text and XML)	Description(Plain text and XML)	XML Tag(XML only)
Supervisor module software version	Top level software version.	/mml/body/fru/swVersion - rme:chassis/Card/SoftwareIdentity
FRU name	Name of the affected FRU generating the event message.	/mml/body/fru/name - rme:chassis/Card/Model
FRU s/n	Serial number of FRU.	/mml/body/fru/serialNo - rme:chassis/Card/SerialNumber
FRU part number	Part number of FRU.	/mml/body/fru/partNo - rme:chassis/Card/PartNumber
FRU slot	Slot number of FRU.	/mml/body/fru/slot - rme:chassis/Card/LocationWithinContainer
FRU hardware version	Hardware version of FRU.	/mml/body/fru/hwVersion - rme:chassis/Card/SoftwareIdentity
FRU software version	Software version(s) running on FRU.	/mml/body/fru/swVersion - rme:chassis/Card/SoftwareIdentity
Command output name	The exact name of the issued command.	/mml/attachments/attachment/name - aml-block:Attachment/Name
Attachment type	Specifically command output.	/mml/attachments/attachment/type - aml-block:Attachment type
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment/mime - aml-block:Attachment/Data encoding
Command output text	Output of command automatically executed after event categories (see Event Triggers, on page 56).	/mml/attachments/attachment/atdata - aml-block:Attachment/Data

Table 14: User-Generated Test Message Format

Data Item (Plain text and XML)	Description(Plain text and XML)	XML Tag(XML only)
Time stamp	Date and time stamp of event in ISO time notation: *YYYY-MM-DDTHH:MM:SS*. **Note** The time zone or daylight savings time (DST) offset from UTC has already been added or subtracted. T is the hardcoded limiter for the time.	/mml/header/time - ch:EventTime
Message name	Name of message. Specifically test message for test type message. Specific event names listed in the Event Triggers, on page 56).	/mml/header/name
Message type	Specifically “Test Call Home.”	/mml/header/type - ch:Type

Data Item (Plain text and XML)	Description(Plain text and XML)	XML Tag(XML only)
Message group	This field should be ignored by the receiving Call Home processing application, but may be populated with either “proactive” or “reactive.”	/mml/header/group
Severity level	Severity level of message, test Call Home message (see Table 10: Severity and Syslog Level Mapping , on page 60).	/mml/header/level - aml-block:Severity
Source ID	Product type for routing.	/mml/header/source - ch:Series
Device ID	Unique device identifier (UDI) for end device generating message. This field should empty if the message is nonspecific to a fabric switch. Format is *type@Sid@seria* l, where: • *type* is the product model number from backplane SEEPROM. • @ is a separator character. • *Sid* is “C” identifying the serial ID as a chassis serial number. • *serial* is the number identified by the Sid field. Example: DS-C9509@C@12345678	/mml/ header /deviceId
Customer ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/header/customerID - ch:CustomerId
Contract ID	Optional user-configurable field used for contract info or other ID by any support service.	/mml/header/contractId - ch:ContractId
Site ID	Optional user-configurable field used for Cisco-supplied site ID or other data meaningful to alternate support service.	/mml/header/siterId - ch:SiteId
Server ID	If the message is generated from the fabric switch, it is the Unique device identifier (UDI) of the switch. Format is *type@Sid@seria* l, where: • *type* is the product model number from backplane SEEPROM. • @ is a separator character. • *Sid* is “C” identifying the serial ID as a chassis serial number. • *serial* is the number identified by the Sid field. Example: “DS-C9509@C@12345678	/mml/header/serverId - -blank-
Message description	Short text describing the error.	/mml/body/msgDesc - ch:MessageDescription

Data Item (Plain text and XML)	Description(Plain text and XML)	XML Tag(XML only)
Device name	Switch that experienced the event.	/mml/body/sysName - ch:SystemInfo/Name
Contact name	Name of person to contact for issues associated with the node experiencing the event.	/mml/body/sysContact - ch:SystemInfo/Contact
Contact email	Email address of person identified as contact for this unit.	/mml/body/sysContacte-mail - ch:SystemInfo/Contact email
Contact phone number	Phone number of the person identified as the contact for this unit.	/mml/body/sysContactPhoneNumber - ch:SystemInfo/ContactPhoneNumber
Street address	Optional field containing street address for RMA part shipments associated with this unit.	/mml/body/sysStreetAddress - ch:SystemInfo/StreetAddress
Model name	Model name of the switch. This is the specific model as part of a product Series name.	/mml/body/chassis/name - rme:Chassis/Model
Serial number	Chassis serial number of the unit.	/mml/body/chassis/serialNo - rme:Chassis/SerialNumber
Chassis part number	Top assembly number of the chassis. For example, 800-xxx-xxxx.	/mml/body/fru/partNo - rme:chassis/Card/PartNumber
Command output text	Output of command automatically executed after event categories listed in Table 9: Event Categories and Executed Commands , on page 58.	/mml/attachments/attachment/atdata - aml-block:Attachment/Data
MIME type	Normally text or plain or encoding type.	/mml/attachments/attachment/mime - aml-block:Attachment/Data encoding
Attachment type	Specifically command output.	/mml/attachments/attachment/type - aml-block:Attachment type
Command output name	The exact name of the issued command.	/mml/attachments/attachment/name - aml-block:Attachment/Name

Guidelines and Limitations

Call Home Database Merger Guidelines

When merging two Call Home databases, follow these guidelines:

- Be aware that the merged database contains the following information:
    - A superset of all the destination profiles from the dominant and subordinate switches that take part in the merge protocol.

- The email addresses and alert groups for the destination profiles.
- Other configuration information (for example, message throttling, periodic inventory) from the switch that existed in the dominant switch before the merge.

See the CFS Merge Support, on page 11 for detailed concepts.

Call Home Configuration Guidelines

When configuring Call Home, follow these guidelines:

- An email server and at least one destination profile (predefined or user-defined) must be configured. The destination profiles used depends on whether the receiving entity is a pager, email, or automated service such as Cisco Smart Call Home.
- Switches can forward events (SNMP traps/informs) up to 10 destinations.
- The contact name (SNMP server contact), phone, and street address information must be configured before Call Home is enabled. This configuration is required to determine the origin of messages received.
- The Cisco MDS 9000 Series switch and the Cisco Nexus 5000 Series switch must have IP connectivity to an email server.
- If Cisco Smart Call Home is used, an active service contract must cover the device being configured.

Default Settings

Table 15: Default Call Home Default Settings , on page 69 lists the default Call Home settings.

Table 15: Default Call Home Default Settings

Parameters	Default
Destination message size for a message sent in full text format.	500,000
Destination message size for a message sent in XML format.	500,000
Destination message size for a message sent in short text format.	4000
DNS or IP address of the SMTP server to reach the server if no port is specified.	25
Alert group association with profile.	All
Format type.	XML
Call Home message level.	0 (zero)
HTTP proxy server use.	Disabled and no proxy server configured.
HTTP proxy server message size for full text destination.	1 MB

Parameters	Default
HTTP proxy server message size for XML.	1 MB

Configuring Call Home

Task Flow for Configuring Call Home

Follow these steps to configure Call Home:

Procedure

Step 1 Configure contact information.

Step 2 Enable or disable Call Home.

Step 3 Configure destination profiles.

Step 4 Associate one or more alert groups to each profile as required by your network. Customize the alert groups, if desired.

Step 5 Configure email options.

Step 6 Test Call Home messages.

Configuring Contact Information

Switch priority is specific to each switch in the fabric. This priority is used by the operations personnel or TAC support personnel to decide which Call Home message they should respond to first. You can prioritize Call Home alerts of the same severity from each switch.

To assign the contact information, follow these steps:

Before you begin

Each switch must include email, phone, and street address information. You can optionally include the contract ID, customer ID, site ID, and switch priority information.

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Configure the SNMP contact name:

switch(config)# **snmp-server contact personname@companyname.com**

Step 3 Enter the Call Home configuration submode:

switch(config)# **callhome**

switch(config-callhome)#

Step 4 Assign the customer's email address. Up to 128 alphanumeric characters are accepted in email address format:

switch(config-callhome)# **email-contact username@company.com**

Note You can use any valid email address. You cannot use spaces.

Step 5 Assign the customer's phone number. Up to 20 alphanumeric characters are accepted in international format:

switch(config-callhome)# **phone-contact +1-800-123-4567**

Note You cannot use spaces. Be sure to use the + prefix before the number.

Step 6 Assign the customer's street address where the equipment is located. Up to 256 alphanumeric characters are accepted in free format:

switch(config-callhome)# **streetaddress 1234 Picaboo Street, Any city, Any state, 12345**

Step 7 Assign the switch priority, with 0 being the highest priority and 7 the lowest:

switch(config-callhome)# **switch-priority 0**

Tip Use this field to create a hierarchical management structure.

Step 8 (Optional) Identify the customer ID:

switch(config-callhome)# **customer-id Customer1234**

Up to 256 alphanumeric characters are accepted in free format.

Step 9 (Optional) Identify the customer site ID:

switch(config-callhome)# **site-id Site1ManhattanNY**

Up to 256 alphanumeric characters are accepted in free format.

Step 10 Assign the customer ID for the switch:

switch(config-callhome)# **contract-id Company1234**

Up to 64 alphanumeric characters are accepted in free format.

Configuring Contact Information Using DCNM-SAN

To assign the contact information using DCNM-SAN, follow these steps:

Procedure

Step 1 Expand **Events** and select **Call Home** from the Physical Attributes pane.

You see the Call Home tabs in the Information pane.

Step 2 In Device Manager, click **Admin** > **Events** > **Call Home**.

Step 3 Click the **General** tab, then assign contact information and enable the Call Home feature. Call Home is not enabled by default. You must enter an email address that identifies the source of Call Home notifications.

Step 4 Click the **Destination(s)** tab to configure the destination email addresses for Call Home notifications. You can identify one or more email addresses that will receive Call Home notifications.

Note Switches can forward events (SNMP traps/informs) up to 10 destinations.

a. Click the **Create** tab to create a new destination. You will see the create destination window.

b. Enter the profile name, ID, and type of destination. You can select email or http in the Type field.

   If you select email, you can enter the email address in the EmailAddress field. The HttpUrl field is disabled.

   If you select http, you can enter the HTTP URL in the HttpUrl field. The EmailAddress field is disabled.

c. Click **Create** to complete the destination profile creation.

Step 5 Click the **e-mail Setup** tab to identify the SMTP server. Identify a message server to which your switch has access. This message server will forward the Call Home notifications to the destinations.

Step 6 In DCNM-SAN, click the **Apply Changes** icon. In Device Manager, click **Apply**.

Enabling Call Home Function

Once you have configured the contact information, you must enable the Call Home function.

To enable the Call Home function, follow these steps:

Procedure

Step 1 Enter configuration mode:

```
switch# configure terminal
```

Step 2 Enter Call Home configuration submode:

```
switch(config)# callhome
```

Step 3 Enable the Call Home function:

```
switch(config-callhome)# enable
Call Home enabled successfully
```

Step 4 (Optional) Disable the Call Home function:

```
switch(config-callhome)# disable
```

Note Even if Call Home is disabled, basic information for each Call Home event is sent.

When you disable the Call Home function, all input events are ignored.

Enabling Call Home Function Using DCNM-SAN

To enable the Call Home function using DCNM-SAN, follow these steps:

Procedure

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Events** and select **Call Home** in the Physical Attributes pane.

You see the Call Home information in the Information pane.

Step 3 Click the **Control** tab.

Step 4 Select a switch in the information pane.

Step 5 Check the **Duplicate Message Throttle** check box.

Step 6 Click the **Apply Changes** icon.

Configuring a Destination Profile

A destination profile contains the required delivery information for an alert notification. Destination profiles are typically configured by the network administrator. You can configure the following attributes for a destination profile:

- Profile name—A string that uniquely identifies each user-defined destination profile and is limited to 32 alphanumeric characters. The format options for a user-defined destination profile are full-txt, short-txt, or XML (default).
- Destination address—The actual address, pertinent to the transport mechanism, to which the alert should be sent.
- Message formatting—The message format used for sending the alert (full text, short text, or XML).

Note If you use the Cisco Smart Call Home service, the XML destination profile is required.

To configure predefined destination profile messaging options, follow these steps:

Note Steps 3, 4, and 5 in this procedure can be skipped or configured in any order.

Before you begin

At least one destination profile is required. You can configure multiple destination profiles of one or more types. You can use one of the predefined destination profiles or define a desired profile. If you define a new profile, you must assign a profile name.

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter the Call Home configuration submode:

switch(config)# **callhome**

switch(config-callhome)#

Step 3 Configure an email address or maximum destination message size for the predefined full-txt-destination profile:

switch(config-callhome)# **destination-profile full-txt-destination** {**email-addr** *email-address* | **message-size** *msg-size-in-bytes*}

The email addresses in this destination profile receives messages in full-txt format. The full-text format provides the complete, detailed explanation of the failure.

Tip Use a standard email address that does not have any text size restrictions.

The valid range is 0 to 1,000,000 bytes and the default is 500,000. A value of 0 implies that a message of any size can be sent.

Note The maximum size of each individual attachment inside the message is 250,000 bytes. If any attachment is more than this maximum size, then the output captured in the attachment will be truncated.

Step 4 Configure an email address or maximum destination message size for the predefined short-txt-destination profile:

switch(config-callhome)# **destination-profile short-txt-destination** {**email-addr** *email-address* | **message-size** *msg-size-in-bytes*}

The email addresses in this destination profile receive messages in short-txt format. This format provides the basic explanation of the failure in the Call Home message.

Tip Use a pager-related email address for this option.

The valid range is 0 to 1,000,000 bytes and the default is 4000. A value of 0 implies that a message of any size can be sent.

Note The maximum size of each individual attachment inside the message is 250,000 bytes. If any attachment is more than this maximum size, then the output captured in the attachment will be truncated.

Step 5 Configure an email address or maximum destination message size for the predefined XML-destination profile:

switch(config-callhome)# **destination-profile XML-destination** {**email-addr** *email-address* | **message-size** *msg-size-in-bytes*}

The email addresses in this destination-profile receives messages in XML format. This format provides information that is compatible with Cisco Systems TAC support.

Tip Do not add a pager-related email address to this destination profile because of the large message size.

The valid range is 0 to 1,000,000 bytes and the default is 500,000. A value of 0 implies that a message of any size can be sent.

Note The maximum size of each individual attachment inside the message is 250,000 bytes. If any attachment is more than this maximum size, then the output captured in the attachment will be truncated.

Configuring a Predefined Destination Profile Using DCNM-SAN

To configure predefined destination profile messaging options using DCNM-SAN, follow these steps:

Procedure

Step 1 Expand **Events** and select **Call Home** in the Physical Attributes pane.

The **Destination** tab is disabled until you click the **Profiles** tab. The profiles have to be loaded for the destination tab to be populated.

Step 2 Click the **Profiles** tab in the Information pane.

You see the Call Home profiles for multiple switches.

Step 3 Set the profile name, message format, message size, and severity level.

Step 4 Click in the Alert Groups column and select or remove an alert group.

Step 5 Click **Apply Changes** icon to create this profile on the selected switches.

Configuring a New Destination Profile

To configure a new destination-profile (and related parameters), follow these steps:

Note Steps 4, 5, and 6 in this procedure can be skipped or configured in any order.

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter the Call Home configuration submode:

switch(config)# **callhome**

Step 3 Configure a new destination profile called test:

```
switch(config-callhome)# destination-profile test
```

Step 4 Configure the email address for the user-defined destination profile (test) sent in default XML format:

```
switch(config-callhome)# destination-profile test e-mail-addr email-address
```

Step 5 Configure a maximum message size for the destination email addresses in the user-defined destination profile (test) sent in default XML format:

```
switch(config-callhome)# destination-profile test message-size msg-size
```

The valid range is 0 to 1,000,000 bytes and the default is 500,000. A value of 0 implies that a message of any size can be sent.

Step 6 Configure message-format for the user-defined destination profile (test) to be full text or short text format:

```
switch(config-callhome)# destination-profile test format {full-txt | short-txt}
```

Configuring a New Destination Profile Using DCNM-SAN

To configure a new destination-profile (and related parameters) using DCNM-SAN, follow these steps:

Procedure

Step 1 Expand **Events** and select **Call Home** in the Physical Attributes pane.

Note The **Destination** tab is disabled until you click the **Profiles** tab. The profiles have to be loaded for the destination tab to be populated.

Step 2 Click the **Profiles** tab in the Information pane.

You see Call Home profiles for multiple switches.

Step 3 Click the **Create Row** icon to add a new profile.

Step 4 Set the profile name, message format, size, and severity level.

Step 5 Click an alert group and select each group that you want sent in this profile.

Step 6 Click a transport method. You can select **email**, **http** or **emailandhttp**.

Step 7 Click **Create** to create this profile on the selected switches.

Associating an Alert Group with a Destination Profile

Different types of Call Home alerts are grouped into different alert groups depending on their type. You can associate one or more alert groups to each profile as required by your network.

The alert group feature allows you to select the set of Call Home alerts to be received by a destination profile (either predefined or user-defined). You can associate multiple alert groups with a destination profile.

To associate an alert group with a destination profile, follow these steps:

Before you begin

A Call Home alert is sent to email destinations in a destination profile only if that Call Home alert belongs to one of the alert groups associated with that destination profile.

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter the Call Home configuration submode:

switch(config)# **callhome**

switch(config-callhome)#

Step 3 (Optional) Configure user-defined destination profile (test1) or predefined short-text destination profile to receive all user-generated Call Home test notifications:

switch(config-callhome)# **destination-profile** {*test1* | **short-txt-destination**} **alert-group** *test*

Step 4 (Optional) Configure user-defined destination profile (test1) to receive Call Home notifications for all events or predefined short-text destination profile to receive Call Home notifications for default events:

switch(config-callhome)# **destination-profile** {*test1* | **short-txt-destination**} **alert-group all**

Step 5 (Optional) Configure user-defined destination profile (test1) or predefined short-text destination profile to receive Call Home notifications for events that are meant only for Cisco TAC or the auto-notify service:

switch(config-callhome)# **destination-profile** {*test1* | **xml-destination**} **alert-group Cisco-TAC**

Step 6 (Optional) Configure user-defined destination profile (test1) or predefined short-text destination profile to receive Call Home notifications for software crash events:

switch(config-callhome)# **destination-profile** {*test1* | **xml-destination**} **alert-group Crash**

Step 7 (Optional) Configure user-defined destination profile (test1) or predefined short-text destination profile to receive Call Home notifications for power, fan, and temperature-related events:

switch(config-callhome)# **destination-profile** {*test1* | **short-txt-destination**} **alert-group environmental**

Step 8 (Optional) Configure user-defined destination profile (test1) or predefined short-text destination profile to receive Call Home notifications for inventory status events:

switch(config-callhome)# **destination-profile** {*test1* | **short-txt-destination**} **alert-group inventory**

Step 9 (Optional) Configure user-defined destination profile (test1) or predefined short-text destination profile to receive Call Home notifications for licensing events:

switch(config-callhome)# **destination-profile** {*test1* | **short-txt-destination**} **alert-group License**

Step 10 (Optional) Configure user-defined destination profile (test1) or predefined short-text destination profile to receive Call Home notifications for module-related events:

switch(config-callhome)# **destination-profile** {*test1* | **short-txt-destination**} **alert-group linecard-hardware**

Step 11 (Optional) Configure user-defined destination profile (test1) or predefined short-text destination profile to receive Call Home notifications for supervisor-related events:

switch(config-callhome)# **destination-profile** {*test1* | **short-txt-destination**} **alert-group supervisor-hardware**

Step 12 (Optional) Configure user-defined destination profile (test1) or predefined short-text destination profile to receive Call Home notifications for software-related events:

switch(config-callhome)# **destination-profile** {*test1* | **short-txt-destination**} **alert-group system**

Associating Alert Group Using DCNM-SAN

To associate an alert group with a destination profile using DCNM-SAN, follow these steps:

Procedure

Step 1 Expand **Events** and select **Call Home** in the Physical Attributes pane.

Step 2 Click the **Profiles** tab in the Information pane.

You see the Call Home profiles for multiple switches.

Step 3 Click the **Alert Groups** column in the row for the profile you want to associate.

You see the alert groups drop-down menu.

Step 4 Click an alert group to select it for association.

Step 5 You see a check next to that alert group.

To deselect it and remove the check, click it again.

Step 6 Click the **Apply Changes** icon.

Customizing Alert Group Messages

To assign show commands to be executed when an alert is sent, you must associate the commands with the alert group. When an alert is sent, Call Home associates the alert group with an alert type and attaches the output of the show commands to the alert message.

Note Make sure the destination profiles for a non-Cisco-TAC alert group, with a predefined show command, and the Cisco-TAC alert group are not the same.

To customize Call Home alert group messages, follow these steps:

Before you begin

- You can assign a maximum of five user-defined show commands to an alert group. Only show commands can be assigned to an alert group.
- Customized show commands are only supported for full text and XML alert groups. Short text alert groups (short-txt-destination) do not support customized show commands because they only allow 128 bytes of text.

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter the Call Home configuration submode:

switch(config)# **callhome**

switch(config-callhome)#

Step 3 Configure a user-defined **show** command for an alert group license:

switch(config-callhome)# **alert-group license user-def-cmd show license usage**

Note Only valid **show** commands are accepted.

Step 4 (Optional) Remove the user-defined **show** command from the alert group:

switch(config-callhome)# **no alert-group license user-def-cmd show license usage**

Customizing Alert Group Messages Using DCNM-SAN

To customize Call Home alert group messages using DCNM-SAN, follow these steps:

Procedure

Step 1 Expand **Events** and select **Call Home** in the Physical Attributes pane.

Step 2 Click the **User Defined Command** tab in the Information pane.

You see the User Defined Command information.

Step 3 Click the **Create Row** icon.

Step 4 Check the check boxes in front of the switches from which you want to receive alerts.

Step 5 Select the alert grouptype from the Alert Group Type drop-down list.

Step 6 Select the ID (1-5) of the CLI command. The ID is used to keep track of the messages.

Step 7 Enter the CLI **show** command in the **CLI Command** field.

Step 8 Click **Create**.

Step 9 Repeat Step 3 through Step 7 for each command you want to associate with the profile.

Step 10 Click **Close** to close the dialog box.

Setting the Call Home Message Levels

To set the message level for each destination profile for Call Home, follow these steps:

Before you begin

The urgency level ranges from 0 (lowest level of urgency) to 9 (highest level of urgency), and the default is 0 (all messages are sent).

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter the Call Home configuration submode:

switch(config)# **callhome**

Step 3 (Optional) Configure the message level urgency as 5 (level) and above for the user-defined profile (test1):

switch(config-callhome)# **destination-profile** *test* **message-level** *level*

Step 4 Remove a previously configured urgency level and reverts it to the default of 0 (all messages are sent):

switch(config-callhome)# **no destination-profile** *oldtest* **message-level** *level*

Configuring the Syslog-Based Alerts

To configure the syslog-group-port alert group, follow these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter the Call Home configuration submode:

switch(config)# **callhome**

switch(config-callhome)#

Step 3 Configure the predefined destination profile (short-txt-destination) to receive Call Home Notifications corresponding to syslog messages for the port facility:

switch(config-callhome)# **destination-profile short-txt-destination alert-group syslog-group-port**

Step 4 (Optional) Configure the predefined destination-profile (short-txt-destination) to send a Call Home message for syslog messages whose severity levels map to Call Home severity level of 5 or greater:

switch(config-callhome)# **destination-profile** *short-txt-destination* **message-level** *level*

The default is message level 0 (all syslog messages).

Configuring the Syslog-Based Alerts Using DCNM-SAN

To configure the syslog-group-port alert group using DCNM-SAN, follow these steps:

Procedure

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Events** and select **Call Home** in the Physical Attributes pane.

You see the Call Home information in the Information pane.

Step 3 Click the **Profiles** tab.

You see the Call Home profiles.

Step 4 Click the **Create Row** icon.

You see the Create Call Home Profile dialog box.

Step 5 Select the switches for which you want to send alerts.

Step 6 Enter the name of the profile in the Name field.

Step 7 Choose the message format, message size, and message severity level.

Step 8 Check the **syslogGroupPort** check box in the AlertGroups section.

Step 9 Click **Create** to create the profile for the syslog-based alerts.

Step 10 Close the dialog box.

Configuring RMON Alerts

To configure RMON alert groups, follow these steps:

Procedure

Step 1 Enter configuration mode:

switch# **config t**

Step 2 Enter Call Home configuration submode:

switch(config)# **callhome**

Step 3 (Optional) Configure a destination message profile (rmon_group) to send Call Home notifications for configured RMON messages:

switch(config-callhome)# **destination-profile**

Configuring RMON Alerts Using DCNM-SAN

To configure RMON alert groups using DCNM-SAN, follow these steps:

Procedure

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Events** and select **Call Home** in the Physical Attributes pane.

You will see the Call Home information in the Information pane.

Step 3 Click the **Profiles** tab.

You will see the Call Home profiles.

Step 4 Select the **Create Row** icon.

You will see the Create Call Home Profile dialog box.

Step 5 Select switches to send alerts.

Step 6 Enter the name of the profile.

Step 7 Select the message format, message size, and message severity level.

Step 8 Check the **RMON** check box in the AlertGroups section.

Step 9 Click **Create** to create the profile for the RMON-based alerts.

Step 10 Close the dialog box.

Configuring Event Trap Notifications

To configure a Call Home event notification trap (except Call Home periodic messages), follow these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter Call Home configuration submode:

switch(config)# **callhome**

Step 3 Enable the SNMP notification trap for Call Home:

switch(config-callhome)# **snmp-server enable**

Configuring General EMail Options

You can configure the from, reply-to, and return-receipt email addresses. While most email address configurations are optional, you must configure the SMTP server address for the Call Home functionality to work.

To configure general email options, follow these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter Call Home configuration submode:

switch(config)# **callhome**

Step 3 Configure the from email address:

switch(config-callhome)#**transport email from** *person@company.com*

Step 4 Configure the reply-to email address to which all responses should be sent:

switch(config-callhome)#**transport email reply-to** *person@company.com*

Step 5 Configure the DNS, IPv4 address, or IPv6 address of the SMTP server to reach the server:

switch(config-callhome)# **transport email smtp-server 192.168.1.1**

switch(config-callhome)# **transport email smtp-server 192.168.1.1 port 30**

The port usage defaults to 25 if no port is specified.

Note The port number is optional and, if necessary, may be changed depending on the server location.

Configuring General EMail Options Using DCNM-SAN

To configure general email options using DCNM-SAN, follow these steps:

Procedure

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Events** and select **Call Home** in the Physical Attributes pane.

You will see the Call Home information in the Information pane.

Step 3 Click the **e-mail Setup** tab.

Step 4 Select a switch in the Information pane.

Step 5 Enter the general email information.

Step 6 Enter the SMTP server IP address type, IP address or name, and port.

Step 7 Click the **Apply Changes** icon to update the email options.

Configuring HTTPS Support

Any predefined or user-defined destination profiles can be configured with the HTTPS URL address.

To configure the HTTPS URL address for any destination profile, follow these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter Call Home configuration submode:

switch(config)# **callhome**

Step 3 (Optional) Configure the predefined full-txt-destination profile with an HTTPS URL address:

switch(config-callhome)# **destination-profile full-txt-destination http**

The Call Home message in full-txt format is uploaded at the configured HTTPS URL address.

Step 4 (Optional) Configure the predefined CiscoTAC-1 profile with an HTTPS URL address:

switch(config-callhome)# **destination-profile CiscoTAC-1 http**

The Call Home message in XML format is uploaded at the configured HTTPS URL address.

Step 5 (Optional) Configure the user-defined destination profile with an HTTPS URL address:

switch(config-callhome)# **destination-profile test1 http**

The Call Home message in the configured format is uploaded at the configured HTTPS URL address.

Enable or Disable Transport Method

Any predefined or user-defined destination profiles can be configured to enable or disable a particular transport method. The transport methods are HTTP and email.

To enable or disable transport method for a destination profile, follow these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter Call Home configuration submode:

switch(config)# **callhome**

Step 3 (Optional) Enable predefined destination profile CiscoTAC-1 for HTTP transport method:

switch(config-callhome)# **destination-profile CiscoTAC-1 transport-method http**

Note For user-defined destination profiles, email is the default. You can enable either or both transport mechanisms. If you disable both methods, email will be enabled.

Step 4 (Optional) Disable predefined destination profile CiscoTAC-1 for email transport method:

switch(config-callhome)# **no destination-profile CiscoTAC-1 transport-method email**

Step 5 (Optional) Enable predefined full-txt-destination profile for HTTP transport method:

switch(config-callhome)# **destination-profile full-txt transport-method http**

Configuring an HTTP Proxy Server

Beginning with Cisco NX-OS Release 5.2, you can configure Smart Call Home to send HTTP messages through an HTTP proxy server. If you do not configure an HTTP proxy server, Smart Call Home sends HTTP messages directly to the Cisco Transport Gateway (TG).

To configure an HTTP proxy server, follow these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter Call Home configuration submode:

switch(config)# **callhome**

Step 3 Configure the HTTP proxy server domain name server (DNS) name, IPv4 address, or IPv6 address:

switch(config-callhome)# **transport http proxy server 192.0.2.1**

Optionally configures the port number. The port range is from 1 to 65535. The default port number is 8080.

Step 4 Enable Smart Call Home to send all HTTP messages through the HTTP proxy server:

switch(config-callhome)# **transport http proxy enable**

Note You can execute this command only after the proxy server address has been configured.

Step 5 (Optional) Display the transport-related configuration for Smart Call Home:

switch(config-callhome)# **show callhome transport**

Note The default value for full text destination and for XML is 1 MB.

Configuring an HTTP Proxy Server Using DCNM-SAN

To configure a Call Home HTTP proxy server using DCNM-SAN, follow these steps:

Procedure

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Events**, select **Call Home**, and **HTTP Proxy Server** in the Physical Attributes pane.
You will see the Call Home HTTP Proxy Server information in the Information pane.

Step 3 Click the **Address Type** tab.
The Address Type options are displayed.

Step 4 Click the **Address** tab and enter the address of the HTTP proxy server.

Step 5 Click the **Port** tab and enter a integer number to specify the port of the HTTP proxy server.

Step 6 Check the **Enable** check box to enable the HTTP proxy configured for Call Home.

Step 7 (Optional) Set an empty value in the **Address** tab to delete the HTTP proxy server from the MDS switch.

Step 8 Choose an address type. You can select ipv4, ipv6, or DNS.

Note If the address is empty, then no proxy server is configured.

Step 9 Click **Apply** to update HTTP Proxy Server options.

Configuring Call Home Wizard

Task Flow for Configuring Call Home Wizard

Follow these steps to configure the Call Home Wizard:

Procedure

Step 1 Configure contact information.

Step 2 Configure SMTP information.

Step 3 Configure the email source and destination information.

Step 4 Use CFS to populate the configuration data.

Step 5 Display the status.

Launching Call Home Wizard

To configure Call Home wizard, follow these steps:

Before you begin

- Enable the global CFS on the switch from DCNM-SAN configuration table.
- Clear the CFS lock on the switch.
- Check the merger status of CFS on the switch. If a merger failure is found, the wizard clears up the merge failure in the backend process while running the wizard.

Procedure

Step 1 Select a fabric in the logical domain tree.

Step 2 Select **Tools**, **Events** and **Call Home**.
The master switch pane is displayed.

Step 3 (Optional) You can also launch the Call Home wizard by clicking the **CallHome Wizard** icon in the Call Home **Control** tab.

Step 4 Select a **Master Switch** and click **Next**.
The contact information pane is displayed.

Step 5 Enter the **Contact**, **Phone Number**, **Email Address** and the **Street Address** information.

Note You must specify all of the four parameters before clicking Next.

Step 6 Click **Next**.
The Email Setup pane is displayed.

Step 7 In the **Email SMTP Servers** tab, enter the **Primary SNTP Server** address.

You can specify up to two SMTP servers if the master switch is version 5.0 or above. However, you cannot specify a secondary SMTP server if the master switch version is below 5.0.

The wizard creates new rows in the SMTP server table.

Step 8 In the **Destination** tab, click **Add** to enter the Call Home destinations.

You can enter up to three Call Home destinations.

Step 9 (Optional) Click **Remove** to delete a Call Home destination entry.

Step 10 From the drop-down list, select **Protocol** and **Profile**.
The Profile drop-down lists three default profiles: xml, short_txt and full_txt.

Step 11 Click **Finish** to configure the wizard.

All major configuration steps and failures are displayed in the Status Dialog window.

The **Status Dialog** window is displayed.

Step 12 Click **Run Test** to perform the Call Home test.

Step 13 Click **Yes** to test the command on all switches in the selected fabric or click **No** to close the window.

Configuring SMTP Server and Ports

To configure the SMTP server and port, follow these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter Call Home configuration submode:

switch(config)# **callhome**

Step 3 Configure the DNS, IPv4 address, or IPv6 address of the SMTP server to reach the server:

switch(config-callhome)# **transport email smtp-server 192.168.1.1**

switch(config-callhome)# **transport email smtp-server 192.168.1.1 port 30**

The port usage defaults to 25 if no port is specified.

Note The port number is optional and, if required, may be changed depending on the server location.

Configuring Multiple SMTP Server Support

To configure multiple SMTP server support, follow these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter Call Home configuration submode:

switch(config)# **callhome**

Step 3 Use one of the following commands:

- Distribute the SMTP server configuration to devices running software releases prior to NX-OS Release 5.0 and earlier:

  switch(config-callhome)# **transport email smtp-server**

- Distribute multiple SMTP server capability:

  switch(config-callhome)# [**no**] **transport email mail-server** {**ipv4** | **IPV6** | **hostname**} [**port** *number*] [**priority** *number*]

Example

Note The **transport email mail-server** command is distributed only to devices running Cisco NX-OS Release 5.0(1a) or later. The **transport email smtp-server** command is distributed only to devices running earlier software releases.

The following example shows how to configure multiple SMTP servers for Call Home messages:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# callhome
switch(config-callhome)# transport email mail-server 192.0.2.10 priority 4
switch(config-callhome)# transport email mail-server 172.21.34.193
switch(config-callhome)# transport email smtp-server 10.1.1.174
switch(config-callhome)# transport email mail-server 64.72.101.213 priority 60
switch(config-callhome)# transport email from person@company.com
switch(config-callhome)# transport email reply-to person@company.com
```

Based on the configuration above, the SMTP servers will be contacted in this order:

10.1.1.174 (priority 0)

192.0.2.10 (priority 4)

172.21.34.193 (priority 50 - default)

64.72.101.213 (priority 60)

Enabling Periodic Inventory Notifications

When you enable this feature without configuring an interval value, the Call Home message is sent every 7 days. This value ranges from 1 to 30 days. By default, this feature is disabled in all switches in the Cisco MDS 9000 Series and Cisco Nexus 5000 Series switches.

To enable periodic inventory notification in a Cisco MDS 9000 Series switch or a Cisco Nexus 5000 Series switch, follow these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter the Call Home configuration submode:

switch(config)# **callhome**

Step 3 Enable the periodic inventory notification feature:

switch(config-callhome)# **periodic-inventory notification**

Disable the periodic inventory notification feature (default):

switch(config-callhome)# **no periodic-inventory notification**

By default, the Call Home message is sent every 7 days.

Step 4 Configure the periodic inventory notification message to be sent every 15 days:

switch(config-callhome)# **periodic-inventory notification interval 15**

Default to using the factory default of sending a Call Home message every 7 days:

switch(config-callhome)# **no periodic-inventory notification interval 15**

This value ranges from 1 to 30 days.

Enabling Periodic Inventory Notifications Using DCNM-SAN

To enable periodic inventory notification in a Cisco MDS 9000 Series switch or a Cisco Nexus 5000 Series switch using DCNM-SAN, follow these steps:

Procedure

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Events** and select **Call Home** in the Physical Attributes pane.
You will see the Call Home information in the Information pane.

Step 3 Click the **Periodic Inventory** tab.
You will see the Call Home periodic inventory information.

Step 4 Select a switch in the Information pane.

Step 5 Check the **Enable** check box.

Step 6 Enter the number of days for which you want the inventory checked.

Step 7 Click the **Apply Changes** icon.

Configuring Duplicate Message Throttle

You can configure a throttling mechanism to limit the number of Call Home messages received for the same event. If the same message is sent multiple times from the switch within a short period of time, you may be swamped with a large number of duplicate messages.

Restrictions

- By default, this feature is enabled in all switches in the Cisco MDS 9000 Series and the Cisco Nexus 5000 Series switches. When enabled, if the number of messages sent exceeds the maximum limit of 30 messages within the 2-hour time frame, then additional messages for that alert type are discarded within that time frame. You cannot modify the time frame or the message counter limit.
- If 2 hours have elapsed since the first such message was sent and a new message has to be sent, then the new message is sent and the time frame is reset to the time when the new message was sent and the count is reset to 1.

To enable message throttling in a Cisco MDS 9000 Series switch or a Cisco Nexus 5000 Series switch, follow these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter the Call Home configuration submode:

switch(config)# **callhome**

Step 3 Disable the duplicate message throttling feature:

switch(config-callhome)# **no duplicate-message throttle**

Step 4 Enable the duplicate message throttling feature (default):

switch(config-callhome)# **duplicate-message throttle**

Configuring Duplicate Message Throttle Using DCNM-SAN

To enable message throttling in a Cisco MDS 9000 Series switch or a Cisco Nexus 5000 Series switch using DCNM-SAN, follow these steps:

Procedure

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Events** and select **Call Home** in the Physical Attributes pane.

You see the Call Home information in the Information pane.

Step 3 Click the **Control** tab.

Step 4 Select a switch in the Information pane.

Step 5 Check the **Duplicate Msg Throttle** check box.

Step 6 Click the **Apply Changes** icon.

Enabling Call Home Fabric Distribution

To enable Call Home fabric distribution, follow these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter the Call Home configuration submode:

switch(config)# **callhome**

Step 3 Enable Call Home configuration distribution to all switches in the fabric:

switch(config-callhome)# **distribute**

Acquires a fabric lock and stores all future configuration changes in the pending database.

Step 4 Disables (default) Call Home configuration distribution to all switches in the fabric.

switch(config-callhome)# **no distribute**

Committing the Call Home Configuration Changes

To commit the Call Home configuration changes, follow these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter the Call Home configuration submode:

switch(config)# **callhome**

Step 3 Distribute the configuration changes to all switches in the fabric and release the lock:

switch(config-callhome)# **commit**

Overwrites the effective database with the changes made to the pending database.

Discarding the Call Home Configuration Changes

To discard the Call Home configuration changes, follow these steps:

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enter the Call Home configuration submode:

switch(config)# **callhome**

Step 3 Discard the configuration changes in the pending database and release the fabric lock:

```
switch(config-callhome)# abort
```

Enabling Call Home Fabric Distribution Using DCNM-SAN

To enable Call Home fabric distribution using DCNM-SAN, follow these steps:

Procedure

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Events** and select **Call Home** in the Physical Attributes pane.

You see the Call Home information in the Information pane.

Step 3 Click the **CFS** tab.

You see the CFS information for Call Home.

Step 4 Select a switch in the Information pane.

Step 5 Select **Enable** from the drop-down list in the Admin column in the row for that switch.

Step 6 Click the **Apply Changes** icon to commit the changes.

Fabric Lock Override

To use administrative privileges and release a locked Call Home session, follow these steps:

Procedure

Use administrative privileges and release a locked Call Home session:

```
switch# clear callhome session
```

Call Home Communications Test

You can test Call Home communications by sending a test message to the configured destination(s) or sending a test inventory message to the configured destination(s).

Use the **test** command to simulate a message generation.

To test the Call Home function, follow these steps:

Procedure

Step 1 Send a test message to the configured destinations:

switch# **callhome test**

Step 2 Send a test inventory message to the configured destinations:

switch(config)# **callhome test inventory**

Call Home Communications Test Using DCNM-SAN

To test the Call Home function and simulate a message generation using DCNM-SAN, follow these steps:

Procedure

Step 1 Select a switch in the Fabric pane.

Step 2 Expand **Events** and select **Call Home** in the Physical Attributes pane.

You see the Call Home information in the Information pane.

Step 3 Click the **Test** tab.

You see the configured tests for the switch and the status of the last testing.

Step 4 Select a switch in the Information pane.

Step 5 From the **TestAction** drop-down list in the row for that switch, select **test** or **testWithInventory**

Step 6 Click the **Apply Changes** icon to run the test.

Table 16: EMC Call Home Traps, on page 94 includes all the traps for EMC Call Home.

Table 16: EMC Call Home Traps

SNMP Trap	Send EMC Call Home When
connUnitStatusChange	operStatus == failed(5)
cefcModuleStatusChange	operStatus != {ok(2), boot(5), selfTest(6), poweredUp(16), syncInProgress(21)}
cefcPowerStatusChange	operStatus = {offDenied(4), offEnvPower(5),offEnvTemp(6),offEnvFan(7),failed(8)}
cefcFRURemoved	all
cefcFanTrayStatusChange	all
cieDelayedLinkUpDown	operStatusReason != {linkFailure, adminDown, portGracefulShutdown}
cefcFRUInserted	all

entSensorThresholdNotification	value >= threshold

Configuring Delayed Traps

The server.callhome.delayedtrap.enable property is added to section 9 Call Home in the server.properties configuration file. The property file can enable the DCNM-SAN server to use delayed traps instead of regular link Down traps for EMC Email Home messages.

Enabling the Delayed Trap Feature

To enable the delayed trap feature, perform this task:

Before you begin

To enable this feature, you need to turn on delayed traps at switch level, and then set the server.callhome.delayedtrap.enable property in the server.properties configuration file to true. By default, the server.callhome.delayedtrap.enable option is disabled and regular linkDown traps are used.

Procedure

Step 1 Enter configuration mode:

switch# **configure terminal**

Step 2 Enable the system-delayed trap feature:

switch(config)# **system delayed-traps enable mode** *FX*

Step 3 Configure the system-delayed trap timeout value:

switch(config)# **system delayed-traps timer** *<1-60>*

If no value is entered, a default value of 4 minutes is used. You can choose any value between 1 to 60 minutes.

Enabling the Delayed Trap Feature Using DCNM-SAN

To enable delayed traps on switches running NX-OS Release 4.1(3) and later using DCNM-SAN, follow these steps:

Procedure

Step 1 Expand **Events** and select **SNMP Traps** in the Physical Attributes pane.

In the table above the map layout in DCNM-SAN, click the **Delayed Traps** tab.

Step 2 Check the **Enable** check box for the switches on which you want to enable delayed traps.

Step 3 Enter the **timer** value in the Delay column.

Step 4 Click **Apply** to save your changes.

Note If no value is entered, the default value of 4 minutes is used.

Enabling Delayed Traps Using Cisco Device Manager

To enable the delayed trap feature using Device Manager, follow these steps:

Procedure

Step 1 In Device Manager, choose **Admin** > **Events** > **Filters** > **Delayed Traps**.

You can see the Events Filters information in the Information pane.

Step 2 Click the **Delayed Traps** tab.

Step 3 Check the **Enable** check box to enable delayed traps.

Delay interval will only be available when the feature is enabled.

Step 4 To disable Delayed Traps, uncheck the **Enable** check box and click **Apply**.

Viewing Event Filter Notification

In Device Manager, choose Admin > Events > Filters to see the descriptive notification.

You can see the Event Filters information in the Information pane.

The Event Filters screen displays the descriptive information about the notification.

Verifying Call Home Configuration

To display the Call Home configuration information, perform one of the following tasks:

Displaying Call Home Information

Use the **show callhome** command to display the configured Call Home information.

Displays Configured Call Home Information

```
switch# show callhome

callhome enabled
Callhome Information:
contact person name:who@where
contact person's email:person@place.com
contact person's phone number:310-408-4000
street addr:1234 Picaboo Street, Any city, Any state, 12345
```

```
site id:Site1ManhattanNewYork
customer id:Customer1234
contract id:Cisco1234
switch priority:0
```

Displays Information for All Destination Profiles (Predefined and User-Defined)

```
switch# show callhome destination-profile

XML destination profile information
maximum message size:500000
message format:XML
message-level:0
email addresses configured:
alert groups configured:
cisco_tac
test destination profile information
maximum message size:100000
message format:full-txt
message-level:5
email addresses configured:
admin@yourcompany.com
alert groups configured:
test
full-txt destination profile information
maximum message size:500000
message format:full-txt
message-level:0
email addresses configured:
alert groups configured:
all
short-txt destination profile information
maximum message size:4000
message format:short-txt
message-level:0
email addresses configured:
alert groups configured:
all
```

Displays Information for a User-defined Destination Profile

```
switch# show callhome destination-profile
 test
test destination profile information
maximum message size:100000
message format:full-txt
message-level:5
email addresses configured:
user
@
company
.com
alert groups configured:
test
```

Displays the Full-Text Profile

```
switch# show callhome destination-profile profile full-txt-destination

full-txt destination profile information
```

```
maximum message size:250000
email addresses configured:
person2@company2.com
```

Displays the Short-Text Profile

```
switch# show callhome destination-profile profile short-txt-destination
Short-txt destination profile information
maximum message size:4000
email addresses configured:
person2@company2.com
```

Displays the XML Destination Profile

```
switch# show callhome destination-profile profile XML-destination
XML destination profile information
maximum message size:250000
email addresses configured:
findout@.cisco.com
```

Displays EMail and SMTP Information

```
switch# show callhome transport-email
from email addr:user@company1.com
reply to email addr:pointer@company.com
return receipt email addr:user@company1.com
smtp server:server.company.com
smtp server port:25
```

Displays the running configuration callhome information

```
switch# show running-config callhome
!Command: show running-config callhome
!Time: Tue Sep  9 12:16:45 2014
version 6.2(9)
logging level callhome 5
callhome
  contract-id contact1
  customer-id cust1
  site-id Site1
  email-contact sakpuri@cisco.com
  phone-contact +1-800-000-0000
  streetaddress 12345 Cisco Way, San Jose, CA
  destination-profile Inventory
  destination-profile Inventory format full-txt
  destination-profile Inventory message-size 1000000
  destination-profile Service
  destination-profile Service format full-txt
  destination-profile Service message-size 1000000
  destination-profile dest1
  destination-profile dest1 format XML
  destination-profile dest1 message-size 500000
  destination-profile full_txt message-size 1000000
  destination-profile httpProf
  destination-profile httpProf format XML
  destination-profile httpProf message-size 0
  destination-profile short_txt message-size 4000
  destination-profile xml message-size 1000000
  destination-profile xml message-size 1000000
  destination-profile Inventory email-addr sakpuri@cisco.com
  destination-profile Service email-addr sakpuri@cisco.com
  destination-profile full_txt email-addr sakpuri@cisco.com
  destination-profile short_txt email-addr sakpuri@cisco.com
  destination-profile xml email-addr sakpuri@cisco.com
  destination-profile Service alert-group environmental
  destination-profile xml alert-group environmental
  destination-profile Inventory alert-group inventory
  destination-profile xml alert-group inventory
  destination-profile Service alert-group linecard-hardware
```

Displays the running configuration for callhome with defaults

```
switch# show running-config callhome all
EG-9506-1-176# show running-config callhome all
!Command: show running-config callhome all
!Time: Tue Sep  9 12:18:22 2014
version 6.2(9)
logging level callhome 5
callhome
  contract-id contact1
  customer-id cust1
  switch-priority 7
  site-id Site1
  email-contact sakpuri@cisco.com
  phone-contact +1-800-000-0000
  streetaddress 12345 Cisco Way, San Jose, CA
  destination-profile Inventory
  destination-profile Inventory format full-txt
  destination-profile Inventory transport-method email
  no destination-profile Inventory transport-method http
  destination-profile Inventory message-size 1000000
  destination-profile Inventory message-level 0
  destination-profile Service
  destination-profile Service format full-txt
  destination-profile Service transport-method email
  no destination-profile Service transport-method http
  destination-profile Service message-size 1000000
  destination-profile Service message-level 0
  destination-profile dest1
  destination-profile dest1 format XML
  destination-profile dest1 transport-method email
  no destination-profile dest1 transport-method http
  destination-profile dest1 message-size 500000
  destination-profile dest1 message-level 0
  destination-profile full_txt
  destination-profile full_txt format full-txt
  destination-profile full_txt transport-method email
  no destination-profile full_txt transport-method http
  destination-profile full_txt message-size 1000000
  destination-profile full_txt message-level 0
  destination-profile httpProf
```

Displays the startup configuration for callhome

```
switch# show startup-config callhome

!Command: show startup-config callhome
!Time: Tue Sep  9 12:19:27 2014
!Startup config saved at: Fri Sep  5 12:13:53 2014
```

```
version 6.2(9)
logging level callhome 5
callhome
  contract-id contact1
  customer-id cust1
  site-id Site1
  email-contact sakpuri@cisco.com
  phone-contact +1-800-000-0000
  streetaddress 12345 Cisco Way, San Jose, CA
  destination-profile Inventory
  destination-profile Inventory format full-txt
  destination-profile Inventory message-size 1000000
  destination-profile Service
  destination-profile Service format full-txt
  destination-profile Service message-size 1000000
  destination-profile dest1
  destination-profile dest1 format XML
  destination-profile dest1 message-size 500000
  destination-profile full_txt message-size 1000000
  destination-profile httpProf
  destination-profile httpProf format XML
  destination-profile httpProf message-size 0
  destination-profile short_txt message-size 4000
  destination-profile xml message-size 1000000
  destination-profile xml message-size 1000000
  destination-profile Inventory email-addr sakpuri@cisco.com
  destination-profile Service email-addr sakpuri@cisco.com
  destination-profile full_txt email-addr sakpuri@cisco.com
  destination-profile short_txt email-addr sakpuri@cisco.com
  destination-profile xml email-addr sakpuri@cisco.com
  destination-profile Service alert-group environmental
  destination-profile xml alert-group environmental
  destination-profile Inventory alert-group inventory
  destination-profile xml alert-group inventory
```

Displaying Delayed Trap Information

Use the **show running-config | in delay** command to display the system-delayed trap state. If no timer value is specified or if the timer value is set to 4 minutes, the following is displayed:

Displays the Delayed Trap Information with No Timer Value (Set to the Default 4 Minutes)

```
switch# show running-config | in delay
system delayed-traps enable mode FX
```

The following example shows the output when the timer value is set to any other value other than 4 minutes:

Displays the Delayed Trap Information with a Timer Value Other Than 4 Minutes

```
switch# show running-config | in delay
system delayed-traps enable mode FX
system delayed-traps timer 5
```

Verifying Alert Group Customization

To verify the alert group customization, use the **show callhome user-def-cmds** command.

```
switch# show callhome user-def-cmds
```

```
User configured commands for alert groups :
alert-group test user-def-cmd "show version"
```

Verifying Event Notification Trap

To verify the SNMP event notification trap, use the show snmp trap | inc callhome command.

```
switch# show snmp trap | inc callhome
callhome : event-notify Yes
callhome : smtp-send-fail No
```

Verifying Call Home Transport

The show callhome transport command displays all of the transport-related configurations for Call Home.

```
switch# show callhome transport
http vrf:management
from email addr:xyz-1@cisco.com
reply to email addr:xyz-1@cisco.com
smtp server:72.163.62.211
smtp server port:25
smtp server vrf:management
smtp server priority:0
http proxy server:10.64.65.52
http proxy server port:8080
http proxy status:Enabled
```

The following example shows how to configure SMTP server port:

```
switch# callhome
switch(config-callhome)# transport email mail-server 192.168.10.23 port 4
switch# config t
```

The following example shows how to configure SMTP server priority:

```
switch(config-callhome)# transport email mail-server 192.168.10.23 priority 60
switch# config t
```

Monitoring Call Home

This section includes the following topics:

Sample Syslog Alert Notification in Full-txt Format

```
source:MDS9000
Switch Priority:7
Device Id:DS-C9506@C@FG@07120011
Customer Id:basu
Contract Id:123
Site Id:San Jose
Server Id:DS-C9506@C@FG@07120011
Time of Event:2004-10-08T11:10:44
Message Name:SYSLOG_ALERT
Message Type:Syslog
Severity Level:2
System Name:10.76.100.177
Contact Name:Basavaraj B
Contact email:admin@yourcompany.com
Contact Phone:+91-80-310-1718
Street Address:#71 , Miller's Road
Event Description:2004 Oct  8 11:10:44 10.76.100.177 %PORT-5-IF_TRUNK_UP: %$VSAN 1%$ Interface
 fc2/5, vsan 1 is up
syslog_facility:PORT
start chassis information:
Affected Chassis:DS-C9506
Affected Chassis Serial Number:FG@07120011
Affected Chassis Hardware Version:0.104
Affected Chassis Software Version:3.1(1)
Affected Chassis Part No:73-8607-01
end chassis information:
```

Sample Syslog Alert Notification in XML Format

```
<?xml version="1.0" encoding="UTF-8" ?>
<soap-env:Envelope xmlns:soap-env="http://www.w3.org/2003/05/soap-envelope">
<soap-env:Header>
<aml-session:Session xmlns:aml-session="http://www.cisco.com/2004/01/aml-session"
soap-env:mustUnderstand="true"
soap-env:role="http://www.w3.org/2003/05/soap-envelope/role/next">
<aml-session:To>http://tools.cisco.com/neddce/services/DDCEService</aml-session:To>
<aml-session:Path>
<aml-session:Via>http://www.cisco.com/appliance/uri</aml-session:Via>
</aml-session:Path>
<aml-session:From>http://www.cisco.com/appliance/uri</aml-session:From>
<aml-session:MessageId>1004:FOX090306QT:3E55A81A</aml-session:MessageId>
</aml-session:Session>
</soap-env:Header>
<soap-env:Body>
<aml-block:Block xmlns:aml-block="http://www.cisco.com/2004/01/aml-block">
<aml-block:Header>
<aml-block:Type>http://www.cisco.com/2005/05/callhome/syslog</aml-block:Type>
<aml-block:CreationDate>2003-02-21 04:16:18 GMT+00:00</aml-block:CreationDate>
<aml-block:Builder>
<aml-block:Name>MDS</aml-block:Name>
<aml-block:Version>4.1</aml-block:Version>
</aml-block:Builder>
<aml-block:BlockGroup>
<aml-block:GroupId>1005:FOX090306QT:3E55A81A</aml-block:GroupId>
<aml-block:Number>0</aml-block:Number>
<aml-block:IsLast>true</aml-block:IsLast>
<aml-block:IsPrimary>true</aml-block:IsPrimary>
<aml-block:WaitForPrimary>false</aml-block:WaitForPrimary>
</aml-block:BlockGroup>
<aml-block:Severity>6</aml-block:Severity>
</aml-block:Header>
<aml-block:Content>
<ch:CallHome xmlns:ch="http://www.cisco.com/2005/05/callhome" version="1.0">
<ch:EventTime>2003-02-21 04:16:18 GMT+00:00</ch:EventTime>
<ch:MessageDescription>LICENSE_VIOLATION 2003 Feb 21 04:16:18 switch %$  %DAEMON-3-SYSTEM_MSG:
 &lt;&lt;%LICMGR-3-LOG_LICAPP_NO_LIC&gt;&gt; License file is missing for feature
SAN_EXTN_OVER_IP</ch:MessageDescription>
<ch:Event>
<ch:Type>syslog</ch:Type>
<ch:SubType>LICENSE_VIOLATION</ch:SubType>
<ch:Brand>Cisco</ch:Brand>
```

```
<ch:Series>MDS9000</ch:Series>
</ch:Event>
<ch:CustomerData>
<ch:UserData>
<ch:email>esajjana@cisco.com</ch:email>
</ch:UserData>
<ch:ContractData>
<ch:CustomerId>eeranna</ch:CustomerId>
<ch:SiteId>Bangalore</ch:SiteId>
<ch:ContractId>123</ch:ContractId>
<ch:DeviceId>DS-C9216I-K9@C@FOX090306QT</ch:DeviceId>
</ch:ContractData>
<ch:SystemInfo>
<ch:Name>switch</ch:Name>
<ch:Contact>Eeranna</ch:Contact>
<ch:Contactemail>esajjana@cisco.com</ch:Contactemail>
<ch:ContactPhoneNumber>+91-80-310-1718</ch:ContactPhoneNumber>
<ch:StreetAddress>#71, Miller&apos;s Road</ch:StreetAddress> </ch:SystemInfo>
</ch:CustomerData> <ch:Device> <rme:Chassis xmlns:rme="http://www.cisco.com/rme/4.0">
<rme:Model>DS-C9216I-K9</rme:Model>
<rme:HardwareVersion>1.0</rme:HardwareVersion>
<rme:SerialNumber>FOX090306QT</rme:SerialNumber>
</rme:Chassis>
</ch:Device>
</ch:CallHome>
</aml-block:Content>
<aml-block:Attachments>
<aml-block:Attachment type="inline">
<aml-block:Name>show logging logfile | tail -n 200</aml-block:Name> <aml-block:Data
encoding="plain">
<![CDATA[syslog_show:: command: 1055 param_count: 0
2003 Feb 21 04:11:48  %KERN-2-SYSTEM_MSG: Starting kernel... - kernel
2003 Feb 21 04:11:48  %KERN-3-SYSTEM_MSG: CMOS: Module initialized - kernel
2003 Feb 21 04:11:48  %KERN-2-SYSTEM_MSG: CARD TYPE: KING BB Index = 2344 - kernel
2003 Feb 21 04:12:04  %MODULE-5-ACTIVE_SUP_OK: Supervisor 1 is active (serial: JAB100700MC)
2003 Feb 21 04:12:04  %PLATFORM-5-MOD_STATUS: Module 1 current-status is MOD_STATUS_ONLINE/OK
2003 Feb 21 04:12:06  %IMAGE_DNLD-SLOT1-5-ADDON_IMG_DNLD_COMPLETE:  Addon module image
download process completed. Addon Image download completed, installing image please wait..
2003 Feb 21 04:12:07  %IMAGE_DNLD-SLOT1-5-ADDON_IMG_DNLD_SUCCESSFUL:  Addon module image
download and install process successful. Addon image installed.
2003 Feb 21 04:12:08  %KERN-3-SYSTEM_MSG: klm_af_xipc: Unknown parameter `start&apos; -
kernel
2003 Feb 21 04:12:08  %KERN-3-SYSTEM_MSG: klm_ips_portcfg: Unknown parameter `start&apos;
- kernel
2003 Feb 21 04:12:08  %KERN-3-SYSTEM_MSG: klm_flamingo: Unknown parameter `start&apos; -
kernel
2003 Feb 21 04:12:10  %PORT-5-IF_UP: Interface mgmt0 is up
2003 Feb 21 04:12:21 switch %LICMGR-3-LOG_LIC_FILE_MISSING: License file(s) missing for
feature ENTERPRISE_PKG.
2003 Feb 21 04:12:21 switch %LICMGR-3-LOG_LIC_FILE_MISSING: License file(s) missing for
feature SAN_EXTN_OVER_IP.
2003 Feb 21 04:12:21 switch %LICMGR-3-LOG_LIC_FILE_MISSING: License file(s) missing for
feature ENTERPRISE_PKG.
2003 Feb 21 04:12:21 switch %LICMGR-3-LOG_LIC_FILE_MISSING: License file(s) missing for
feature SAN_EXTN_OVER_IP.
2003 Feb 21 04:12:23 switch %PLATFORM-5-MOD_STATUS: Module 1 current-status is
MOD_STATUS_ONLINE/OK
2003 Feb 21 04:12:23 switch %MODULE-5-MOD_OK: Module 1 is online (serial: JAB100700MC)
2003 Feb 21 04:12:25 switch %PORT-5-IF_DOWN_ADMIN_DOWN: %$VSAN 1%$ Interface fc1/1 is down
 (Administratively down)
2003 Feb 21 04:12:25 switch %PORT-5-IF_DOWN_ADMIN_DOWN: %$VSAN 1%$ Interface fc1/2 is down
 (Administratively down)
2003 Feb 21 04:12:25 switch %PORT-5-IF_DOWN_ADMIN_DOWN: %$VSAN 1%$ Interface fc1/3 is down
 (Administratively down)
```

```
2003 Feb 21 04:12:25 switch %PORT-5-IF_DOWN_ADMIN_DOWN: %$VSAN 1%$ Interface fc1/4 is down
 (Administratively down)
2003 Feb 21 04:12:26 switch %PLATFORM-5-PS_STATUS: PowerSupply 1 current-status is PS_FAIL
2003 Feb 21 04:12:26 switch %PLATFORM-2-PS_FAIL: Power supply 1 failed or shut down (Serial
 number QCS1007109F)
2003 Feb 21 04:12:26 switch %PLATFORM-5-PS_FOUND: Power supply 2 found (Serial number
QCS1007109R)
2003 Feb 21 04:12:26 switch %PLATFORM-2-PS_OK: Power supply 2 ok (Serial number QCS1007109R)
2003 Feb 21 04:12:26 switch %PLATFORM-5-PS_STATUS: PowerSupply 2 current-status is PS_OK
2003 Feb 21 04:12:26 switch %PLATFORM-2-PS_FANOK: Fan in Power supply 2 ok
2003 Feb 21 04:12:26 switch %PLATFORM-5-FAN_DETECT: Fan module 1 (Serial number NWG0901031X)
 ChassisFan1 detected
2003 Feb 21 04:12:26 switch %PLATFORM-2-FAN_OK: Fan module ok
2003 Feb 21 04:12:26 switch %PLATFORM-2-CHASSIS_CLKMODOK: Chassis clock module A ok
2003 Feb 21 04:12:26 switch %PLATFORM-2-CHASSIS_CLKSRC: Current chassis clock source is
clock-A
2003 Feb 21 04:12:26 switch %PORT-5-IF_DOWN_ADMIN_DOWN: %$VSAN 1%$ Interface fc1/5 is down
 (Administratively down)
2003 Feb 21 04:12:26 switch %PORT-5-IF_DOWN_ADMIN_DOWN: %$VSAN 1%$ Interface fc1/6 is down
 (Administratively down)
2003 Feb 21 04:12:26 switch %PORT-5-IF_DOWN_ADMIN_DOWN: %$VSAN 1%$ Interface fc1/7 is down
 (Administratively down)
2003 Feb 21 04:12:26 switch %PORT-5-IF_DOWN_ADMIN_DOWN: %$VSAN 1%$ Interface fc1/8 is down
 (Administratively down)
2003 Feb 21 04:12:26 switch %PORT-5-IF_DOWN_ADMIN_DOWN: %$VSAN 1%$ Interface fc1/9 is down
 (Administratively down)
2003 Feb 21 04:12:26 switch %PORT-5-IF_DOWN_ADMIN_DOWN: %$VSAN 1%$ Interface fc1/10 is down
 (Administratively down)
2003 Feb 21 04:12:27 switch %PORT-5-IF_DOWN_ADMIN_DOWN: %$VSAN 1%$ Interface fc1/11 is down
 (Administratively down)
2003 Feb 21 04:12:27 switch %PORT-5-IF_DOWN_ADMIN_DOWN: %$VSAN 1%$ Interface fc1/12 is down
 (Administratively down)
2003 Feb 21 04:12:27 switch %PORT-5-IF_DOWN_ADMIN_DOWN: %$VSAN 1%$ Interface fc1/13 is down
 (Administratively down)
2003 Feb 21 04:12:27 switch %PORT-5-IF_DOWN_ADMIN_DOWN: %$VSAN 1%$ Interface fc1/14 is down
 (Administratively down)
2003 Feb 21 04:12:30 switch %PLATFORM-2-MOD_DETECT: Module 2 detected (Serial number
JAB0923016X) Module-Type IP Storage Services Module Model DS-X9304-SMIP
2003 Feb 21 04:12:30 switch %MODULE-2-MOD_UNKNOWN: Module type [25] in slot 2 is not supported
2003 Feb 21 04:12:45 switch %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty by root on
console0
2003 Feb 21 04:14:06 switch %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty by admin on
console0
2003 Feb 21 04:15:12 switch %VSHD-5-VSHD_SYSLOG_CONFIG_I: Configured from vty by admin on
console0
2003 Feb 21 04:15:52 switch %SYSMGR-3-BASIC_TRACE: core_copy: PID 1643 with message Core
not generated by system for licmgr(0). WCOREDUMP(9) returned zero .
2003 Feb 21 04:15:52 switch %SYSMGR-2-SERVICE_CRASHED: Service \"licmgr\" (PID 2272)
hasn&apos;t caught signal 9 (no core).
2003 Feb 21 04:16:18 switch %LICMGR-3-LOG_LIC_FILE_MISSING: License file(s) missing for
feature ENTERPRISE_PKG.
2003 Feb 21 04:16:18 switch %LICMGR-3-LOG_LIC_FILE_MISSING: License file(s) missing for
feature SAN_EXTN_OVER_IP.
2003 Feb 21 04:16:18 switch %LICMGR-3-LOG_LIC_FILE_MISSING: License file(s) missing for
feature ENTERPRISE_PKG.
2003 Feb 21 04:16:18 switch %LICMGR-3-LOG_LIC_FILE_MISSING: License file(s) missing for
feature SAN_EXTN_OVER_IP.
2003 Feb 21 04:16:18 switch %CALLHOME-2-EVENT: LICENSE_VIOLATION
2003 Feb 21 04:16:18 switch %CALLHOME-2-EVENT: LICENSE_VIOLATION
2003 Feb 21 04:16:18 switch %CALLHOME-2-EVENT: LICENSE_VIOLATION
2003 Feb 21 04:16:18 switch %CALLHOME-2-EVENT: LICENSE_VIOLATION ]]> </aml-block:Data>
</aml-block:Attachment> <aml-block:Attachment type="inline"> <aml-block:Name>show license
usage</aml-block:Name> <aml-block:Data encoding="plain">
<![CDATA[Feature                      Ins  Lic   Status Expiry Date Comments
```

```
                                 Count
--------------------------------------------------------------------------------
DMM_184_PKG                   No    0   Unused             Grace expired
FM_SERVER_PKG                 No    -   Unused             Grace expired
MAINFRAME_PKG                 No    -   Unused             Grace expired
ENTERPRISE_PKG                Yes   -   Unused never       license missing
DMM_FOR_SSM_PKG               No    0   Unused             Grace expired
SAN_EXTN_OVER_IP              Yes   8   Unused never       8 license(s) missing
PORT_ACTIVATION_PKG           No    0   Unused             -
SME_FOR_IPS_184_PKG           No    0   Unused             Grace expired
STORAGE_SERVICES_184          No    0   Unused             Grace expired
SAN_EXTN_OVER_IP_18_4         No    0   Unused             Grace expired
SAN_EXTN_OVER_IP_IPS2         No    0   Unused             Grace expired
SAN_EXTN_OVER_IP_IPS4         No    0   Unused             Grace expired
STORAGE_SERVICES_SSN16        No    0   Unused             Grace expired
10G_PORT_ACTIVATION_PKG       No    0   Unused             -
STORAGE_SERVICES_ENABLER_PKG  No    0   Unused             Grace expired
--------------------------------------------------------------------------------
**** WARNING: License file(s) missing. **** ]]> </aml-block:Data> </aml-block:Attachment>
</aml-block:Attachments> </aml-block:Block> </soap-env:Body> </soap-env:Envelope>
```

Sample RMON Notification in XML Format

```
<?xml version="1.0" encoding="UTF-8" ?>
<soap-env:Envelope xmlns:soap-env="http://www.w3.org/2003/05/soap-envelope">
<soap-env:Header>
<aml-session:Session xmlns:aml-session="http://www.cisco.com/2004/01/aml-session"
soap-env:mustUnderstand="true"
soap-env:role="http://www.w3.org/2003/05/soap-envelope/role/next">
<aml-session:To>http://tools.cisco.com/neddce/services/DDCEService</aml-session:To>
<aml-session:Path>
<aml-session:Via>http://www.cisco.com/appliance/uri</aml-session:Via>
</aml-session:Path>
<aml-session:From>http://www.cisco.com/appliance/uri</aml-session:From>
<aml-session:MessageId>1086:FHH0927006V:48BA26BD</aml-session:MessageId>
</aml-session:Session>
</soap-env:Header>
<soap-env:Body>
<aml-block:Block xmlns:aml-block="http://www.cisco.com/2004/01/aml-block">
<aml-block:Header>
<aml-block:Type>http://www.cisco.com/2005/05/callhome/diagnostic</aml-block:Type>
<aml-block:CreationDate>2008-08-31 05:06:05 GMT+00:00</aml-block:CreationDate>
<aml-block:Builder>
<aml-block:Name>MDS</aml-block:Name>
<aml-block:Version>4.1</aml-block:Version>
</aml-block:Builder>
<aml-block:BlockGroup>
<aml-block:GroupId>1087:FHH0927006V:48BA26BD</aml-block:GroupId>
<aml-block:Number>0</aml-block:Number>
<aml-block:IsLast>true</aml-block:IsLast>
<aml-block:IsPrimary>true</aml-block:IsPrimary>
<aml-block:WaitForPrimary>false</aml-block:WaitForPrimary>
</aml-block:BlockGroup>
<aml-block:Severity>2</aml-block:Severity>
</aml-block:Header>
<aml-block:Content>
<ch:CallHome xmlns:ch="http://www.cisco.com/2005/05/callhome" version="1.0">
<ch:EventTime>2008-08-31 05:06:05 GMT+00:00</ch:EventTime>
<ch:MessageDescription>RMON_ALERT WARNING(4) Falling:iso.3.6.1.4.1.9.9.305.1.1.1.0=1 &lt;=
 89:1, 4</ch:MessageDescription>
<ch:Event>
<ch:Type>environment</ch:Type>
```

```
<ch:SubType>minor</ch:SubType>
<ch:Brand>Cisco</ch:Brand>
<ch:Series>MDS9000</ch:Series>
</ch:Event>
<ch:CustomerData>
<ch:UserData>
<ch:email>mchinn@cisco.com</ch:email>
</ch:UserData>
<ch:ContractData>
<ch:CustomerId>12ss</ch:CustomerId>
<ch:SiteId>2233</ch:SiteId>
<ch:ContractId>rrr55</ch:ContractId>
<ch:DeviceId>DS-C9513@C@FHH0927006V</ch:DeviceId>
</ch:ContractData>
<ch:SystemInfo>
<ch:Name>sw172-22-46-174</ch:Name>
<ch:Contact>Mani</ch:Contact>
<ch:Contactemail>mchinn@cisco.com</ch:Contactemail>
<ch:ContactPhoneNumber>+1-800-304-1234</ch:ContactPhoneNumber>
<ch:StreetAddress>1234 wwee</ch:StreetAddress>
</ch:SystemInfo>
</ch:CustomerData>
<ch:Device>
<rme:Chassis xmlns:rme="http://www.cisco.com/rme/4.0">
<rme:Model>DS-C9513</rme:Model>
<rme:HardwareVersion>0.205</rme:HardwareVersion>
<rme:SerialNumber>FHH0927006V</rme:SerialNumber>
</rme:Chassis>
</ch:Device>
</ch:CallHome>
</aml-block:Content>
</aml-block:Block>
</soap-env:Body>
</soap-env:Envelope>
```

Field Descriptions for Call Home

This section displays the field descriptions for this feature:

Call Home General

Field	Description
Contact	The contact person for this switch, together with information on how to contact this person.
PhoneNumber	The phone number of the contact person. The phone number must start with '+' and contains only numeric characters except for space and '-'. Some valid phone numbers are +44 20 8332 9091 +45 44886556 +81-46-215-4678 +1-650-327-2600.
EmailAddress	The email address of the contact person. Some valid email addresses are raj@helpme.com, bob@service.com, mtom@abc.caview.ca.us.
StreetAddress	The mailing address of this switch.
CustomerId	A string, in whatever format is appropriate, to identify the customer.

Field	Description
ContractId	A string, in whatever format is appropriate, to identify the support contract between the customer and support partner.
SiteId	A location identifier of this device.
DeviceServicePriority	The service priority of the device. This determines how fast the device has to be serviced.
Enable	Enables/disables the Call Home infrastructure on the local device.

Related Topics

Information About Call Home, on page 49

Call Home Destinations

Field	Description
EmailAddress	The email address associated this destination profile. Some examples are raj@helpme.com, bob@service.com, mtom@abc.caview.ca.us.

Related Topics

Call Home Destination Profiles, on page 52

Call Home SMTP Servers

Field	Description
Address Type, Address	IP address of the SMTP server.
Port	TCP port of the SMTP server.
Priority	Priority value.

Call Home Email Setup

Field	Description
From	The email address that is to be used in the From field when sending the email using SMTP. Some examples are raj@helpme.com, bob@service.com, mtom@abc.caview.ca.us.
ReplyTo	The email address that is to be used in the Reply-To field when sending the email using SMTP. Some examples are raj@helpme.com, bob@service.com, mtom@abc.caview.ca.us.
IP Address Type	The IP address type (IPv4, IPv6, or DNS).

Field	Description
Name or IP Address	Name or IP address of the SMTP server.
Port	TCP port of the SMTP server.

Related Topics

General EMail Options Using HTTPS Support, on page 54

Call Home Alerts

Field	Description
Action	Test — Sends a Call Home message TestWithInventory — Sends a message with inventory details.
Status	The status of the last Call Home action invocation.
FailureCause	The failure cause for the last Call Home test invocation.
LastTimeSent	When the last Call Home alert was sent.
NumberSent	The number of Call Home alerts sent.
Interval	Time frame for sending the periodic software inventory Call Home message.
Throttling Enable	If checked, enables the message throttling mechanism implemented on the system, to limit the number of Call Home messages for an alert type within a time frame. The maximum is 30 in a 2-hour time frame, and any further messages for that alert type are discarded.
Enable	If checked, enables the sending of periodic software inventory Call Home messages on the system.

Related Topics

Call Home Alert Groups, on page 52

Call Home Message Level Feature, on page 53

Call Home User Defined Command

Field	Description
User Defined Command	Configures user-defined commands for the Call Home alert group types.

Delayed Traps

Field	Description
Enable	Enables or disables delay traps.

Field	Description
Delay	Delays interval in minutes (valid values are between 1 to 60).

Call Home Profiles

Field	Description
MsgFormat	XML, full text, or short text.
MaxMsgSize	Maximum message size that can be sent to destination pointed to by this destination profile.
MsgLevel	Threshold level, used for filtering alert messages sent to a destination. Callhome alert message with severity level lower than the configured threshold level would not be sent. The default threshold level is debug (1), which means all the alert messages will be sent.
AlertGroups	The list of configured alert groups for this destination profile.

Event Destinations Addresses

Field	Description
Address/Port	IP address and port to send event.
Security Name	The SNMP parameters to be used when generating messages to be sent to this address.
Security Model	Is used when generating SNMP messages using this entry.
Inform Type	• Trap — Unacknowledged event • Inform — Acknowledged event.
Inform Timeout	This expected maximum round-trip time for communicating with the address.
RetryCount	The number of retries to be attempted when a response is not received for a generated message.

Event Destinations Security (Advanced)

Field	Description
MPModel	The message processing model to be used when generating SNMP messages using this entry.
SecurityModel	The security model to be used when generating SNMP messages using this entry.
SecurityName	Identifies the principal on whose behalf SNMP messages will be generated using this entry.
SecurityLevel	The level of security to be used when generating SNMP messages using this entry.

Event Filters General

Field	Description
FSPF - Nbr State Changes	Specifies whether or not the local switch should issue notification when the local switch learns of a change in the neighbor's state (state in the FSPF neighbor finite state machine) on an interface on a VSAN.
Domain Mgr - ReConfig Fabrics	Specifies whether or not the local switch should issue a notification on sending or receiving ReConfigureFabric (RCF) on a VSAN.
Zone Server - Request Rejects	Specifies if the zone server should issue a notification on rejects.
Zone Server - Merge Failures	Specifies if the zone server should issue a notification on merge failures.
Zone Server - Merge Successes	Specifies if the zone server should issue a notification on merge successes.
Zone Server - Default Zone Behavior Change	Specifies if the zone server should issue a notification if the propagation policy changes.
Zone Server - Unsupp Mode	Specifies if the zone server should issue a notification on unsupp mode changes
FabricConfigServer - Request Rejects	Specifies if the fabric configuration server should issue a notification on rejects.
RSCN - ILS Request Rejects	Specifies if the RSCN module should generate notifications when a SW_RSCN request is rejected.
RSCN - ILS RxRequest Rejects	Specifies if the RSCN module should generate notifications when a SW_RSCN request is rejected.
RSCN - ELS Request Rejects	Specifies if the RSCN module should generate notifications when a SCR or RSCN request is rejected.
FRU Changes	A false value will prevent field replaceable unit (FRU) notifications from being generated by this system.
SNMP - Community Auth Failure	Indicates whether the SNMP entity is permitted to generate authenticationFailure traps.
VRRP	Indicates whether the VRRP-enabled router will generate SNMP traps for events defined in this MIB.
FDMI	Specifies if the FDMI should generate notifications when a registration request is rejected.
License Manager	Indicates whether the system should generate notifications.
Port/Fabric Security	Specifies if the system should generate notifications when a port/fabric security issue arises.
FCC	Specifies whether the agent should generate notifications.

Field	Description
Name Server	If checked, the name server generates a notification when a request is rejected. If false, the notification is not generated.

Event Filters Interfaces

Field	Description
EnableLinkTrap	Indicates whether linkUp/linkDown traps should be generated for this interface.

Event Filters Control

Field	Description
Variable	Represents the notification to be controlled.
Descr	Description about the notification.
Enabled	Check to enable notification of the control. Shows the status of the control.

Note You see the Descr column only on switches that run Cisco NX-OS Release 5.0 or later.

Additional References

For additional information related to implementing Call Home, see the following section:

MIBs

MIBs	MIBs Link
• CISCO-CALLHOME-CAPABILITY-MIB • CISCO-CALLHOME-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_

Feature History for Call Home

Feature History for Call Home, on page 111 lists the release history for this feature. Only features that were introduced or modified in Release 3.x or a later release appear in the table.

Table 17: Feature History for Call Home

Feature Name	Releases	Feature Information
Call Home HTTP Proxy Server	5.2	Added the Call Home HTTP Proxy Server support details.
Call Home Wizard	5.2	Added the Call Home Wizard configuration details.
Call Home HTTP Proxy Server	5.2	Added the Call Home HTTP Proxy Server support details. Added Verifying Callhome Transport commands.
Multiple SMTP Server Support	5.0(1a)	Added Multiple SMTP Server Support details. Added Verifying Callhome Transport commands.
Notification Enhancements	5.0(1a)	Added the enhancement in Notification in the Event Filter Using Device Manager.
Call Home	4.1(1b)	Added the HTTPS support for Call Home.
Call Home - Delayed Traps for EMC Call Home configuration window in DCNM-SAN.	4.1(1a)	Added the delayed traps enhancements for EMC Call Home.
Call Home Destination tab	4.2(1)	Added the enhancement in Destination tab.
Call Home HTTPs support	4.2(1)	Added Call Home HTTPs enhancement.
EMC Email Home	3.3(3)	EMC Email Home configuration information was added to this chapter.
EMC Call Home	3.0(1)	Enables the forwarding of traps as XML data using email, according to EMC specifications.
Call Home enhancement	3.0(1)	Enables customization of alert group messages.

CHAPTER 6

Scheduling Maintenance Jobs

The Cisco MDS command scheduler feature helps you schedule configuration and maintenance jobs in any switch in the Cisco MDS 9000 Family. You can use this feature to schedule jobs on a one-time basis or periodically.

Information About the Command Scheduler

The Cisco NX-OS command scheduler provides a facility to schedule a job (set of CLI commands) or multiple jobs at a specified time in the future. The job(s) can be executed once at a specified time in the future or at periodic intervals.

You can use this feature to schedule zone set changes, make QoS policy changes, back up data, save the configuration and do other similar jobs.

Scheduler Terminology

The following terms are used in this chapter:

- Job—A job is a set of NX-OS CLI commands (EXEC and config mode) that are executed as defined in the schedule.

- Schedule—A schedule determines the time when the assigned jobs must be executed. Multiple jobs can be assigned to a schedule. A schedule executes in one of two modes: one-time or periodic.
- Periodic mode—A job is executed at the user-specified periodic intervals, until it is deleted by the administrator. The following types of periodic intervals are supported:
    - Daily—The job is executed once a day.
    - Weekly—The job is executed once a week.
    - Monthly—The job is executed once a month.
    - Delta—The job is executed beginning at the specified start time and thereafter at user-specified intervals (days:hours:minutes).
- One-time mode—The job is executed once at a user-specified time.

Licensing Requirements for Command Scheduler

To use the command scheduler, you do not need to obtain any license.

Guidelines and Limitations

Before scheduling jobs on a Cisco MDS switch, note the following guidelines:

- Prior to Cisco MDS SAN-OS Release 3.0(3), only users local to the switch could perform scheduler configuration. As of Cisco MDS SAN-OS Release 3.0(3), remote users can perform job scheduling using AAA authentication.
- Be aware that the scheduled job can fail if it encounters one of the following situations when executing the job:
    - If the license has expired for a feature at the time when a job containing commands pertaining to that feature is scheduled.
    - If a feature is disabled at the time when a job containing commands pertaining to that feature is scheduled.
    - If you have removed a module from a slot and the job has commands pertaining to the interfaces for that module or slot.
- Verify that you have configured the time. The scheduler does not have any default time configured. If you create a schedule and assign job(s) and do not configure the time, that schedule is not launched.
- While defining a job, verify that no interactive or disruptive commands (for example, **copy bootflash:** *file* **ftp:** *URI*, **write erase**, and other similar commands) are specified as part of a job because the job is executed noninteractively at the scheduled time.

Default Settings

Table 18: Default Command Scheduler Parameters , on page 115 lists the default settings for command scheduling parameters.

Table 18: Default Command Scheduler Parameters

Parameters	Default
Command scheduler	Disabled.
Log file size	16 KB.

Configuring the Command Scheduler

The Cisco NX-OS command scheduler provides a facility to schedule a job (set of CLI commands) or multiple jobs at a specified time in the future.

Task Flow for Configuring the Command Scheduler

Follow these steps to configure the Command Scheduler:

Procedure

Step 1 Enable the scheduler.

Step 2 Authorize remote user access (optional).

Step 3 Define the job.

Step 4 Specify the schedule and assign jobs to the schedule.

Step 5 Specify the time for the schedule(s).

Step 6 Verify the scheduled configuration.

Enabling the Command Scheduler

To use the scheduling feature, you must explicitly enable this feature on the required switches in the fabric. By default, this feature is disabled in all switches in the Cisco MDS 9000 Family.

The configuration and verification commands for the command scheduler feature are only available when this feature is enabled on a switch. When you disable this feature, all related configurations are automatically discarded.

To enable the command scheduling feature, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **feature scheduler**

Enables the command scheduler.

Step 3 switch(config)# **no feature scheduler**

Discards the scheduler configuration and disables the command scheduler (default).

Examples

To display the command schedule status, use the **show scheduler config** command.

```
switch# show scheduler config
  config terminal
  feature scheduler
  scheduler logfile size 16
  end
```

Configuring Remote User Authentication

Prior to Cisco MDS SAN-OS Release 3.0(3), only users local to the switch could perform scheduler configuration. As of Cisco MDS SAN-OS Release 3.0(3), remote users can perform job scheduling using AAA authentication.

To configure remote user authentication, follow these steps:

Before you begin

AAA authentication requires the clear text password of the remote user before creating and configuring command scheduler jobs.

Procedure

Step 1 switch# **configuration terminal**

Enters configuration mode.

Step 2 switch(config)# **scheduler aaa-authentication password X12y34Z56a**

Configures a clear text password for remote users.

Step 3 switch(config)# **scheduler aaa-authentication password 0 X12y34Z56a**

Configures a clear text password for remote users.

Step 4 switch(config)# **no scheduler aaa-authentication password**

Removes the clear text password for remote users.

Step 5 switch(config)#**scheduler aaa-authentication user newuser password Z98y76X54b**

Configures a clear text password for remote user newuser.

Step 6 switch(config)#**scheduler aaa-authentication user newuser password 0 Z98y76X54b**

Configures a clear text password for remote user newuser.

Step 7 switch(config)# **no scheduler aaa-authentication password user newuser**

Removes the clear text password for remote user newuser.

Defining a Job

To define a job, you must specify the job name. This action places you in the job definition (config-job) submode. In this submode, you can define the sequence of CLI commands that the job has to perform. Be sure to exit the config-job submode to complete the job definition.

- Job configuration files created using MDS NX-OS or SAN-OS releases before Cisco MDS NX-OS Release 4.1(1b) are not supported. However, you can edit the job configuration file and combine the commands within a job into a single line using a semicolon (;).
- You must exit the config-job submode for the job definition to be complete.
- You cannot modify or remove a command after exiting the config-job submode. To make changes, you must explicitly delete the defined job name and then reconfigure the job with new commands.

To define a job for the command scheduler, follow these steps:

Procedure

Step 1 switch# **configuration terminal**

Enters the configuration mode.

Step 2 switch(config)# **scheduler job name addMemVsan99**

switch(config-job)#

Defines a job name and enters the job definition submode.

Step 3 switch(config-job)# *command1* ;[*command2* ;*command3* ;...]

switch(config-job-submode)# **end**

Example:

```
switch(config-job)# configure terminal;vsan database;vsan 99 interface fc1/1 4
switch(config-job-config-vsan-db)# end
switch#
```

Specifies a sequence of actions for the specified job. The defined commands are checked for validity and stored for future use.

Note Be sure you exit the config-job submode.

Example:

```
switch(config)# scheduler job name offpeakQOS
switch(config-job)# configuration terminal; qos class-map offpeakbackupcmap match-all ;
match source-wwn 23:15:00:05:30:00:2a:1f ; match destination-wwn 20:01:00:05:30:00:28:df
;exit ; qos policy-map offpeakbackuppolicy ; class offpeakbackupcmap ; priority high ; exit
 ; exit ; qos service policy offpeakbackuppolicy vsan 1
switch(config-job)# end
switch#
```

Provides example of scheduling a set of configuration commands.

Step 4 **exit**

Example:

```
switch(config-job)# exit
switch(config)#
```

Exits the job configuration mode and saves the job.

Step 5 **show scheduler job** [*name*]

Example:

```
switch(config)# show scheduler job
```

(Optional) Displays the job information.

Step 6 **copy running-config startup-config**

Example:

```
switch(config)# copy running-config startup-config
```

(Optional) Saves this configuration change.

Deleting a Job

To delete a job for the command scheduler, follow these steps:

Procedure

Step 1 switch# **configuration terminal**

Enters the configuration mode.

Step 2 switch(config)# **no scheduler job name addMemVsan99**

Deletes a defined job and all commands defined within that job.

Specifying a Schedule

After defining jobs, you can create schedules and assign jobs to the schedule. Subsequently, you can configure the time of execution. The execution can be one-time or periodic depending on your requirements. If the time for the schedule is not configured, then it will never be executed.

You can specify a periodic job execution at the specified (daily, weekly, monthly, or delta) intervals.

To specify a periodic job for the command scheduler, follow these steps:

Procedure

Step 1 switch# **configuration terminal**

Enters the configuration mode.

Step 2 switch(config)# **scheduler schedule name weekendbackupqos**

switch(config-schedule)#

Defines a job schedule (weekendbackup) and enters the submode for that schedule.

Step 3 switch(config)# **no scheduler schedule name weekendbackup**

Deletes the defined schedule.

Step 4 switch(config-schedule)# **job name offpeakZoning**

switch(config-schedule)# **job name offpeakQOS**

Assigns two jobs (offpeakZoning and offpeakQOS) for this schedule.

Step 5 switch(config-schedule)# **no job name addMem99**

Deletes the job assigned for this schedule.

Examples

The following examples are for reference:

Command	Purpose
switch(config-schedule)# **time daily 23:00**	Executes the specified jobs at 11 p.m. every day.
switch(config-schedule)# **time weekly Sun:23:00**	Specifies a weekly execution every Sunday at 11 p.m.
switch(config-schedule)# **time monthly 28:23:00**	Specifies a monthly execution at 11 p.m on the 28th of each month. If you specify the date as either 29, 30, or 31, the command is automatically executed on the last day of each month.
switch(config-schedule)# **time start now repeat 48:00**	Specifies a job to be executed every 48 hours beginning 2 minutes from *now* —if today is September 24, 2004, and the time is now 2:00 p.m., the command begins executing at 2 minutes past 2:00 p.m. on September 24, 2004, and continues to execute every 48 hours after that.
switch(config-schedule)# **time start 14:00 repeat 14:00:00**	If today is September 24, 2004, (Friday), this command specifies the job to be executed every alternate Friday at 2 p.m. (every 14 days).

The most significant fields in the **time** parameter are optional. If you omit the most significant fields, the values are assumed to be the same as the current time. For example, if the current time is September 24, 2004, 22:00 hours, then the commands are executed as follows:

- The **time start 23:00 repeat 4:00:00** command implies a start time of September 24, 2004, 23:00 hours.
- The **time daily 55** command implies every day at 22:55 hours.
- The **time weekly 23:00** command implies every Friday at 23:00 hours.
- The **time monthly 23:00** command implies the 24th of every month at 23:00 hours.

Note If the time interval configured for any schedule is smaller than the time taken to execute its assigned job(s), then the subsequent schedule execution occurs only after the configured interval amount of time has elapsed following the completion time of the last iteration of the schedule. For example, a schedule is executed at 1-minute intervals and a job assigned to it takes 2 minutes to complete. If the first schedule is at 22:00 hours, the job finishes at 22:02 after which the 1-minute interval is observed, and the next execution occurs at 22:03 and finishes at 22:05.

Specifying a One-Time Schedule

When you specify a one-time job execution, that job is only executed once.

To specify a one-time job for the command scheduler, follow these steps:

Procedure

Step 1 switch# **configuration terminal**

Enters the configuration mode.

Step 2 switch(config)# **scheduler schedule name configureVsan99**

switch(config-schedule)#

Defines a job schedule (configureVsan99) and enters the submode for that schedule.

Step 3 switch(config-schedule)# **job name addMemVsan99**

Assigns a predefined job name (addMemVsan99) for this schedule.

Step 4 switch(config-schedule)# **time start 2004:12:14:23:00**

Specifies a one-time execution on December 14, 2004, at 11 p.m.

Step 5 switch(config-schedule)# **no time**

Deletes the time assigned for this schedule.

Deleting a Schedule

To delete a schedule, follow these steps:

Procedure

Step 1 switch# **configuration terminal**

Enters the configuration mode.

Step 2 switch(config)# **no scheduler schedule name weekendbackup**

Deletes the defined schedule.

Removing an Assigned Job

To remove an assigned job, follow these steps:

Procedure

Step 1 switch# **configuration terminal**

Enters the configuration mode.

Step 2 switch(config)# **scheduler schedule name weekendbackupqos**

switch(config-schedule)#

Specifies a job schedule (weekendbackupqos) and enters the submode for that schedule.

Step 3 switch(config-schedule)# **no job name addMem99**

Removes a job (addMem99) assigned to this schedule.

Deleting a Schedule Time

To delete the schedule time, follow these steps:

Procedure

Step 1 switch# **configuration terminal**

Enters the configuration mode.

Step 2 switch(config)# **scheduler schedule name weekendbackupqos**

switch(config-schedule)#

Defines a job schedule (weekendbackup) and enters the submode for that schedule.

Step 3 switch(config-schedule)# **no time**

Deletes the schedule time configuration.The schedule will not be run until the time is configured again.

Configuring Execution Logs

The command scheduler maintains a log file. While you cannot modify the contents of this file, you can change the file size. This log file is a circular log that contains the output of the job executed. If the output of the job is greater than the log file, then the output stored in this file remains truncated.

You can configure the log file size to be a maximum of 1024 KB. The default size of the execution log file is 16 KB.

To configure the execution log file size, follow these steps:

Procedure

Step 1 switch# **configuration terminal**

Enters the configuration mode.

Step 2 switch(config)# **scheduler logfile size 1024**

Configures the log file to be a maximum of 1024 KB

Step 3 switch(config)# **no scheduler logfile size**

Defaults to the log size of 16 KB.

Clearing the Execution Log File Contents

To clear the contents of the scheduler execution log file, issue the clear scheduler logfile command in EXEC mode.

```
switch# clear scheduler logfile
```

Verifying Scheduler Configuration

To display the command scheduler configuration information, perform one of the following tasks:

Command	Purpose
show scheduler config	Displays the scheduler configuration
show scheduler schedule	Verifies the command scheduler execution status
show scheduler job	Verifies the job definition
show scheduler logfile	Displays the execution log for all jobs executed in the system
clear scheduler logfile	Clear the contents of the scheduler execution log file

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference* .

Verifying the Command Scheduler Configuration

To display the scheduler configuration, use the **show scheduler config** command.

```
switch# show scheduler config
config terminal
  feature scheduler
  scheduler logfile size 512
end
config terminal
   scheduler job name addMemVsan99
     config terminal
       vsan database
       vsan 99 interface fc1/1
       vsan 99 interface fc1/2
       vsan 99 interface fc1/3
       vsan 99 interface fc1/4
end
config terminal
  scheduler schedule name configureVsan99
    time start 2004:8:10:9:52
    job name addMemVsan99
end
```

Verifying the Command Scheduler Execution Status

To verify the command scheduler execution status, use the **show scheduler schedule** command.

```
switch# show scheduler schedule configureVsan99
Schedule Name      : configureVsan99
------------------------------------
User Name          : admin
Schedule Type      : Run once on Tue Aug 10 09:48:00 2004
Last Execution Time: Tue Aug 10 09:48:00 2004
-----------------------------------------------
          Job Name                  Status
```

Verifying the Job Definition

To verify the job definition, use the **show scheduler job** command.

```
switch# show scheduler job addMemVsan99
Job Name: addMemVsan99
----------------------
  config terminal
   vsan database
    vsan 99 interface fc1/1
    vsan 99 interface fc1/2
    vsan 99 interface fc1/3
    vsan 99 interface fc1/4
```

Displaying Execution Log File Contents

To display the execution log for all jobs executed in the system, use the **show scheduler logfile** command.

```
switch# show scheduler logfile
Job Name       : addMemVsan99              Job Status: Success (0)
Schedule Name  : configureVsan99           User Name : admin
Completion time: Tue Aug 10 09:48:00 2004
--------------------------- Job Output ---------------------------
`config terminal`
`vsan database`
`vsan 99 interface fc1/1`
`vsan 99 interface fc1/2`
`vsan 99 interface fc1/3`
`vsan 99 interface fc1/4`
```

To display the scheduler password configuration for remote users, use the **show running-config** command.

```
switch# show running-config | include “scheduler aaa-authentication”
scheduler aaa-authentication username newuser password 7 "C98d76S54e"
```

Note The scheduler remote user passwords are always displayed in encrypted form in the **show running-config** command output. The encrypted option (**7**) in the command exists to support applying the ASCII configuration to the switch.

To display the execution log file configuration, use the **show scheduler config** command.

```
switch# show scheduler config
config terminal
  feature scheduler
  scheduler logfile size 1024
end
```

Clearing the Execution Log File Contents

To clear the contents of the scheduler execution log file, issue the **clear scheduler logfile** command in EXEC mode.

```
switch# clear scheduler logfile
-----------------------------------------------
addMemVsan99                       Success (0)
```

Configuration Examples for Scheduler

```
configure terminal

scheduler job name start
 configure
 no cli var name time
 exit
 echo $(TIMESTAMP) | sed 's/^/cli var name time /' | vsh
 show switchname > debug-$(time)-1
 show switchname > debug-$(time)-2
 exit

scheduler job name part1
 show clock >> debug-$(time)-1
 show interface mgmt 0 >> debug-$(time)-1
 sleep 60
 show clock >> debug-$(time)-1
 show interface mgmt 0 >> debug-$(time)-1
 sleep 200
 gzip debug-$(time)-1
 exit

scheduler job name part2
 show clock >> debug-$(time)-2
 show processes cpu history >> debug-$(time)-2
 sleep 60
 show clock >> debug-$(time)-2
 show processes cpu history >> debug-$(time)-2
 show clock >> debug-$(time)-2
 gzip debug-$(time)-2
 exit

scheduler schedule name cpu-stats
  job name start
  job name part1
  job name part2
  time start 2001:12:31:01:00
  exit

end
```

CHAPTER 7

System Status Monitoring

This chapter provides details on monitoring the health of the switch.

Feature History for System Status Monitoring

Table 19: Feature History for System Status Monitoring , on page 127 lists the release history for this feature. Only features that were introduced or modified in Release 3.x or a later release appear in the table.

Table 19: Feature History for System Status Monitoring

Feature Name	Releases	Feature Information
Kernel Core Logging	8.4(2c)	Core files are created when NX-OS encounters an unrecoverable fault. The core files can be used by Cisco to diagnose the fault.
Common Information Model	3.3(1a)	Added commands for displaying Common Information Model.
On-line system health maintenance (OHMS) enhancements	3.0(1)	Includes the following OHMS enhancements: • Configuring the global frame length for loopback test for all modules on the switch. • Specifying frame count and frame length on for the loopback test on a specific module. • Configuring source and destination ports for external loopback tests. • Providing serdes loopback test to check hardware.

Feature Name	Releases	Feature Information
On-board failure logging (OBFL)	3.0(1)	Describes OBFL, how to configure it for Generation 2 modules, and how to display the log information.

Information About System Status Monitoring

Online Health Management System

The Online Health Management System (OHMS) (system health) is a hardware fault detection and recovery feature. It ensures the general health of switching, services, and supervisor modules in any switch in the Cisco MDS 9000 series.

The OHMS monitors system hardware in the following ways:

- The OHMS component running on the active supervisor maintains control over all other OHMS components running on the other modules in the switch.
- The system health application running in the standby supervisor module only monitors the standby supervisor module, if that module is available in the HA standby mode.

The OHMS application launches a daemon process in all modules and runs multiple tests on each module to test individual module components. The tests run at preconfigured intervals, cover all major fault points, and isolate any failing component in the MDS switch. The OHMS running on the active supervisor maintains control over all other OHMS components running on all other modules in the switch.

On detecting a fault, the system health application attempts the following recovery actions:

- Performs additional testing to isolate the faulty component.
- Attempts to reconfigure the component by retrieving its configuration information from persistent storage.
- If unable to recover, sends Call Home notifications, system messages and exception logs; and shuts down and discontinues testing the failed module or component (such as an interface).
- Sends Call Home and system messages and exception logs as soon as it detects a failure.
- Shuts down the failing module or component (such as an interface).
- Isolates failed ports from further testing.
- Reports the failure to the appropriate software component.
- Switches to the standby supervisor module, if an error is detected on the active supervisor module and a standby supervisor module exists in the Cisco MDS switch. After the switchover, the new active supervisor module restarts the active supervisor tests.
- Reloads the switch if a standby supervisor module does not exist in the switch.
- Provides CLI support to view, test, and obtain test run statistics or change the system health test configuration on the switch.
- Performs tests to focus on the problem area.

Each module is configured to run the test relevant to that module. You can change the default parameters of the test in each module as required.

Loopback Test Configuration Frequency

Loopback tests are designed to identify hardware errors in the data path in the module(s) and the control path in the supervisors. One loopback frame is sent to each module at a preconfigured frequency—it passes through each configured interface and returns to the supervisor module.

The loopback tests can be run at frequencies ranging from 5 seconds (default) to 255 seconds. If you do not configure the loopback frequency value, the default frequency of 5 seconds is used for all modules in the switch. Loopback test frequencies can be altered for each module.

Loopback Test Configuration Frame Length

Loopback tests are designed to identify hardware errors in the data path in the module(s) and the control path in the supervisors. One loopback frame is sent to each module at a preconfigured size—it passes through each configured interface and returns to the supervisor module.

The loopback tests can be run with frame sizes ranging from 0 bytes to 128 bytes. If you do not configure the loopback frame length value, the switch generates random frame lengths for all modules in the switch (auto mode). Loopback test frame lengths can be altered for each module.

Hardware Failure Action

The failure-action command controls the Cisco NX-OS software from taking any action if a hardware failure is determined while running the tests.

By default, this feature is enabled in all switches in the Cisco MDS 9000 Family—action is taken if a failure is determined and the failed component is isolated from further testing.

Failure action is controlled at individual test levels (per module), at the module level (for all tests), or for the entire switch.

Performing Test Run Requirements

Enabling a test does not guarantee that the test will run.

Tests on a specific interface or module only run if you enable system health for all of the following items:

- The entire switch
- The required module
- The required interface

Tip The test will not run if system health is disabled in any combination. If system health is disabled to run tests, the test status shows up as disabled.

Tip If the specific module or interface is enabled to run tests, but is not running the tests due to system health being disabled, then tests show up as enabled (not running).

Tests for a Specified Module

The system health feature in the NX-OS software performs tests in the following areas:

- Active supervisor’s in-band connectivity to the fabric.
- Standby supervisor’s arbiter availability.
- Bootflash connectivity and accessibility on all modules.
- EOBC connectivity and accessibility on all modules.
- Data path integrity for each interface on all modules.
- Management port’s connectivity.
- User-driven test for external connectivity verification, port is shut down during the test (Fibre Channel ports only).
- User-driven test for internal connectivity verification (Fibre Channel and iSCSI ports).

Note In Cisco MDS 9700 Series Switches, iSCSI ports are not applicable.

Clearing Previous Error Reports

You can clear the error history for Fibre Channel interfaces, iSCSI interfaces, an entire module, or one particular test for an entire module. By clearing the history, you are directing the software to retest all failed components that were previously excluded from tests.

If you previously enabled the failure-action option for a period of time (for example, one week) to prevent OHMS from taking any action when a failure is encountered and after that week you are now ready to start receiving these errors again, then you must clear the system health error status for each test.

Tip The management port test cannot be run on a standby supervisor module.

Interpreting the Current Status

The status of each module or test depends on the current configured state of the OHMS test in that particular module (see Table 20: OHMS Configured Status for Tests and Modules , on page 130).

Table 20: OHMS Configured Status for Tests and Modules

Status	Description
Enabled	You have currently enabled the test in this module and the test is not running.
Disabled	You have currently disabled the test in this module.
Running	You have enabled the test and the test is currently running in this module.
Failing	This state is displayed if a failure is imminent for the test running in this module—possibility of test recovery exists in this state.
Failed	The test has failed in this module—and the state cannot be recovered.
Stopped	The test has been internally stopped in this module by the Cisco NX-OS software.

Status	Description
Internal failure	The test encountered an internal failure in this module. For example, the system health application is not able to open a socket as part of the test procedure.
Diags failed	The startup diagnostics has failed for this module or interface.
On demand	The system health external-loopback or the system health internal-loopback tests are currently running in this module. Only these two commands can be issued on demand.
Suspended	Only encountered in the MDS 9100 Series due to one oversubscribed port moving to a E or TE port mode. If one oversubscribed port moves to this mode, the other three oversubscribed ports in the group are suspended.

The status of each test in each module is visible when you display any of the **show system health** commands. See the Displaying System Health , on page 147.

On-Board Failure Logging

The Generation 2 Fibre Channel switching modules provide the facility to log failure data to persistent storage, which can be retrieved and displayed for analysis. This on-board failure logging (OBFL) feature stores failure and environmental information in nonvolatile memory on the module. The information will help in post-mortem analysis of failed cards.

OBFL data is stored in the existing CompactFlash on the module. OBFL uses the persistent logging (PLOG) facility available in the module firmware to store data in the CompactFlash. It also provides the mechanism to retrieve the stored data.

The data stored by the OBFL facility includes the following:

- Time of initial power-on
- Slot number of the card in the chassis
- Initial temperature of the card
- Firmware, BIOS, FPGA, and ASIC versions
- Serial number of the card
- Stack trace for crashes
- CPU hog information
- Memory leak information
- Software error messages
- Hardware exception logs
- Environmental history
- OBFL specific history information
- ASIC interrupt and error statistics history
- ASIC register dumps

Core Files

Core files are created when NX-OS encounters an unrecoverable fault. These are a bundle of files in *tar.gz* format and can be used by Cisco to diagnose the fault.

NX-OS can generate process and kernel core file from both supervisor and modules. Process core files are uploaded from the module they occur on to the active supervisor at the time of the fault. The core files are volatile and will be lost if the supervisor resets. Kernel core files are stored on the supervisor they were created on and are retained after a supervisor reset.

First and Last Core

Generally, the first and the most recent cores that are generated by a process have the most useful information for debugging. If the core files are generated on the active supervisor module, the first and last core feature automatically deletes the intermediate cores if a new one is generated for the same process to conserve space in the core repository.

Default Settings

Table 21: Default System Status Monitoring, on page 132 lists the default settings.

Table 21: Default System Status Monitoring

Parameters	Default
Kernel core collection	Disabled
System health	Enabled
Loopback frequency	5 seconds
Failure action	Enabled

Configuring System Health

The Online Health Management System (OHMS) (system health) is a hardware fault detection and recovery feature. It ensures the general health of switching, services, and supervisor modules in any switch in the Cisco MDS 9000 Family.

Task Flow for Configuring System Health

Follow these steps to configure system health:

Procedure

Step 1 Enable System Health Initiation.
Step 2 Configure Loopback Test Configuration Frequency.
Step 3 Configure Loopback Test Configuration Frame Length.
Step 4 Configure Hardware Failure Action.
Step 5 Perform Test Run Requirements.
Step 6 Clear Previous Error Reports.

Step 7 Perform Internal Loopback Tests.

Step 8 Perform External Loopback Tests.

Step 9 Perform Serdes Loopbacks.

Configuring System Health Initiation

By default, the system health feature is enabled in each switch in the Cisco MDS 9000 Family.

To disable or enable this feature in any switch in the Cisco MDS 9000 Family, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **no system health**

System Health is disabled.

Disables system health from running tests in this switch.

Step 3 switch(config)# **system health**

System Health is enabled.

Enables (default) system health to run tests in this switch.

Step 4 switch(config)# **no system health interface fc8/1**

System health for interface fc8/13 is disabled.

Disables system health from testing the specified interface.

Step 5 switch(config)# **system health interface fc8/1**

System health for interface fc8/13 is enabled.

Enables (default) system health to test for the specified interface.

Configuring Loopback Test Configuration Frequency

To configure the frequency of loopback tests for all modules on a switch, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **system health loopback frequency 50**

```
The new frequency is set at 50 Seconds.
```

Configures the loopback frequency to 50 seconds. The default loopback frequency is 5 seconds. The valid range is from 5 to 255 seconds.

Configuring Loopback Test Configuration Frame Length

To configure the frame length for loopback tests for all modules on a switch, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **system health loopback frame-length 128**

Configures the loopback frame length to 128 bytes. The valid range is 0 to 128 bytes.

Step 3 switch(config)# **system health loopback frame-length auto**

Configures the loopback frame length to automatically generate random lengths (default).

Configuring Hardware Failure Action

To configure failure action in a switch, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **system health failure-action**

```
System health global failure action is now enabled.
```

Enables the switch to take failure action (default).

Step 3 switch(config)# **no system health failure-action**

```
System health global failure action now disabled.
```

Reverts the switch configuration to prevent failure action being taken.

Step 4 switch(config)# **system health module 1 failure-action**

```
System health failure action for module 1 is now enabled.
```

Enables switch to take failure action for failures in module 1.

Step 5 switch(config)# **no system health module 1 loopback failure-action**

```
System health failure action for module 1 loopback test is now disabled.
```

Prevents the switch from taking action on failures determined by the loopback test in module 1.

Performing Test Run Requirements

To perform the required test on a specific module, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Note The following steps can be performed in any order.

Note The various options for each test are described in the next step. Each command can be configured in any order. The various options are presented in the same step for documentation purposes.

Step 2 switch(config)# **system health module 8 bootflash**

Enables the bootflash test on module in slot 8.

Step 3 switch(config)# **system health module 8 bootflash frequency 200**

Sets the new frequency of the bootflash test on module 8 to 200 seconds.

Step 4 switch(config)# **system health module 8 eobc**

Enables the EOBC test on module in slot 8.

Step 5 switch(config)# **system health module 8 loopback**

Enables the loopback test on module in slot 8.

Step 6 switch(config)# **system health module 5 management**

Enables the management test on module in slot 5.

Clearing Previous Error Reports

Use the EXEC-level **system health clear-errors** command at the interface or module level to erase any previous error conditions logged by the system health application. The **bootflash**, the **eobc**, the **inband**, the **loopback**, and the **mgmt** test options can be individually specified for a given module.

The following example clears the error history for the specified Fibre Channel interface:

```
switch# system health clear-errors interface fc 3/1
```

The following example clears the error history for the specified module:

```
switch# system health clear-errors module 3
```

The following example clears the management test error history for the specified module:

```
switch# system health clear-errors module 1 mgmt
```

Performing Internal Loopback Tests

You can run manual loopback tests to identify hardware errors in the data path in the switching or services modules, and the control path in the supervisor modules. Internal loopback tests send and receive FC2 frames to and from the same ports and provide the round-trip time taken in microseconds. These tests are available for Fibre Channel, IPS, and iSCSI interfaces.

Use the EXEC-level **system health internal-loopback** command to explicitly run this test on demand (when requested by the user) within ports for the entire module.

```
switch# system health internal-loopback interface iscsi 8/1
Internal loopback test on interface iscsi8/1 was  successful.
Sent 1 received 1 frames
Round trip time taken is 79 useconds
```

Use the EXEC-level **system health internal-loopback** command to explicitly run this test on demand (when requested by the user) within ports for the entire module and override the frame count configured on the switch.

```
switch# system health internal-loopback interface iscsi 8/1 frame-count 20
Internal loopback test on interface iscsi8/1 was  successful.
Sent 1 received 1 frames
Round trip time taken is 79 useconds
```

Use the EXEC-level **system health internal-loopback** command to explicitly run this test on demand (when requested by the user) within ports for the entire module and override the frame length configured on the switch.

```
switch# system health internal-loopback interface iscsi 8/1 frame-count 32
Internal loopback test on interface iscsi8/1 was  successful.
Sent 1 received 1 frames
Round trip time taken is 79 useconds
```

Note If the test fails to complete successfully, the software analyzes the failure and prints the following error: External loopback test on interface fc 7/2 failed. Failure reason: Failed to loopback, analysis complete Failed device ID 3 on module 1

Performing External Loopback Tests

You can run manual loopback tests to identify hardware errors in the data path in the switching or services modules, and the control path in the supervisor modules. External loopback tests send and receive FC2 frames to and from the same port or between two ports.

You need to connect a cable (or a plug) to loop the Rx port to the Tx port before running the test. If you are testing to and from the same port, you need a special loop cable. If you are testing to and from different ports, you can use a regular cable. This test is only available for Fibre Channel interfaces.

Use the EXEC-level **system health external-loopback interface** *interface* command to run this test on demand for external devices connected to a switch that is part of a long-haul network.

```
switch# system health external-loopback interface fc 3/1
This will shut the requested interfaces Do you want to continue (y/n)?  [n] y
External loopback test on interface fc3/1 was  successful.
Sent 1 received 1 frames
```

Use the EXEC-level **system health external-loopback source** *interface* **destination interface** *interface* command to run this test on demand between two ports on the switch.

```
switch# system health external-loopback source interface fc 3/1 destination interface fc
3/2
This will shut the requested interfaces Do you want to continue (y/n)?  [n] y
External loopback test on interface fc3/1 and interface fc3/2 was  successful.
Sent 1 received 1 frames
```

Use the EXEC-level **system health external-loopback** *interface* **frame-count** command to run this test on demand for external devices connected to a switch that is part of a long-haul network and override the frame count configured on the switch.

```
switch# system health external-loopback interface fc 3/1 frame-count 10
This will shut the requested interfaces Do you want to continue (y/n)?  [n] y
External loopback test on interface fc3/1 was  successful.
Sent 1 received 1 frames
```

Use the EXEC-level **system health external-loopback** *interface* **frame-length** command to run this test on demand for external devices connected to a switch that is part of a long-haul network and override the frame length configured on the switch.

```
switch# system health external-loopback interface fc 3/1 frame-length 64
This will shut the requested interfaces Do you want to continue (y/n)?  [n] y
External loopback test on interface fc3/1 was  successful.
Sent 1 received 1 frames
```

Use the **system health external-loopback** *interface* **force** command to shut down the required interface directly without a back out confirmation.

```
switch# system health external-loopback interface fc 3/1 force
External loopback test on interface fc3/1 was  successful.
Sent 1 received 1 frames
```

Note If the test fails to complete successfully, the software analyzes the failure and prints the following error: External loopback test on interface fc 7/2 failed. Failure reason: Failed to loopback, analysis complete Failed device ID 3 on module 1

Performing Serdes Loopbacks

Serializer/Deserializer (serdes) loopback tests the hardware for a port. These tests are available for Fibre Channel interfaces.

Use the EXEC-level **system health serdes-loopback** command to explicitly run this test on demand (when requested by the user) within ports for the entire module.

```
switch# system health serdes-loopback interface fc 3/1
This will shut the requested interfaces Do you want to continue (y/n)?  [n] y
Serdes loopback test passed for module 3 port 1
```

Use the EXEC-level **system health serdes-loopback** command to explicitly run this test on demand (when requested by the user) within ports for the entire module and override the frame count configured on the switch.

```
switch# system health serdes-loopback interface fc 3/1 frame-count 10
This will shut the requested interfaces Do you want to continue (y/n)?  [n] y
Serdes loopback test passed for module 3 port 1
```

Use the EXEC-level **system health serdes-loopback** command to explicitly run this test on demand (when requested by the user) within ports for the entire module and override the frame length configured on the switch.

```
switch# system health serdes-loopback interface fc 3/1 frame-length 32
This will shut the requested interfaces Do you want to continue (y/n)?  [n] y
Serdes loopback test passed for module 3 port 1
```

Note If the test fails to complete successfully, the software analyzes the failure and prints the following error: External loopback test on interface fc 3/1 failed. Failure reason: Failed to loopback, analysis complete Failed device ID 3 on module 3.

Configuring On-Board Failure Logging

Each hardware module logs failure data to on-module persistent storage, which can be retrieved and displayed for analysis. This on-board failure logging (OBFL) feature stores failure and environmental information in nonvolatile memory on the module. The information will help in post-mortem analysis of failed cards.

Configuring OBFL for a Switch

To configure OBFL for all the modules on the switch, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **hw-module logging onboard**

Enables all OBFL features.

Note This CLI only enable OBFL features that are disabled by no hw-module logging onboard command. For OBFL features that were individually disabled, please enable those using hw-module logging onboard obfl-feature command.

Step 3 switch(config)# **hw-module logging onboard cpu-hog**

Enables the OBFL CPU hog events.

Step 4 switch(config)# **hw-module logging onboard environmental-history**

Enables the OBFL environmental history.

Step 5 switch(config)# **hw-module logging onboard error-stats**

Enables the OBFL error statistics.

Step 6 switch(config)# **hw-module logging onboard interrupt-stats**

Enables the OBFL interrupt statistics.

Step 7 switch(config)# **hw-module logging onboard mem-leak**

Enables the OBFL memory leak events.

Step 8 switch(config)# **hw-module logging onboard miscellaneous-error**

Enables the OBFL miscellaneous information.

Step 9 switch(config)# **hw-module logging onboard obfl-log**

Enables the boot uptime, device version, and OBFL history.

Step 10 switch(config)# **no hw-module logging onboard**

Disables all OBFL features.

Configuring OBFL for a Module

To configure OBFL for specific modules on the switch, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **hw-module logging onboard module 1**

Enables all OBFL features on a module.

Step 3 switch(config)# **hw-module logging onboard module 1 cpu-hog**

Enables the OBFL CPU hog events on a module.

Step 4 switch(config)# **hw-module logging onboard module 1 environmental-history**

Enables the OBFL environmental history on a module.

Step 5 switch(config)# **hw-module logging onboard module 1 error-stats**

Enables the OBFL error statistics on a module.

Step 6 switch(config)# **hw-module logging onboard module 1 interrupt-stats**

Enables the OBFL interrupt statistics on a module.

Step 7 switch(config)# **hw-module logging onboard module 1 mem-leak**

Enables the OBFL memory leak events on a module.

Step 8 switch(config)# **hw-module logging onboard module 1 miscellaneous-error**

Enables the OBFL miscellaneous information on a module.

Step 9 switch(config)# **hw-module logging onboard module 1 obfl-log**

Enables the boot uptime, device version, and OBFL history on a module.

Step 10 switch(config)# **no hw-module logging onboard module 1**

Disables all OBFL features on a module.

Clearing the Module Counters

Note The module counters cannot be cleared using Device Manager or DCNM-SAN.

To reset the module counters, follow these steps:

Procedure

Step 1 switch# **attach module 1**

ModuleX#

Attaches module 1 to the chasiss.

Step 2 ModuleX# **clear asic-cnt all**

Clears the counters for all the devices in the module.

Step 3 ModuleX# **clear asic-cnt list-all-devices**

ModuleX# **clear asic-cnt device-id** *device-id*

Clears the counters for only the specified device ID. The device ID can vary from 1 through 255.

Resetting Counters for All Modules

To reset the counters for all the modules, follow these steps:

Procedure

switch# **debug system internal clear-counters all**

Clears the counters for all the modules in the switch.

Configuring Alerts, Notifications, and Monitoring of Counters

This section provides information on how to configure alerts, notification, and monitor counters.

Monitoring the CPU Utilization

To display the system CPU utilization, use the **show processes cpu** command.

This example shows how to display processes and CPU usage in the current VDC:

```
switch# show processes cpu
PID    Runtime(ms)  Invoked   uSecs  1Sec    Process
-----  -----------  --------  -----  ------  -----------
    4       386829  67421866      5    0.9%  ksoftirqd/0
3667       270567    396229    682    9.8%  syslogd
3942          262       161   1632    7.8%  netstack
4006    106999945  354495641    301    28.2%  snmpd
4026      4454796    461564   9651    0.9%  sac_usd
4424        84187    726180    115    0.9%  vpc
4426       146378    919073    159    0.9%  tunnel
CPU util  :   25.0% user,   30.5% kernel,   44.5% idle
```

Obtaining RAM Usage Information

You can obtain the processor RAM usage by using this SNMP variable: ceExtProcessorRam.

```
ceExtProcessorRam OBJECT-TYPE
    SYNTAX  Unsigned32
    UNITS   "bytes"
        MAX-ACCESS  read-only
        STATUS  current
        DESCRIPTION
        "Total number of bytes of RAM available on the
         Processor."
        ::= { ceExtPhysicalProcessorEntry 1 }
```

Monitoring Rx and Tx Traffic Counters

When monitoring Rx and Tx traffic counters, you should include the Rx counter OID:

```
ifHCInOctets
```

Monitoring Status of Interfaces

To monitor status of interfaces, use the IETF extended-linkDown trap, which has ifAlias (this trap can set interface description) and ifDescr, which shows port name in the ASCII format as shown below:

```
switch (config)# snmp-server enable traps link
  cieLinkDown                 Cisco extended link state down notification
  cieLinkUp                   Cisco extended link state up notification
  cisco-xcvr-mon-status-chg   Cisco interface transceiver monitor status change
                             notification
  delayed-link-state-change  Delayed link state change
  extended-linkDown           IETF extended link state down notification
  extended-linkUp             IETF extended link state up notification
  linkDown                   IETF Link state down notification
  linkUp                      IETF Link state up notification
switch (config)#
```

The following is an example of the trap:

```
[+]          10           16:41:39.79         IF-MIB:linkDown trap:SNMPv2c   from
```

```
[172.25.234.200  Port: 162 Community: public]
SNMPv2-MIB:sysUpTime.0 :  (35519336)      Syntax: TimeTicks
SNMPv2-MIB:snmpTrapOID.0 :  (IF-MIB:linkDown)      Syntax: ObjectID
IF-MIB:ifIndex.440414208 :   (440414208)  Syntax: INTEGER,  Instance IDs: (440414208)
IF-MIB:ifAdminStatus.440414208 :   (down)  Syntax: INTEGER,  Instance IDs: (440414208)
IF-MIB:ifOperStatus.440414208 :   (down)  Syntax: INTEGER,  Instance IDs: (440414208)
IF-MIB:ifDescr.440414208 :   (Ethernet9/4)  Syntax: RFC1213-MIB:DisplayString,  Instance
IDs: (440414208)
IF-MIB:ifAlias.440414208 :   (eth9/4)  Syntax: SNMPv2-TC:DisplayString,  Instance IDs:
(440414208)
SNMPv2-MIB:snmpTrapEnterprise.0 :  (IF-MIB:linkDown)      Syntax: ObjectID
```

Monitoring Transceiver Thresholds

Use the cisco-xcvr-mon-status-chg trap way to monitor digital diagnostics statistics for thresholds as shown below:

```
switch (config)# snmp-server enable traps link cisco-xcvr-mon-status-chg
switch (config)#
```

The trap MIB is as show below:

```
cIfXcvrMonStatusChangeNotif NOTIFICATION-TYPE
    OBJECTS         {
                        ifName,
                        cIfXcvrMonDigitalDiagTempAlarm,
                        cIfXcvrMonDigitalDiagTempWarning,
                        cIfXcvrMonDigitalDiagVoltAlarm,
                        cIfXcvrMonDigitalDiagVoltWarning,
                        cIfXcvrMonDigitalDiagCurrAlarm,
                        cIfXcvrMonDigitalDiagCurrWarning,
                        cIfXcvrMonDigitalDiagRxPwrAlarm,
                        cIfXcvrMonDigitalDiagRxPwrWarning,
                        cIfXcvrMonDigitalDiagTxPwrAlarm,
                        cIfXcvrMonDigitalDiagTxPwrWarning,
                        cIfXcvrMonDigitalDiagTxFaultAlarm
                    }
    STATUS          current
```

This example shows how to display transceiver details:

```
switch(config)# show interface ethernet 1/17 transceiver details
Ethernet1/17
    transceiver is present
    type is 10Gbase-SR
    name is CISCO-AVAGO
    part number is SFBR-7702SDZ
    revision is G2.3
    serial number is AGA1427618P
    nominal bitrate is 10300 MBit/sec
    Link length supported for 50/125um OM2 fiber is 82 m
    Link length supported for 62.5/125um fiber is 26 m
    Link length supported for 50/125um OM3 fiber is 300 m
    cisco id is --
    cisco extended id number is 4
           SFP Detail Diagnostics Information (internal calibration)
  ----------------------------------------------------------------------------
                Current              Alarms                  Warnings
                Measurement     High        Low         High          Low
  ---------------------------------------------------------------------------
  Temperature   27.65 C        75.00 C     -5.00 C     70.00 C        0.00 C
  Voltage        3.29 V         3.63 V      2.97 V      3.46 V        3.13 V
  Current        5.42 mA       10.50 mA     2.50 mA    10.50 mA       2.50 mA
  Tx Power      -2.51 dBm       1.69 dBm  -11.30 dBm   -1.30 dBm     -7.30 dBm
  Rx Power      -2.64 dBm       1.99 dBm  -13.97 dBm   -1.00 dBm     -9.91 dBm
  Transmit Fault Count = 0
  ---------------------------------------------------------------------------
  Note: ++  high-alarm; +  high-warning; --  low-alarm; -  low-warning
switch(config)#
```

Configuring Supervisor Switchover Notification

The supervisor switchover notification can be monitored by listening for the ciscoRFSwactNotif trap:

```
                 ciscoRFSwactNotif NOTIFICATION-TYPE
                 OBJECTS {
                 cRFStatusUnitId,
                 sysUpTime,
                 cRFStatusLastSwactReasonCode
               }
```

Configuring a Counter to Include CRC and FCS Errors

You can include CRC and FCS errors of interfaces by polling dot3StatsFCSErrors counter as shown in the following example:

dot3StatsFCSErrors Counter32

```
            Dot3StatsEntry ::= SEQUENCE {
       dot3StatsIndex                      InterfaceIndex,
       dot3StatsAlignmentErrors            Counter32,
       dot3StatsFCSErrors                  Counter32,
       dot3StatsSingleCollisionFrames      Counter32,
       dot3StatsMultipleCollisionFrames    Counter32,
       dot3StatsSQETestErrors              Counter32,
       dot3StatsDeferredTransmissions      Counter32,
       dot3StatsLateCollisions             Counter32,
       dot3StatsExcessiveCollisions        Counter32,
       dot3StatsInternalMacTransmitErrors  Counter32,
       dot3StatsCarrierSenseErrors         Counter32,
       dot3StatsFrameTooLongs              Counter32,
       dot3StatsInternalMacReceiveErrors   Counter32,
       dot3StatsEtherChipSet               OBJECT IDENTIFIER,
       dot3StatsSymbolErrors               Counter32,
       dot3StatsDuplexStatus               INTEGER,
       dot3StatsRateControlAbility         TruthValue,
       dot3StatsRateControlStatus          INTEGER
      }
```

Configuring Call Home for Alerts

The Call Home feature enables you to receive a Call Home email when exceptions occur in the system. Use the following CLI or SNMP to set up the Call Home configurations and to enable all alert-groups:

```
switch (config)# callhome
switch-FC-VDC(config-callhome)# destination-profile full-txt-destination alert-group
```

```
All                     This alert group consists of all of the callhome
                        messages
  Cisco-TAC             Events which are meant for Cisco TAC only
  Configuration         Events related to Configuration
  Diagnostic            Events related to Diagnostic
  EEM                   EEM events
  Environmental         Power,fan,temperature related events
  Inventory             Inventory status events
  License               Events related to licensing
  Linecard-Hardware     Linecard related events
  Supervisor-Hardware   Supervisor related events
  Syslog-group-port     Events related to syslog messages filed by port manager
  System                Software related events
  Test                  User generated test events
switch-FC-VDC(config-callhome)#
```

Monitoring User Authentication Failures

You can monitor any user authentication failures by listening the authenticationFailure trap:

```
SNMPv2-MIB: authenticationFailure trap
```

Configuring Cores

Core files can be saved either manually by a user or automatically at the time of the fault. If a core file is created, preserve it by copying it to nonvolatile file space (such as to a host) and report it to Cisco for diagnosis.

Cores may be copied multiple times. Both IPv4, IPv6, and many protocols are supported for copying cores to file space on remote hosts. This includes passwordless SSH which is convenient for automatic copying in secure environments. For more information about configuring passwordless access to remote hosts, see the 'Passwordless File Copy and SSH' section in the 'Configuring SSH Services and Telnet' chapter of Cisco MDS 9000 Series Security Configuration Guide, Release 8.x.

There is no upper limit on the total number of core files in the active supervisor module.

Tip Before copying a core, ensure that you create the destination directory with write permission for the user.

Configuring Kernel Core Collection

To configure kernel core collection, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **system kernel core**

Enables collection of kernel core if there is a kernel crash.

Step 3 switch(config)# **no system kernel core**

(Optional) Disables collection of kernel cores.

Copying Cores Manually

The supported on-switch destination is slot0. Supported protocols to transfer cores to remote destinations are TFTP, SFTP, and SCP.

To configure saving cores manually, follow these steps:

Procedure

switch# **copy** *core://module/process-id[/instance] destination://[[user@]host/][directory]*

Copies a core of process to the specified location.

Copying Cores Automatically

Supported on-switch destinations are bootflash, slot0, and usb1. Supported protocols to transfer cores to remote destinations are HTTP, HTTPS, TFTP, FTP, SFTP, and SCP.

To configure saving cores automatically, follow these steps:

Procedure

Step 1 switch# **configure**

Enters configuration mode.

Step 2 switch(config)# **system cores** *destination://[[user@]host/][directory]*

Save core files to the specified destination as soon as they are created.

Step 3 switch(config)# **no system cores**

(Optional) Disables saving core files automatically.

Deleting Cores

Core files are not automatically deleted after copying. After a core is copied, delete it from the switch core repository to reclaim the space and report it to Cisco support for analysis.

Use the **clear core_file** command to delete a single core from the switch core repository.

switch# **clear core_file module** *module* **pid** *pid*

Use the **clear cores** command to clear all cores in the switch core repository.

switch# **clear cores**

Example: Configuring Cores

The following example copies a core of process with PID 1524 generated on slot 5 to the *cores* directory on a host with HTTPS as user *mdsadmin*:

```
switch# copy core://5/1524 https://mdsadmin@192.168.1.2/cores
```

The following example automatically copies any core files immediately after they are created to the /tftpboot/cores directory on a host with SCP as user mdsadmin. Configure passwordless SSH first for this to work.

```
switch# configure
switch(config)# system cores scp://mdsadmin@192.168.1.2/tftpboot/cores
```

The following example deletes the core generated from module 1 for the process with PID 1234.

```
switch# clear core_file module 1 pid 1234
```

Verifying System Status Monitoring Configuration

To display the system status monitoring configuration information, perform one of the following tasks:

Displaying System Health

Use the **show system health** command to display system-related status information (see Current Health of All Modules in the Switch, on page 147 to Loopback Test Time Log for a Specified Module, on page 150).

Current Health of All Modules in the Switch

The following example displays the current health of all modules in the switch:

```
switch# show system health

Current health information for module 2.
Test                       Frequency    Status          Action
-----------------------------------------------------------------
Bootflash                 5 Sec         Running         Enabled
EOBC                      5 Sec         Running         Enabled
Loopback                  5 Sec         Running         Enabled
-----------------------------------------------------------------
Current health information for module 6.
Test                       Frequency    Status          Action
-----------------------------------------------------------------
InBand                    5 Sec         Running         Enabled
Bootflash                 5 Sec         Running         Enabled
EOBC                      5 Sec         Running         Enabled
```

```
Management Port           5 Sec         Running         Enabled
-----------------------------------------------------------------
```

Current Health of a Specified Module

The following example displays the current health of a specified module:

```
switch# show system health module 8
Current health information for module 8.
Test                       Frequency    Status          Action
-----------------------------------------------------------------
Bootflash                 5 Sec         Running         Enabled
EOBC                      5 Sec         Running         Enabled
Loopback                  5 Sec         Running         Enabled
-----------------------------------------------------------------
```

Health Statistics for All Modules

The following example displays health statistics for all modules:

```
switch# show system health statistics
Test statistics for module # 1
------------------------------------------------------------------------------
Test Name           State            Frequency Run    Pass    Fail CFail Errs
------------------------------------------------------------------------------
Bootflash           Running             5s    12900   12900       0     0    0
EOBC                Running             5s    12900   12900       0     0    0
Loopback            Running             5s    12900   12900       0     0    0
------------------------------------------------------------------------------
Test statistics for module # 3
------------------------------------------------------------------------------
Test Name           State            Frequency Run    Pass    Fail CFail Errs
------------------------------------------------------------------------------
Bootflash           Running             5s    12890   12890       0     0    0
EOBC                Running             5s    12890   12890       0     0    0
Loopback            Running             5s    12892   12892       0     0    0
------------------------------------------------------------------------------
Test statistics for module # 5
------------------------------------------------------------------------------
Test Name           State            Frequency Run    Pass    Fail CFail Errs
------------------------------------------------------------------------------
InBand              Running             5s    12911   12911       0     0    0
Bootflash           Running             5s    12911   12911       0     0    0
EOBC                Running             5s    12911   12911       0     0    0
Management Port     Running             5s    12911   12911       0     0    0
------------------------------------------------------------------------------
Test statistics for module # 6
------------------------------------------------------------------------------
Test Name           State            Frequency Run    Pass    Fail CFail Errs
------------------------------------------------------------------------------
InBand              Running             5s    12907   12907       0     0    0
Bootflash           Running             5s    12907   12907       0     0    0
EOBC                Running             5s    12907   12907       0     0    0
------------------------------------------------------------------------------
Test statistics for module # 8
------------------------------------------------------------------------------
Test Name           State            Frequency Run    Pass    Fail CFail Errs
------------------------------------------------------------------------------
Bootflash           Running             5s    12895   12895       0     0    0
```

```
EOBC                Running             5s    12895   12895       0     0    0
Loopback            Running             5s    12896   12896       0     0    0
------------------------------------------------------------------------------
```

Displays Statistics for a Specified Module

The following example displays statistics for a specified module:

```
switch# show system health statistics module 3
Test statistics for module # 3
------------------------------------------------------------------------------
Test Name           State            Frequency Run    Pass    Fail CFail Errs
------------------------------------------------------------------------------
Bootflash           Running             5s    12932   12932       0     0    0
EOBC                Running             5s    12932   12932       0     0    0
Loopback            Running             5s    12934   12934       0     0    0
------------------------------------------------------------------------------
```

Loopback Test Statistics for the Entire Switch

The following example displays loopback test statistics for the entire switch:

```
switch# show system health statistics loopback
-----------------------------------------------------------------
Mod Port Status                Run     Pass     Fail   CFail Errs
  1   16 Running             12953    12953        0       0    0
  3   32 Running             12945    12945        0       0    0
  8    8 Running             12949    12949        0       0    0
-----------------------------------------------------------------
```

Loopback Test Statistics for a Specified Interface

The following example displays loopback test statistics for a specified interface:

```
switch# show system health statistics loopback interface fc 3/1
-----------------------------------------------------------------
Mod Port Status                Run     Pass     Fail   CFail Errs
  3    1 Running                 0        0        0       0    0
-----------------------------------------------------------------
```

Note Interface-specific counters will remain at zero unless the module-specific loopback test reports errors or failures.

Loopback Test Time Log for All Modules

The following example displays loopback test time log for all modules:

```
switch# show system health statistics loopback timelog
-----------------------------------------------------------------
Mod        Samples     Min(usecs)     Max(usecs)     Ave(usecs)
```

```
  1          1872           149           364           222
  3          1862           415           743           549
  8          1865           134           455           349
-----------------------------------------------------------------
```

Loopback Test Time Log for a Specified Module

The following example displays the loopback test time log for a specified module:

```
switch# show system health statistics loopback module 8 timelog
-----------------------------------------------------------------
Mod        Samples     Min(usecs)     Max(usecs)     Ave(usecs)
  8          1867           134           455           349
-----------------------------------------------------------------
```

Verifying Loopback Test Configuration Frame Length

To verify the loopback frequency configuration, use the **show system health loopback frame-length** command.

```
switch# show system health loopback frame-length
Loopback frame length is set to auto-size between 0-128 bytes
```

Verifying OBFL for the Switch

Use the **show logging onboard status** command to display the configuration status of OBFL.

```
switch# show logging onboard status
    Switch OBFL Log:                                    Enabled
    Module:  6 OBFL Log:                                Enabled
    error-stats                                         Enabled
    exception-log                                       Enabled
    miscellaneous-error                                 Enabled
    obfl-log (boot-uptime/device-version/obfl-history)  Enabled
    system-health                                       Enabled
    stack-trace                                         Enabled
```

Verifying OBFL for a Module

Use the **show logging onboard status** command to display the configuration status of OBFL.

```
switch# show logging onboard status
    Switch OBFL Log:                                    Enabled
    Module:  6 OBFL Log:                                Enabled
    error-stats                                         Enabled
    exception-log                                       Enabled
    miscellaneous-error                                 Enabled
    obfl-log (boot-uptime/device-version/obfl-history)  Enabled
    system-health                                       Enabled
    stack-trace                                         Enabled
```

Verifying Kernel Core Collection

The kernel core collection configuration may be verified by checking the running configuration.

```
switch# show running-config | include 'kernel core'
system kernel core
```

Verifying Automatic Core Copying

Use the **show system cores** command to display the configuration of the automatic core copying feature.

```
switch# show system cores
Cores are transferred to scp://mdsadmin@192.168.1.2/tftpboot/cores
```

Displaying OBFL Logs

To display OBFL information stored on a module, use the following commands:

Command	Purpose
show logging onboard boot-uptime	Displays the boot and uptime information.
show logging onboard counter-stats	Displays counter statistics. **Note** In Cisco MDS 9132T and Cisco MDS 9396T switches, the output of this command displays information about removed LEM ports.
show logging onboard cpu-hog	Displays information for CPU hog events.
show logging onboard device-version	Displays device version information.
show logging onboard endtime	Displays OBFL logs to an end time.
show logging onboard environmental-history	Displays environmental history.
show logging onboard error-stats	Displays error statistics.
show logging onboard exception-log	Displays exception log information.
show logging onboard interrupt-stats	Displays interrupt statistics.
show logging onboard mem-leak	Displays memory leak information.
show logging onboard miscellaneous-error	Displays miscellaneous error information.
show logging onboard module *slot*	Displays OBFL information for a specific module.
show logging onboard obfl-history	Displays history information.

Command	Purpose
show logging onboard register-log	Displays register log information.
show logging onboard stack-trace	Displays kernel stack trace information.
show logging onboard starttime	Displays OBFL logs from a specified start time.
show logging onboard system-health	Displays system health information.

Displaying the Module Counters Information

This example shows the device IDs of all the devices in a module:

```
switch# attach  module 4
Attaching to module 4 ...
To exit type 'exit', to abort type '$.'
Linux lc04 2.6.10_mvl401-pc_target #1 Tue Dec 16 22:58:32 PST 2008 ppc GNU/Linux

module-4# clear asic-cnt list-all-devices
           Asic Name  |        Device ID
        Stratosphere  |              63
         transceiver  |              46
        Skyline-asic  |              57
          Skyline-ni  |              60
        Skyline-xbar  |              59
         Skyline-fwd  |              58
        Tuscany-asic  |              52
        Tuscany-xbar  |              54
         Tuscany-que  |              55
         Tuscany-fwd  |              53
       Fwd-spi-group  |              73
          Fwd-parser  |              74
                eobc  |              10
            X-Bus IO  |               1
   Power Mngmnt Epld  |              25
```

Displaying System Processes

Use the **show processes** command to obtain general information about all processes (see CPU Utilization Information, on page 153 to Memory Information About Processes , on page 155).

Displays System Processes

The following example displays system processes

```
switch# show processes

PID    State  PC        Start_cnt    TTY   Process
-----  -----  --------  -----------  ----  -------------
  868      S  2ae4f33e           1      -  snmpd
  869      S  2acee33e           1      -  rscn
  870      S  2ac36c24           1      -  qos
  871      S  2ac44c24           1      -  port-channel
  872      S  2ac7a33e           1      -  ntp
    -     ER         -           1      -  mdog
    -     NR         -            0     -  vbuilder
```

Where:

- ProcessId = Process ID
- State = process state.
    - D = uninterruptible sleep (usually I/O).
    - R = runnable (on run queue).
    - S = sleeping.
    - T = traced or stopped.
    - Z = defunct (“zombie”) process.
- NR = not running.
- ER = should be running but currently not-running.
- PC = current program counter in hex format.
- Start_cnt = number of times a process has been started (or restarted).
- TTY = terminal that controls the process. A hyphen usually means a daemon not running on any particular TTY.
- Process Name = name Name of the process.

CPU Utilization Information

The following example displays CPU Utilization Information

```
switch# show processes cpu
PID    Runtime(ms)  Invoked   uSecs  1Sec   Process
-----  -----------  --------  -----  -----  -----------
  842         3807    137001     27    0.0  sysmgr
 1112         1220     67974     17    0.0  syslogd
 1269          220     13568     16    0.0  fcfwd
 1276         2901     15419    188    0.0  zone
 1277          738     21010     35    0.0  xbar_client
 1278         1159      6789    170    0.0  wwn
 1279          515     67617      7    0.0  vsan
```

Where:

- MemAllocated = Sum of all the dynamically allocated memory that this process has received from the system, including memory that may have been returned
- Runtime CPU Time (ms) = CPU time the process has used, expressed in milliseconds.microseconds
- Invoked = number of times the process has been invoked.
- uSecs = microseconds of CPU time on average for each process invocation.
- 1Sec = CPU utilization in percentage for the last one second.

Process Log Information

The following example displays process log information:

```
switch# show processes log
Process           PID     Normal-exit  Stack-trace  Core     Log-create-time
----------------  ------  -----------  -----------  -------  ---------------
```

```
fspf             1339             N            Y        N  Jan  5 04:25
lcm              1559             N            Y        N  Jan  2 04:49
rib              1741             N            Y        N  Jan  1 06:05
```

Where:

- Normal-exit = whether or not the process exited normally.
- Stack-trace = whether or not there is a stack trace in the log.
- Core = whether or not there exists a core file.
- Log-create-time = when the log file got generated.

Detail Log Information About a Process

The following example displays detail log information about a process

```
switch# show processes log pid 1339

Service: fspf
Description: FSPF Routing Protocol Application
Started at Sat Jan  5 03:23:44 1980 (545631 us)
Stopped at Sat Jan  5 04:25:57 1980 (819598 us)
Uptime: 1 hours 2 minutes 2 seconds
Start type: SRV_OPTION_RESTART_STATELESS (23)
Death reason: SYSMGR_DEATH_REASON_FAILURE_SIGNAL (2)
Exit code: signal 9 (no core)
CWD: /var/sysmgr/work
Virtual Memory:
    CODE      08048000 - 0809A100
    DATA      0809B100 - 0809B65C
    BRK       0809D988 - 080CD000
    STACK     7FFFFD20
    TOTAL     23764 KB
Register Set:
    EBX 00000005         ECX 7FFFF8CC         EDX 00000000
    ESI 00000000         EDI 7FFFF6CC         EBP 7FFFF95C
    EAX FFFFFDFE         XDS 8010002B         XES 0000002B
    EAX 0000008E (orig)  EIP 2ACE133E         XCS 00000023
    EFL 00000207         ESP 7FFFF654         XSS 0000002B
Stack: 1740 bytes. ESP 7FFFF654, TOP 7FFFFD20
0x7FFFF654: 00000000 00000008 00000003 08051E95 ................
0x7FFFF664: 00000005 7FFFF8CC 00000000 00000000 ................
0x7FFFF674: 7FFFF6CC 00000001 7FFFF95C 080522CD ........\....".. 
0x7FFFF684: 7FFFF9A4 00000008 7FFFFC34 2AC1F18C ........4......*
```

All Process Log Details

The following example displays all process log details

```
switch# show processes log details
======================================================
Service: snmpd
Description: SNMP Agent
Started at Wed Jan  9 00:14:55 1980 (597263 us)
Stopped at Fri Jan 11 10:08:36 1980 (649860 us)
Uptime: 2 days 9 hours 53 minutes 53 seconds
Start type: SRV_OPTION_RESTART_STATEFUL (24)
Death reason: SYSMGR_DEATH_REASON_FAILURE_SIGNAL (2)
Exit code: signal 6 (core dumped)
```

```
CWD: /var/sysmgr/work
Virtual Memory:
    CODE      08048000 - 0804C4A0
    DATA      0804D4A0 - 0804D770
    BRK       0804DFC4 - 0818F000
    STACK     7FFFFCE0
    TOTAL     26656 KB
...
```

Memory Information About Processes

The following example displays memory information about processes

```
switch# show processes memory
PID    MemAlloc  MemLimit    MemUsed     StackBase/Ptr      Process
-----  --------  ----------  ----------  -----------------  ----------------
    1    147456  0      1667072     7ffffe50/7ffff950      init
    2         0  0           0                  0/0         ksoftirqd/0
    3         0  0           0                  0/0         desched/0
    4         0  0           0                  0/0         events/0
    5         0  0           0                  0/0         khelper
```

Where:

- MemAlloc = total memory allocated by the process.
- StackBase/Ptr = process stack base and current stack pointer in hex format.

Displaying System Status

Use the **show system** command to display system-related status information (see Default Switch Port States, on page 155 to System Related CPU and Memory Information, on page 156).

Default Switch Port States

The following example displays default switch port states:

```
switch# show system default switchport
System default port state is down
System default trunk mode is on
```

Error Information for a Specified ID

The following example displays error information for a specified ID:

```
switch# show system error-id 0x401D0019
Error Facility: module
Error Description: Failed to stop Linecard Async Notification.
```

System Reset Information

The following example displays the System Reset Information:

```
switch# Show system reset-reason module 5
----- reset reason for module 5 -----
1) At 224801 usecs after Fri Nov 21 16:36:40 2003
    Reason: Reset Requested by CLI command reload
    Service:
    Version: 1.3(1)
2) At 922828 usecs after Fri Nov 21 16:02:48 2003
    Reason: Reset Requested by CLI command reload
    Service:
    Version: 1.3(1)
3) At 318034 usecs after Fri Nov 21 14:03:36 2003
    Reason: Reset Requested by CLI command reload
    Service:
    Version: 1.3(1)
4) At 255842 usecs after Wed Nov 19 00:07:49 2003
    Reason: Reset Requested by CLI command reload
    Service:
    Version: 1.3(1)
```

The **show system reset-reason** command displays the following information:

- In a Cisco MDS 9700 Director, the last four reset-reason codes for the supervisor modules are displayed. If either supervisor module is absent, the reset-reason codes for that supervisor module are not displayed..
- In a Cisco MDS 9000 Series Fabric Switch, the last four reset-reason codes for the supervisor are displayed. The supervisor on a fabric switch is represented as module in slot 1.
- The show system reset-reason module number command displays the last four reset-reason codes for a specific module in a given slot. If a module is absent, then the reset-reason codes for that module are not displayed.

Use the **clear system reset-reason** command to clear the reset-reason information stored in NVRAM and volatile persistent storage.

- In a Cisco MDS 9700 Director, this command clears the reset-reason information stored in NVRAM in the active and standby supervisor modules.
- In a Cisco MDS 9000 Series Fabric Switch, this command clears the reset-reason information stored in NVRAM in the active supervisor module.

System Uptime

The following example displays system uptime:

```
switch# show system uptime
Start Time: Sun Oct 13 18:09:23 2030
Up Time:    0 days, 9 hours, 46 minutes, 26 seconds
```

Use the **show system resources** command to display system-related CPU and memory statistics (see System Related CPU and Memory Information, on page 156).

System Related CPU and Memory Information

The following example displays system related CPU and memory information:

```
switch# show system resources
Load average:   1 minute: 0.43   5 minutes: 0.17   15 minutes: 0.11
```

```
Processes   :   100 total, 2 running
CPU states  :   0.0% user,   0.0% kernel,   100.0% idle
Memory usage:   1027628K total,    313424K used,    714204K free
                   3620K buffers,   22278K cache
```

Where:

- Load average—Displays the number of running processes. The average reflects the system load over the past 1, 5, and 15 minutes.
- Processes—Displays the number of processes in the system, and how many are actually running when the command is issued.
- CPU states—Displays the CPU usage percentage in user mode, kernel mode, and idle time in the last one second.
- Memory usage—Displays the total memory, used memory, free memory, memory used for buffers, and memory used for cache in KB. Buffers and cache are also included in the *used* memory statistics.

Displaying Process Fault Logs

Displaying the Process Fault Log Summary

The history of fatal process faults and the logs collected of each event can be displayed on a per-module basis. Use the **slot** command to run the **show processes log** command on a particular module.

The following example displays the process fault log summary on module 2:

```
switch# slot 2 show processes log
Process            PID     Normal-exit  Stack  Core   Log-create-time
----------------   ------  -----------  -----  -----  ---------------
ExceptionLog       2862              N      Y      N  Wed Aug  6 15:08:34 2003
acl                2299              N      Y      N  Tue Oct 28 02:50:01 2003
bios_daemon        2227              N      Y      N  Mon Sep 29 15:30:51 2003
```

The following example displays the process cores of a system in Device Manager:

Figure 4: Show Cores Dialog Box

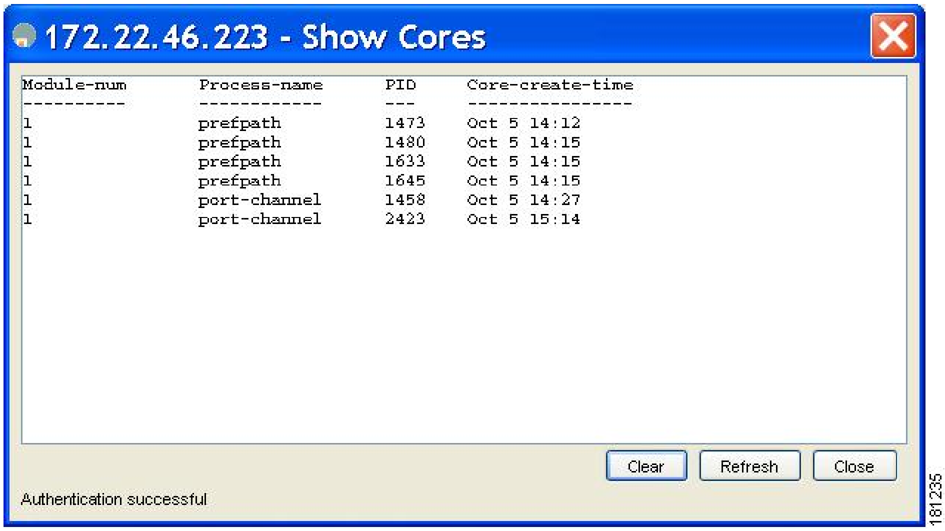

172.22.46.223 - Show Cores

```
Module-num     Process-name     PID    Core-create-time
----------     ------------     ---    ----------------
1              prefpath         1473   Oct 5 14:12
1              prefpath         1480   Oct 5 14:15
1              prefpath         1633   Oct 5 14:15
1              prefpath         1645   Oct 5 14:15
1              port-channel     1458   Oct 5 14:27
1              port-channel     2423   Oct 5 15:14
```

Authentication successful

Displaying Process Cores

The following example displays all cores stored on the active supervisor module:

```
switch# show cores
Module-num  Process-name    PID     Core-create-time
----------  ------------    ---     ----------------
5           fspf            1524    Nov 9 03:11
6           fcc             919     Nov 9 03:09
8           acltcam         285     Nov 9 03:09
8           fib             283     Nov 9 03:08
```

Additional References

For additional information related to implementing System Processes and Logs, see the following section:

MIBs

MIBs	MIBs Link
• CISCO-SYSTEM-EXT-MIB • CISCO-SYSTEM-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

CHAPTER 8

About Embedded Event Manager

This chapter describes how to configure the EEM to detect and handle critical events on a device.

Feature History for EEM

Table 22: Feature History for EEM , on page 159 lists the release history for this feature. Only features that were introduced or modified in Release 3.x or a later release appear in the table.

Table 22: Feature History for EEM

Feature Name	Releases	Feature Information
EEM	8.1(1)	Added **cli** keyword to the **action***number* command.
Zone, FCNS, and FLOGI	6.2(11)	This feature enables users to configure custom limits for default Zone, FCNS, and FLOGI system policies.
EEM	4.1(3)	New chapter on configuring Embedded Event Manager (EEM) has been added.

Information About EEM

Embedded Event Manager monitors events that occur on your device and takes action to recover or troubleshoot these events, based on your configuration.

EEM Overview

EEM consists of three major components:

- Event statements—Events to monitor from another Cisco NX-OS component that may require some action, workaround, or notification.
- Action statements —Actions that EEM can take, such as sending an e-mail, or disabling an interface, to recover from an event.
- Policies—A combination of an event statement and an action statement. When the specified event occurs, the configured action is executed.

Policies

An EEM policy consists of an event statement and one or more action statements. The event statement defines the event to look for as well as the filtering characteristics for the event. The action statement defines the action EEM takes when the event occurs.

Figure 5: EEM Policy Statements, on page 160 shows the two basic statements in an EEM policy.

Figure 5: EEM Policy Statements

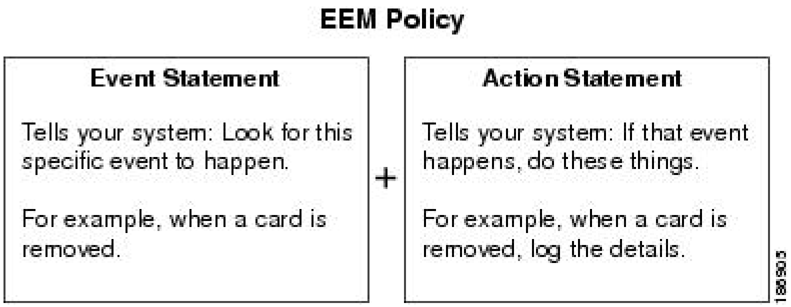
EEM Policy
Event Statement
Tells your system: Look for this specific event to happen.
For example, when a card is removed.
+
Action Statement
Tells your system: If that event happens, do these things.
For example, when a card is removed, log the details.
186905

You can configure EEM policies using the CLI or using a VSH script.

Note EEM policy matching is not supported on MDS switches.

EEM maintains event logs on the supervisor.

Cisco NX-OS has a number of preconfigured system policies. These system policies define many common events and actions for the device. System policy names begin with two underscore characters (__).

The following are some of the preconfigured system policies available in Cisco MDS 9000 Series Switches:

- Zone
  - __zone_dbsize_max_per_vsan : Syslog warning when Zone database size exceeds the max limit of 4000000 bytes for a vsan.
  - __zone_members_max_per_sw: Syslog warning when Zone member count exceeds the max limit of 32000 for the switch.
  - __zone_zones_max_per_sw: Syslog warning when Zone count exceeds the max limit of 16000 for the switch.

  - __zone_zonesets_max_per_sw: Syslog warning when Zoneset count exceeds the max limit of 1000 for the switch.
  - __zone_member_fan_out_ratio: Syslog warning when the number of devices exceed the specified fan-out-ratio limit.

- Fabric Login (FLOGI)
  - __flogi_fcids_max_per_switch: Syslog warning when the number of flogis in the switch exceeds 2000.
  - __flogi_fcids_max_per_module: Syslog warning when the number of flogis in the module exceeds 400.
  - __flogi_fcids_max_per_intf: Syslog warning when the number of flogis on the interface exceeds 256.

Note All the above three FLOGI policies are over ridable.

- Fibre Channel Name Server (FCNS)
  - __fcns_entries_max_per_switch : Configuring max limit for Name server entries verified across all VSANs per switch.

Action: Display a syslog

Note User should not configure an event for a different component's policy.

You can create user policies to suit your network. Actions defined by the user policies are executed along with the actions defined by the system policies. To configure a user policy, see the Defining a User Policy Using the CLI, on page 165.

You can also override some system policies. The override policies replace the system policies. You can override the event or the actions.

Use the **show event manager system-policy** command to view the preconfigured system policies and determine which policies that you can override.

To configure an overriding policy, see the Overriding a Policy, on page 174.

Note You should use the **show running-config eem** command to check the configuration of each policy. An override policy that consists of an event statement and no action statement triggers no action and no notification of failures.

Note Your override policy should always include an event statement. An override policy without an event statement overrides all possible events in the system policy.

Event Statements

An event is any device activity for which some action, such as a workaround or a notification, should be taken. In many cases, these events are related to faults in the device such as when an interface or a fan malfunctions.

Figure 6: EEM Overview, on page 162 EEM defines event filters so only critical events or multiple occurrences of an event within a specified time period trigger an associated action.

shows events that are handled by EEM.

Figure 6: EEM Overview

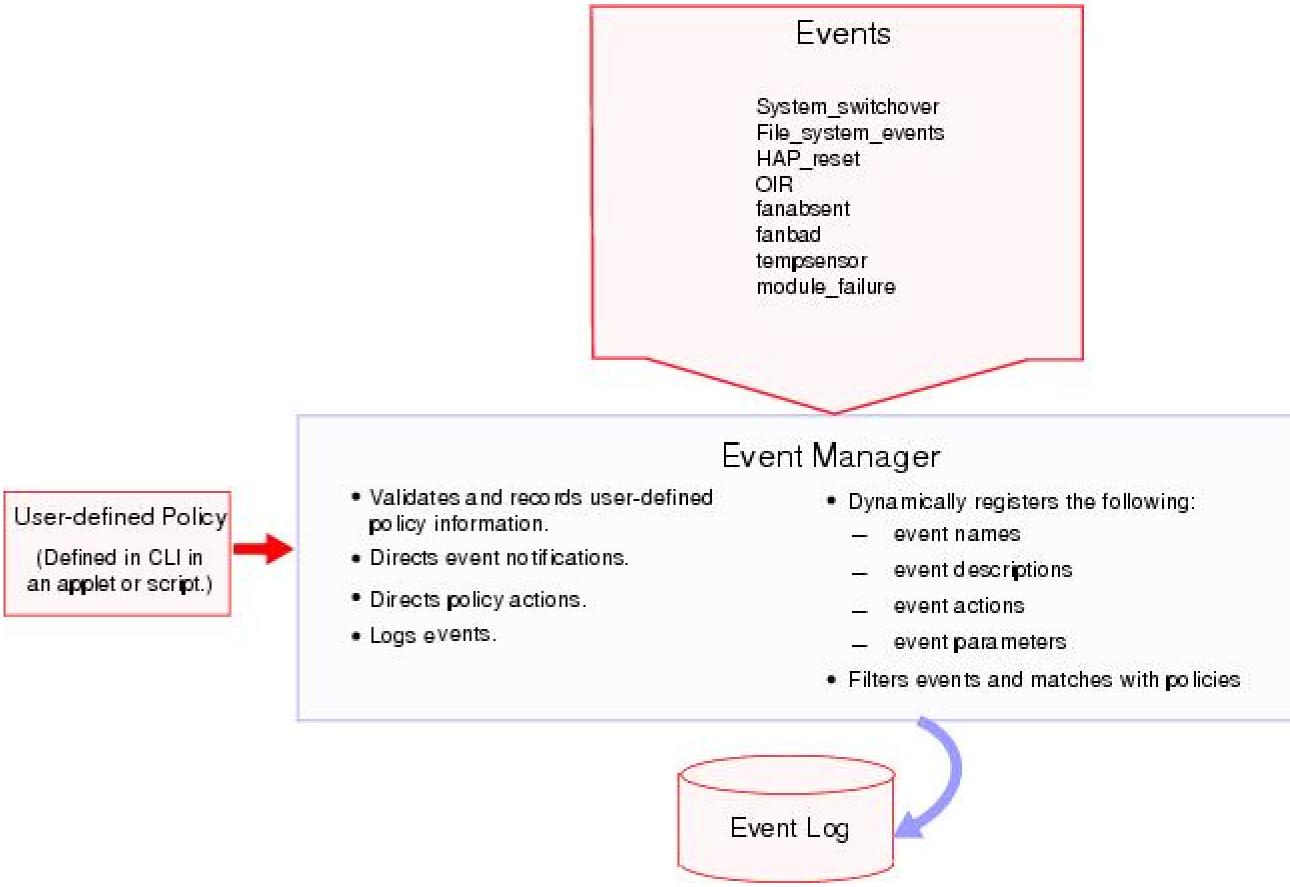
Events
System_switchover
File_system_events
HAP_reset
OIR
fanabsent
fanbad
tempsensor
module_failure
Event Manager
• Validates and records user-defined policy information.
• Directs event notifications.
• Directs policy actions.
• Logs events.
• Dynamically registers the following:
– event names
– event descriptions
– event actions
– event parameters
• Filters events and matches with policies
User-defined Policy
(Defined in CLI in an applet or script.)
Event Log

Event statements specify the event that triggers a policy to run. You can configure only one event statement per policy.

EEM schedules and runs policies on the basis of event statements. EEM examines the event and action commands and runs them as defined.

Action Statements

Action statements describe the action triggered by a policy. Each policy can have multiple action statements. If no action is associated with a policy, EEM still observes events but takes no actions.

EEM supports the following actions in action statements:

- Execute any CLI commands.
- Update a counter.
- Log an exception.
- Force the shut down of any module.
- Reload the device.
- Shut down specified modules because the power is over budget.
- Generate a syslog message.
- Generate a Call Home event.
- Generate an SNMP notification.
- Use the default action for the system policy.

Note If you want to allow the triggered event to process the default actions also, you must explicitly configure an EEM action with event-default or policy-default, based on the type of policy. For example, if you match a CLI command in a match statement, you must add the event-default action statement to the EEM policy. If the event-default action statement is not added, EEM will not allow the CLI command to execute.

Note Verify that your action statements within your user policy or overriding policy do not negate each other or adversely affect the associated system policy.

VSH Script Policies

You can also write policies in a VSH script, using a text editor. These policies have an event statement and action statement(s) just as other policies, and these policies can either augment or override system polices. After you write your script policy, copy it to the device and activate it. To configure a policy in a script, see the Defining a Policy Using a VSH Script, on page 173.

Environment Variables

You can define environment variables for EEM that are available for all policies. Environment variables are useful for configuring common values that you can use in multiple policies. For example, you can create an environment variable for the IP address of an external e-mail server.

You can use an environment variable in action statements by using the parameter substitution format.

Action Statement

The following example shows a sample action statement to force a module 1 shutdown, with a reset reason of “EEM action.”

```
switch (config-eem-policy)# action 1.0 forceshut module 1 reset-reason “EEM action”
```

Action Statement with Environment Variable

If you define an environment variable for the shutdown reason, called default-reason, you can replace that reset reason with the environment variable, as shown in following example.

```
switch (config-eem-policy)# action 1.0 forceshut module 1 reset-reason $default-reason
```

You can reuse this environment variable in any policy. For more information on environment variables, see the Defining an Environment Variable, on page 175.

EEM Event Correlation

Beginning with Cisco NX-OS Release 5.2, you can trigger an EEM policy based on a combination of events. First, you use the **tag** keyword to create and differentiate multiple events in the EEM policy. Then using a set of boolean operators (**and**, **or**, and **not**), along with the count and time, you can define a combination of these events to trigger a custom action.

High Availability

Cisco NX-OS supports stateless restarts for EEM. After a reboot or supervisor switchover, Cisco NX-OS applies the running configuration.

Licensing Requirements for EEM

The following table shows the licensing requirements for this feature:

Product	License Requirement
NX-OS	EEM requires no license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided at no extra charge to you.

Prerequisites for EEM

EEM has the following prerequisites:

- You must have network-admin user privileges to configure EEM.

Guidelines and Limitations

EEM has the following configuration guidelines and limitations:

- Action statements within your user policy or overriding policy should not negate each other or adversely affect the associated system policy.
- If you want to allow the triggered event to process the default actions also, you must explicitly configure an EEM action with event-default or policy-default, based on the type of policy. For example, if you match a CLI command in a match statement, you must add the event-default action statement to the EEM policy or EEM will not allow the CLI command to execute.
- An override policy that consists of an event statement and no action statement triggers no action and no notification of failures.
- An override policy without an event statement overrides all possible events in the system policy.
- When more than one event statement is included in an EEM policy, each event statement must have a **tag** keyword with a unique tag argument.

Default Settings

Table 23: Default EEM Parameters , on page 165 lists the default settings for EEM parameters.

Table 23: Default EEM Parameters

Parameters	Default
system policies	active

Configuring Embedded Event Manager

Defining a User Policy Using the CLI

You can define a user policy using the CLI.

To define a user policy using the CLI, follow these steps:

Procedure

Step 1 **configure terminal**

Enters configuration mode.

Step 2 **event manager applet** *applet-name*

Registers the applet with EEM and enters applet configuration mode. The *applet-name* can be any case-sensitive alphanumeric string up to 29 characters.

Step 3 **description** *policy-description*

(Optional) Configures a descriptive string for the policy. The string can be any alphanumeric string up to 80 characters. Enclose the string in quotation marks.

Step 4 **event** *event-statement*

Configures the event statement for the policy. See the Configuring Event Statements, on page 166.

Step 5 Do one of the following:

- **tag** *tagname1* {**and** | **andnot**} *tagname2* [{**and** | **andnot**} *tagname3* [{**and** | **andnot**} *tagname4*]] **happens** *occurs* **in** *seconds*

(Optional) Correlates multiple events in the policy.

The range for *occurs* is from 1 to 4294967295. The range for *seconds* is from 0 to 4294967295 seconds.

Step 6 **action** *action-statement*

Configures an action statement for the policy. See the Configuring Action Statements, on page 170.

Repeat Step 5 for multiple action statements.

Step 7 **show event manager policy internal** *name*

(Optional) Displays information about the configured policy.

Step 8 **copy running-config startup-config**

(Optional) Saves this configuration change.

Configuring Event Statements

To configure an event statement, use one the following commands in EEM configuration mode:

Command	Purpose
event cli [**tag** *tag_name* **match** *expression*] [**count** *repeats* |**time** *seconds*]	Triggers an event if you enter a CLI command that matches the regular expression. The **tag** *tag_name* keyword-argument pair identifies this specific event when multiple events are included in the policy. The *repeats* range is from 1 to 65000. The time range, in seconds, is from 0 to 4294967295, where 0 indicates no time limit.
event counter name *counter* **entry-val** *entry* **entry-op** {**eq**|**ge**|**gt**|**le**|**lt**|**ne** }[**exit-val** *exit* **exit-op** *exit* {**eq**|**ge**|**gt**|**le**|**lt**|**ne** }]	Triggers an event if the counter crosses the entry threshold (based on the entry operation—greater than, less than, and so on.) The event resets immediately. Optionally, you can configure the event to reset after the counter passes the exit threshold. The *counter* name can be any case-sensitive, alphanumeric string up to 28 characters. The *entry* and *exit* value ranges are from 0 to 2147483647.

Command	Purpose
event fanabsent [**fan** *number*] **time** *seconds*	Triggers an event if a fan is removed from the device for more than the configured time, in seconds. The fan number range is dependent on different switches (for example for 9513 switches the range is from1 to 2, for 9506/9509 switches the range is 1). The seconds range is from 10 to 64000.
event fanbad [**fan** *number*] **time** *seconds*	Triggers an event if a fan fails for more than the configured time, in seconds. The fan number range is dependent on different switches (for example for 9513 switches the range is from1 to 2, for 9506/9509 switches the range is 1). The seconds range is from 10 to 64000.
event memory {**critical** | **minor** |**severe**}	Triggers an event if a memory threshold is crossed.
event module-failure type *failure-type* **module** { *slot* | **all** { *slot* | **count** *repeats* [**time** *seconds*]	Triggers an event if a module experiences the failure type configured. The *slot* range is dependent on different switches (for example for 9513 switches the range is from1 to 13, for 9509 switches the range is 1 to 9). The *repeats* range is from 0 to 4294967295. The *seconds* range is from 0 to 4294967295.
event oir {**fan** | **module** | **powersupply**} {**anyoir** | **insert** | **remove** [*number*]}	Triggers an event if the configured device element (fan, module, or power supply) is inserted or removed from the device. You can optionally configure a specific fan, module, or power supply number. The *number* range is as follows: • Fan number is dependent on different switches. • Module number is dependent on different switches. • Power supply number range is from 1 to 2.
event policy-default count *repeats* [**time** *seconds*]	Uses the event configured in the system policy. Use this option for overriding policies. The *repeats* range is from 1 to 65000. The *seconds* range is from 0 to 4294967295.
event poweroverbudget	Triggers an event if the power budget exceeds the capacity of the configured power supplies.

Command	Purpose
event snmp oid *oid* **get-type**{**exact** | **next**} **entry-op** {**eq** | **ge** | **gt** | **le** | **lt** |**ne**} **entry-val** *entry* [**exit-comb** {**and** | **or**}] **exit-op** {**eq** | **ge** | **gt** | **le** | **lt** |**ne**} **exit-val** *exit***exit-time** *time***polling-interval** *interval*	Triggers an event if the SNMP OID crosses the entry threshold (based on the entry operation—greater than, less than, and so on.) The event resets immediately, or optionally you can configure the event to reset after the counter passes the exit threshold. The OID is in dotted decimal notation. The *entry* and *exit* value ranges are from 0 to 18446744073709551615. The time range is from 0 to 2147483647. The interval range is from 1 to 2147483647.

Command	Purpose
event syslog {**occurs occurs number** | **pattern syslog** *pattern* | **period** *time intervals* | **priority syslog** *priority* | **tag** *tag_name* }	Triggers an event based on a message logged in the syslog logfile. occurs occurs number-Specifies the number of occurrences. The range is from 1 to 65000. pattern syslog pattern-Specifies the syslog pattern. Normal regular expression pattern matching is used. The maximum size is 256 alphanumerical characters. period time intervals-Specifies the maximum time interval between messages. The range is from 0 to 4294967295 seconds. priority syslog priority-Specifies the syslog priority. • alerts—Specifies the alert log message • critical—Specifies the critical log message • debugging—Specifies the debugging message • emergencies—Specifies the emergency log message • errors—Specifies the error log message • informational—Specifies the informational log message • notification—Specifies the notification log message • pattern—Specifies the pattern matching • warnings—Specifies the warning message tag tag_name-Specifies the tag name. Maximum size is 29 alphanumerical characters. The tag tag_name keyword argument pair identifies this specific event when multiple events are included in the policy.
event temperature [**module** *slot*] [**sensor** *sensor number*]**threshold** {**any** | **major** | **minor**}	Triggers an event if the temperature sensor exceeds the configured threshold. The *slot* range is dependent on different switches. The sensor range is from 1 to 8 on MDS modules, but current MDS modules use the range from 1 to 3 only, some modules use the range from 1 to 2.

Configuring Action Statements

Use the following commands in EEM configuration mode to configure action statements:

Command	Purpose
action *number* **add** *variable-name*	Adds variable values to the **action** command when an EEM applet is triggered. To undo the add action, use the **no** form of this command.
action *number* **append** *variable-name*	Appends a variable value to an existing variable string when an EEM applet is triggered. To undo the append action, use the **no** form of this command.
action *number* **break**	Exits from a loop of action when an EEM applet is triggered. To disable the break action, use the **no** form of this command.
action *number* **cli command** *command-name*	Executes the configured VSH CLI commands when an EEM applet is triggered. To disable the CLI command action, use the **no** form of this command. Valid value for the VSH command name is 256 characters. From Cisco MDS NX-OS Release 8.1(1) the **command** keyword was added. The **command** keyword specifies the message to be sent to the Cisco NX-OS CLI. Add the command name within double quotation marks.
action *number* **cli local command** *command-name*	Executes the action commands in the same card on which an event is triggered. To disable the**action cli local command**, use the **no** form of this command. Valid value for the VSH command name is 256 characters. From Cisco MDS NX-OS Release 8.1(1), the **command** keyword was added. The **command** keyword specifies the message to be sent to the Cisco NX-OS CLI. Add the command name within double quotation marks.
action *number* **comment** *string*	Specifies an action of adding comments to an applet when an EEM applet is triggered. To disable the comment action, use the **no** form of this command. Valid value for string sequence is 256 characters.
action *number* **continue**	Specifies an action of continuing with loop of actions when an EEM applet is triggered. To disable the comment action, use the **no** form of this command.
action *number* [**.** *number*] **counter name** *counter* **value** *val* **op** {**dec** | **inc** | **nop** | **set**}	Modifies the counter by the configured value and operation. The action label is in the format number1.number2. *number* can be any number up to 16 digits. The range for *number2* is from 0 to 9. The counter name can be any case-sensitive, alphanumeric string up to 29 characters. The *val* can be an integer from 0 to 2147483647 or a substituted parameter.
action *number* **decrement** *decrement-name*	Specifies the action of decrementing the value of a variable, when an EEM applet is triggered. To remove the action from the applet, use the **no** form of this command.

Command	Purpose
action *number* **divide** *divide-name*	Divides the dividend value by the given divisor value when an EEM applet is triggered. To remove the calculation process, use the **no** form of this command.
action *number* **eem**	Specifies the EEM action command when an EEM applet is triggered. To remove the EEM action command, use the **no** form of this command.
action *number* **else**	Specifies the beginning of an else conditional action block in an if/else conditional action block when an EEM applet is triggered. To remove the else conditional action block, use the **no** form of this command.
action *number* **elseif**	Specifies the beginning of an elseif conditional action block in an if/else conditional action block when an EEM applet is triggered. To remove the else conditional action block, use the **no** form of this command.
action *number* **end**	Specifies the end of a conditional action block in the if/else and while conditional action block when an EEM applet is triggered. To remove the end conditional action block, use the **no** form of this command.
action *number* [**.** *number*] **event-default**	Executes the default action for the associated event. The action label is in the format number1.number2. *number* can be any number up to 16 digits. The range for *number2* is from 0 to 9.
action *number* **exit**	Exits from the running applet configuration when an EEM applet is triggered. To cancel the process of immediate exit from the running applet, use the **no** form of this command.
action *number* **file** {**close** | **delete** | **gets** | **open** | **puts** | **read** | **write**}	Configures the EEM applet file operations, use the action file command in applet configuration mode. To disable the configuration, use the **no** form of this command.
action *number* **foreach** *foreach-name*	Specifies the iteration of an input string using the delimiter as a tokenizing pattern. To remove iteration of the input string, use the **no** form of this command.
action *number* **if** *if-name*	Identifies the beginning of an if conditional block when an EEM applet is triggered, use the action if command in applet configuration mode. To remove the **if** conditional action block, use the **no** form of this command.
action *number* **increment** *increment-name*	Specifies the action of incrementing the value of a variable, when an EEM applet is triggered. To remove the action from the applet, use the **no** form of this command.
action *number* **multiply** *multiply-name*	Specifies the action of multiplying the variable value with a specified given integer value when an EEM applet is triggered. To remove the calculation process, use the **no** form of this command.
action *number* **overbudgetshut** [**module** *module-name*]	Forces one or more modules or the entire system to shut down because of a power overbudget issue.

Command	Purpose
action *number* **policy-default**	Executes the default action for the policy that you are overriding. To remove the action policy command from the configuration, use the **no** form of this command.
action *number* **publish-event**	Specifies the action of publishing an application-specific event when the event specified for an EEM applet is triggered. To remove the action of publishing an application-specific event, use the **no** form of this command.
action *number* **puts**	Enables the action of printing data directly to the local TTY when an EEM applet is triggered. To disable this function, use the **no** form of this command.
action *number* **regexp***regexp-name*	Matches a regular expression pattern on an input string when an EEM applet is triggered. To disable this function, use the **no** form of this command.
action *number* **reload**	Forces one or more modules or the entire system to reload.
action *number* **set***set-name*	Sets the value of a variable when an EEM applet is triggered. To remove the value of an EEM applet variable, use the **no** form of this command.
action *number* [**.** *number2*] **snmp-trap** {[**intdata1** *data* [**intdata2** *data* [**strdata string**]]]}	Sends an SNMP trap with the configured data. *number* can be any number up to 16 digits. The range for *number2* is from 0 to 9. The *data* arguments can by any number up to 80 digits. The *string* can be any alphanumeric string up to 80 characters.
action *number* **string**	Specifies the string action command for an EEM applet. To remove the action of string operation, use the **no** form of this command.
action *number* **wait** *wait-value*	Specifies the wait time for an action for an EEM applet. To disable this function, use the **no** form of this command.
action *number* **while** *while-number*	Identifies the beginning of a loop of a conditional block when an EEM applet is triggered. To disable this function, use the **no** form of this command.
action *number* [**.** *number2*] **exceptionlog module** *module* **syserr** *error* **devid** *id* **errtype** *type* **errcode** *code* **phylayer** *layer* **ports** *list* **harderror** *error* [**desc** *string*]	Logs an exception if the specific conditions are encountered when an EEM applet is triggered.
action *number* [**.** *number number2*] **forceshut** [**module** *slot* | **xbar** *xbar number*] **reset-reason** *seconds*	Forces a module, crossbar, or the entire system to shut down. The action label is in the format number1.number2. *number* can be any number up to 16 digits. The range for *number2* is from 0 to 9. The *slot* range is dependent on different switches. The *xbar-number* range is from 1 to 2 and is only available on MDS 9513 modules. The reset reason is a quoted alphanumeric string up to 80 characters.

Command	Purpose
action *number* [**.** *number*] **overbudgetshut** [**module** *slot* [**-** *slot*]]	Forces one or more modules or the entire system to shut down because of a power overbudget issue. *number* can be any number up to 16 digits. The range for *number2* is from 0 to 9. The *slot* range is dependent on different switches.
action *number* [**.** *number*] **policy-default**	Executes the default action for the policy that you are overriding. The action label is in the format number1.number2. *number* can be any number up to 16 digits. The range for *number2* is from 0 to 9.
action *number* [**.** *number*] **reload** [**module** *slot* [**-** *slot*]]	Forces one or more modules or the entire system to reload. *number* can be any number up to 16 digits. The range for *number2* is from 0 to 9. The *slot* range is dependent on different switches.
action *number* [**.** *number2*] **syslog** [**priority** *prio-val*] **msg** *error message*	Sends a customized syslog message at the configured priority.*number* can be any number up to 16 digits. The range for *number2* is from 0 to 9. The *error-message* can be any quoted alphanumeric string up to 256 characters.

Note If you want to allow the triggered event to process the default actions also, you must explicitly configure an EEM action with event-default or policy-default, based on the type of policy. For example, if you match a CLI command in a match statement, you must add the event-default action statement to the EEM policy or EEM will not allow the CLI command to execute. You can bypass all CLI-based EEM policies using the **terminal event-manager bypass** command. To revert use the **terminal no event-manager bypass** command.

Defining a Policy Using a VSH Script

To define a policy using a VSH script, follow these steps:

Procedure

Step 1 In a text editor, list the CLI commands that define the policy.

Step 2 Name the text file and save it.

Step 3 Copy the file to the following system directory:

bootflash://eem/user_script_policies

Registering and Activating a VSH Script Policy

To register and activate a policy defined in a VSH script, follow these steps:

Procedure

Step 1 **configure terminal**

Enters configuration mode.

Step 2 **event manager policy** *policy-script*

Registers and activates an EEM script policy. The *policy-script* can be any case-sensitive alphanumeric string up to 29 characters.

Step 3 **show event manager internal policy** *name*

(Optional) Displays information about the configured policy.

Step 4 **copy running-config startup-config**

(Optional) Saves this configuration change.

Overriding a Policy

To override a system policy, follow these steps:

Procedure

Step 1 **configure terminal**

Enters configuration mode.

Step 2 **show event manager policy-state** *system-policy*

(Optional) Displays information about the system policy that you want to override, including thresholds. Use the **show event manager system-policy** command to find the system policy names.

Step 3 [**no**] **event manager applet** *applet-name* **override** *system-policy*

Overrides a system policy and enters applet configuration mode. The *applet-name* can be any case-sensitive alphanumeric string up to 29 characters. The *system-policy* must be one of the existing system policies.

Step 4 **description** *policy-description*

(Optional) Configures a descriptive string for the policy. The string can be any alphanumeric string up to 80 characters. Enclose the string in quotation marks.

Step 5 [**no**] **event** *event-statement*

Configures the event statement for the policy. See the Configuring Event Statements, on page 166. Using the **no** keyword deletes the overridden event, if any.

- Deleting an overridden policy does not remove the default system policy.
- You can modify an overridden policy by changing the respective Zone, FCNS, or FLOGI limit values.

Step 6 **action** *action-statement*

Configures an action statement for the policy. See the Configuring Action Statements, on page 170.

Repeat Step 6 for multiple action statements.

- Zone, FLOGI, and FCNS support only syslog message generation as the action.
- If an action is not configured, the default action associated with the default system policy is executed. If an action is configured, both the configured and default actions are executed. This functionality is applicable only to Zone, FLOGI, and FCNS system policies.

Step 7 **show event manager policy-state** *name*

(Optional) Displays information about the configured policy.

Step 8 **copy running-config startup-config**

(Optional) Saves this configuration change.

Note Multiple overrides for Zone, FLOGI, and FCNS EEM policies are not allowed.

Defining an Environment Variable

To define a variable to serve as a parameter in an EEM policy, follow these steps:

Procedure

Step 1 **configure terminal**

Enters configuration mode.

Step 2 **event manager environment** *variable-name variable-value*

Creates an environment variable for EEM. The *variable-name* can be any case-sensitive alphanumeric string up to 29 characters. The *variable-value* can be any quoted alphanumeric string up to 39 characters.

Step 3 **show event manager environment**

(Optional) Displays information about the configured environment variables.

Step 4 **copy running-config startup-config**

(Optional) Saves this configuration change.

Verifying the EEM Configuration

To display EEM configuration information, perform one of the following tasks:

Command	Purpose
show event manager environment [**variable-name** | **all**]	Displays information about the event manager environment variables.
show event manager event-types [*event* | **all** | **module** *slot*]	Displays information about the event manager event types.
show event manager history events [detail] [maximum *num-events*] **[severity {catastrophic | minor | moderate | severe}]**	Displays the history of events for all policies.
show event manager policy internal [*policy-name*] [**inactive**]	Displays information about the configured policies.
show event manager policy-state *policy-name*	Displays information about policy state, including thresholds.
show event manager script system [*policy-name*}**all**]	Displays information about the script policies.
show event manager system-policy [**all**]	Displays information about the predefined system policies.
show running-config eem	Displays information about the running configuration for EEM.
show startup-config eem	Displays information about the startup configuration for EEM.

Configuration Examples for EEM

This example overrides the __lcm_module_failure system policy by changing the threshold for just module 3 hitless upgrade failures. The following example also sends a syslog message. The settings in the system policy, __lcm_module_failure, apply in all other cases.

```
event manager applet example2 override __lcm_module_failure
 event module-failure type hitless-upgrade-failure module 3 count 2
 action 1 syslog priority errors msg module 3 “upgrade is not a hitless upgrade!”
 action 2 policy-default
```

The following example modifies an overridden policy by changing the number of FCNS database entries to 1500. It also generates both the configured and the default syslog messages of the default system policy

```
event manager applet fcns_policy override __fcns_entries_max_per_switch
event fcns entries max-per-switch 1500
```

```
      action 1.0 syslog priority warnings msg  FCNS DB entries have reached the EEM limit
```

The following example deletes the event of an overridden policy:

```
no event manager applet zone_policy
```

The following example creates an EEM policy that allows the CLI command to execute but triggers an SNMP notification when a user enters configuration mode on the device:

```
event manager applet TEST
event cli match "conf t"
action 1.0 snmp-trap strdata "Configuration change"
action 2.0 event-default
```

Note You must add the event-default action statement to the EEM policy or EEM will not allow the CLI command to execute.

The following example shows how to confiure a VSH command string to be executed when an EEM applet is triggered:

```
switch# configure terminal
switch(config)# event manager applet cli-applet
switch(config-applet)# action 1.0 cli command “show interface e 3/1‟
```

Additional References

For additional information related to implementing EEM, see the following section:

MIBs

MIBs	MIBs Link
- CISCO-EMBEDDED-EVENT-MGR-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

CHAPTER 9

Configuring RMON

RMON is an Internet Engineering Task Force (IETF) standard monitoring specification that allows various network agents and console systems to exchange network monitoring data. You can use the RMON alarms and events to monitor Cisco MDS 9000 Family switches running the Cisco SAN-OS Release 2.0(1b) or later or Cisco NX-OS Release 4.1(3) or later software.

Information About RMON

RMON is disabled by default, and no events or alarms are configured in the switch.

All switches in the Cisco MDS 9000 Family support the following RMON functions (defined in RFC 2819):

- Alarm—Each alarm monitors a specific management information base (MIB) object for a specified interval. When the MIB object value exceeds a specified value (rising threshold), the alarm condition is set and only one event is triggered regardless of how long the condition exists. When the MIB object value falls below a certain value (falling threshold), the alarm condition is cleared. This allows the alarm to trigger again when the rising threshold is crossed again.
- Event—Determines the action to take when an event is triggered by an alarm. The action can be to generate a log entry, an SNMP trap, or both.

For agent and management information, see the *Cisco MDS 9000 Family MIB Quick Reference.*

For information on an SNMP-compatible network management station, see the System Management Configuration Guide, Cisco DCNM for SAN.

For SNMP security-related CLI configurations, see .

RMON Configuration Information

RMON is disabled by default and no events or alarms are configured in the switch. You can configure your RMON alarms and events by using the CLI or an SNMP-compatible network management station.

Tip We recommend an additional, generic RMON console application on the network management station (NMS) to take advantage of RMON's network management capabilities. Refer to the System Management Configuration Guide, Cisco DCNM for SAN.

RMON Configuration Using Threshold Manager

RMON is disabled by default and no events or alarms are configured in the switch. You can configure your RMON alarms and events by using the CLI or by using Threshold Manager in Device Manager.

The Threshold Monitor allows you to trigger an SNMP event or log a message when the selected statistic goes over a configured threshold value. RMON calls this a rising alarm threshold. The configurable settings are as follows:

- Variable—The statistic you want to set the threshold value on.
- Value—The value of the variable that you want the alarm to trigger at. This value is the difference (delta) between two consecutive polls of the variable by Device Manager.
- Sample—The sample period (in seconds) between two consecutive polls of the variable. Select your sample period such that the variable does not cross the threshold value you set under normal operating conditions.
- Warning—The warning level used by Device Manager to indicate the severity of the triggered alarm. This is a DCNM-SAN and Device Manager enhancement to RMON.

Note To configure any type of RMON alarm (absolute or delta, rising or falling threshold) click More on the Threshold Manager dialog box. You should be familiar with how RMON defines these concepts before configuring these advanced alarm types. Refer to the RMON-MIB (RFC 2819) for information on how to configure RMON alarms.

Note You must also configure SNMP on the switch to access RMON MIB objects.

RMON Alarm Configuration Information

Threshold Manager provides a list of common MIB objects to set an RMON threshold and alarm on. The alarm feature monitors a specific MIB object for a specified interval, triggers an alarm at a specified value (rising threshold), and resets the alarm at another value (falling threshold).

You can also set an alarm on any MIB object. The specified MIB must be an existing SNMP MIB object in standard dot notation (1.3.6.1.2.1.2.2.1.14.16777216 16 16777216 for ifInOctets.16777216161616777216).

Use one of the following options to specify the interval to monitor the MIB variable (ranges from 1 to 4294967295 seconds):

- Use the **delta** option to test the change between samples of a MIB variable.
- Use the **absolute** option to test each MIB variable directly.
- Use the **delta** option to test any MIB objects that are counters.

The range for the **rising threshold** and **falling threshold** values is -2147483647 to 2147483647.

⚠

Caution The **falling threshold** must be less than the **rising threshold**.

You can optionally specify the following parameters:

- The event-number to trigger if the rising or falling threshold exceeds the specified limit.
- The owner of the alarm.

Default Settings

Table 24: Default RMON Settings , on page 181 lists the default settings for all RMON features in any switch.

Table 24: Default RMON Settings

Parameters	**Default**
RMON alarms	Disabled
RMON events	Disabled

Configuring RMON

RMON is disabled by default, and no events or alarms are configured in the switch.

Configuring the RMON Traps in SNMP

To enable RMON traps in the SNMP configuration, follow these steps:

Before you begin

You must enable the RMON traps in the SNMP configuration for the RMON configuration to function correctly.

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server enable traps rmon**

Enables the RMON trap types.

Note You must also configure SNMP on the switch to access RMON MIB objects.

Configuring the RMON Alarm

To enable RMON alarms, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **rmon alarm 20 1.3.6.1.2.1.2.2.1.14.16777216 2900 delta rising-threshold 15 1 falling-threshold 0 owner test**

Configures RMON alarm number 20 to monitor the 1.3.6.1.2.1.2.2.1.14.16777216 once every 900 seconds until the alarm is disabled and checks the change in the variables rise or fall. If the value shows a MIB counter increase of 15 or more, the software triggers an alarm. The alarm in turn triggers event number 1, which is configured with the RMON event command. Possible events can include a log entry or an SNMP trap. If the MIB value changes by 0, the alarm is reset and can be triggered again.

Note You can also configure the following rmon events:

- Event 1: Fatal
- Event 3: Error
- Event 4: Warning
- Event 5: Information

Step 3 switch(config)# **no rmon alarm 2**

Deletes the specified entry from the alarm table.

Configuring the RMON Event

To enable RMON events, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **rmon event 2 log trap eventtrap description CriticalErrors owner Test2**

Creates RMON event number 2 to define CriticalErrors and generates a log entry when the event is triggered by the alarm. The user Test2 owns the row that is created in the event table by this command. This example also generates an SNMP trap when the event is triggered.

Note You can also configure the following rmon events:

- Event 1: Fatal
- Event 3: Error
- Event 4: Warning
- Event 5: Information

Step 3 switch(config)# **no rmon event 5**

Deletes an entry from the RMON event table.

Verifying the RMON Configuration

To display the RMON configuration information, perform one of the following tasks:

Command	Purpose
show rmon alarms	Displays Configured RMON Alarms
show rmon hcalarms	Displays Configured RMON High Capacity Alarms
show rmon events	Displays Configured RMON Events

For detailed information about the fields in the output from these commands, refer to the Cisco MDS 9000 NX-OS Command Reference.

Use the **show rmon** and **show snmp** commands to display configured RMON and SNMP information (see Configured RMON Alarms, on page 183 and Configured RMON Events, on page 184).

Configured RMON Alarms

The following example displays configured RMON alarms:

```
switch# show rmon alarms
Alarm 1 is active, owned by admin
 Monitors 1.3.6.1.2.1.2.2.1.16.16777216 every 1 second(s)
 Taking delta samples, last value was 0
 Rising threshold is 1, assigned to event 0
 Falling threshold is 0, assigned to event 0
 On startup enable rising or falling alarm
```

Confirmed RMON High Capacity Alarms

The following example displays confirmed RMON high capacity alarms:

```
switch# show rmon hcalarms
High Capacity Alarm 10 is active, owned by Testuser
 Monitors 1.3.6.1.2.1.31.1.1.1.6.16785408 every 300 second(s)
 Taking absolute samples, last value was 0 (valuePositive)
 Rising threshold low is 4294967295 & high is 15 (valuePositive)
 Rising threshold assigned to event 1
 Falling threshold low is 0 & high is 0 (valueNotAvailable)
 Falling threshold assigned to event 0
 On startup enable rising alarm
 Number of Failed Attempts is 0
```

Note High capacity RMON alarms can be configured using the CISCO-HC-ALARM-MIB. See the *Cisco MDS 9000 Series MIB Quick Reference*.

Configured RMON Events

The following example displays configured RMON events:

```
switch# show rmon events
Event 2 is active, owned by Test2
 Description is CriticalErrors
 Event firing causes log and trap to community eventtrap, last fired 0
Event 500 is active, owned by admin
 Description is
 Event firing causes log, last fired 138807208
```

Additional References

For additional information related to implementing RMON, see the following section:

MIBs

MIBs	MIBs Link
• CISCO-RMON-CAPABILITY.my • CISCO-RMON-CONFIG-CAPABILITY.my • CISCO-RMON-CONFIG-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

Feature History for RMON

The following table lists the release history for this feature. Only features that were introduced or modified in Release 3.x or a later release appear in the table.

Table 25: Feature History for RMON

Feature Name	Releases	Feature Information
RMON high capacity alarms	3.0(1)	Provides the show rmon high capacity alarms command to display RMON high capacity alarm values.

CHAPTER 10

Configuring Online Diagnostics

Beginning with Cisco MDS NX-OS Release 6.2, the Cisco MDS 9700 Series supports the GOLD (Generic Online Diagnostics) feature. GOLD is a diagnostic service which is also supported on the Cisco Nexus 7000 and 7700 series switches. This chapter describes how to configure the GOLD feature on a Cisco MDS 9700 Series switch.

Information About Online Diagnostics

Online diagnostics verifies the hardware and data paths and identifies faulty devices.

Online Diagnostic Overview

The GOLD (Generic Online Diagnostics) framework tests and verifies the hardware devices and data path in a live system.

The GOLD tests can be executed in three modes:

- Bootup
- Health-monitoring (also called Runtime)
- On-demand

The following explains the diagnostics test suite attributes:

- B/C/* - Bypass bootup level test / Complete bootup level test / NA
- P/* - Per port test / NA
- M/S/* - Only applicable to active / standby unit / NA
- D/N/* - Disruptive test / Non-disruptive test / NA
- H/O/* - Always enabled monitoring test / Conditionally enabled test / NA

- F/* - Fixed monitoring interval test / NA
- X/* - Not a health monitoring test / NA
- E/* - Sup to line card test / NA
- L/* - Exclusively run this test / NA
- T/* - Not an ondemand test / NA
- A/I/* - Monitoring is active / Monitoring is / NA

Bootup Diagnostics

Bootup diagnostics run during bootup and detect faulty hardware before a Cisco MDS 9700 series switch brings a module online. For example, if there is a faulty module in the device, the appropriate bootup diagnostics test fails to indicate the fault.

Note The bootup diagnostics tests are triggered during bootup.

Table 26: Bootup Diagnostics , on page 186 describes the bootup diagnostic tests for a module and a supervisor.

Table 26: Bootup Diagnostics

Diagnostic	Attributes	Description
Linecard		
EOBCPortLoopback	C**D**X**T*	Verifies the health of EOBC (Ethernet Out-of-Band Connectivity) interface.
OBFL	C**N**X**T*	Verifies the integrity of the OBFL (Onboard Failure Logging) flash.
BootupPortLoopback	CP*N**XE*T*	PortLoopback test that runs only during module bootup.
Supervisor		
USB	C**N**X**T*	Verifies the USB controller initialization on a module.
ManagementPortLoopback	C**D**X**T*	Verifies the health of management interface of a module.
EOBCPortLoopback	C**D**X**T*	Verifies the health of EOBC (Ethernet Out-of-Band Connectivity) interface.
OBFL	C**N**X**T*	Verifies the integrity of the OBFL (Onboard Failure Logging) flash.

When the **show module** command is executed, the result of bootup diagnostics is displayed as Online Diag Status. The result of individual test is displayed when the **show diagnostic result** command is executed for appropriate module and test ID or test name.

The Cisco MDS 9700 Family switch can be configured to either bypass the bootup diagnostics or run the complete set of bootup diagnostics. See the Setting the Bootup Diagnostic Level, on page 191.

Health Monitoring Diagnostics

Health Monitoring (HM) diagnostics is enabled by default to verify the health of a live system at periodic intervals. The monitoring interval (within an allowed range) can be configured by the user, which is different for each test. See the Activating a Health Monitoring Diagnostic Test, on page 192 for more information. The diagnostic tests detect hardware errors and data path issues.

Health Monitoring diagnostics are non-disruptive (does not disrupt the data or control traffic). The Health Monitoring tests can be disabled by the user. See the Deactivating a Health Monitoring Diagnostic Test, on page 193 for more information.

The following table describes the health monitoring diagnostics for a supervisor.

Diagnostic	Default Testing Interval	Attributes	Description
Supervisor			
ASICRegisterCheck	20 seconds	***N******A	Verifies read or write access to scratch registers for the ASICs on the supervisor.
NVRAM	5 minutes	***N******A	Verifies the sanity of the NVRAM blocks on a supervisor.
RealTimeClock	5 minutes	***N******A	Verifies that the real-time clock on the supervisor is ticking.
PrimaryBootROM	30 minutes	***N******A	Verifies the integrity of the primary boot device on the supervisor.
SecondaryBootROM	30 minutes	***N******A	Verifies the integrity of the secondary boot device on the supervisor.
CompactFlash	30 minutes	***N******A	Verifies access to the compact flash devices.
ExternalCompactFlash	30 minutes	***N******A	Verifies access to the external compact flash devices.
PwrMgmtBus	30 seconds	**MN******A	Verifies the standby power management control bus.
SystemMgmtBus	30 seconds	**MN******A	Verifies the availability of the standby system management bus.
StatusBus	30 seconds	**MN******A	Verifies the status transmitted by the status bus for the supervisor, modules, and fabric cards.
StandbyFabricLoopback	30 seconds	**SN******A	Verifies the connectivity of the standby supervisor to the fabric modules.

Table 27: Health Monitoring Diagnostics, on page 188 describes the health monitoring diagnostics for the .

Table 27: Health Monitoring Diagnostics

Diagnostic	Default Testing Interval	Attributes	Description
Linecard			
ASICRegisterCheck	1 minute	***N******A	Verifies read or write access to scratch registers for the ASICs on a module.
PrimaryBootROM	30 minutes	***N******A	Verifies the integrity of the primary boot device on a module.
SecondaryBootROM	30 minutes	***N******A	Verifies the integrity of the secondary boot device on a module.
SnakeLoopback	20 minutes	*P*N***E**	Verifies connectivity from sup to all the ports in the Linecard. It checks the integrity of the data path up to the MAC component in a progressive manner (a single run of tests covers all the ports). It runs on all the ports irrespective of their states. This is a non-disruptive test.
IntPortLoopback	5 minutes	*P*N***E***	Verifies connectivity from sup to all the ports in the Linecard (one port at a time). It checks the integrity of the data path up to the MAC component. This test runs in Health Monitoring (HM) mode as well as it can be triggered in "on-demand mode." This test is Non-disruptive. **Note** The IntPortLoopback test is supported from Cisco MDS NX-OS Release 6.2(7).
RewriteEngine Loopback	1 minute	*P*N***E**A	Verifies the integrity of each link on the fabric module from sup to linecard.

On-Demand Diagnostics

All the Health Monitoring tests can be evoked on demand also. On-demand diagnostics runs only when invoked by the user.

Note The data paths (PHY and SFP) which are not verified by other Health Monitoring tests can be verified by the PortLoopback and ExtPortLoopback tests.

You can run on-demand diagnostics whenever required. See the Starting or Stopping an On-Demand Diagnostic Test, on page 194 for more information.

On , both the PortLoopback and ExtPortLoopback tests are available in on-demand mode only as they are disruptive.

Table 28: On-Demand Diagnostics, on page 189 describes the on-demand diagnostics (for module only) on the

Table 28: On-Demand Diagnostics

Diagnostic	Attributes	Description
Linecard		
PortLoopback	*P*D**XE***	Verifies connectivity from sup to all the ports in the module. It checks the integrity of the data path up to PHY. This test is available only in “on-demand mode.” The test runs on all the ports irrespective of the port state. **Note** Portloopback test is equivalent to the Serdes Loopback test of OHMS.
ExtPortLoopback	*P*D**XE***	Identifies hardware errors in the entire data path up to PHY including the SFP. **Note** Connect a loopback plug to loop the Tx of the port to the Rx of the port before running the test. If the loopback plug is not connected this test fails. **Note** The ExtPortLoopback test is supported from Cisco MDS NX-OS Release 6.2(11c).

Caution The PortLoopback and ExtPortLoopback tests are disruptive as they bring down the port for the purpose of diagnostic operation.

Recovery Actions on Specified Health Monitoring Diagnostics

When the Health Monitoring Diagnostic test fails consecutively for a threshold number of up to 10 times, it takes default action through EEM, which includes generating alerts (callhome, syslog) and logging (OBFL, exception logs), and the diagnostic test gets disabled on the failed instance (port or fabric or device).

These actions are informative, but they do not remove faulty devices from the live system, which can lead to network disruption, traffic black holing, and so forth.

Note Restart the Health Monitoring tests on failed instances by clearing the test result, deactivating, and then activate the test on the same module. For more information see Clearing Diagnostic Results, on page 196, Deactivating a Health Monitoring Diagnostic Test, on page 193, and Activating a Health Monitoring Diagnostic Test, on page 192.

Beginning with the Cisco MDS NX-OS Release 6.2(11), the system can be configured to take corrective (recovery) actions in addition to the default actions after reaching the threshold number of consecutive failures for any of the following Health Monitoring tests:

- PortLoopback test (supported only on Cisco MDS 48-Port 10-Gbps FCoE Module)

- RewriteEngineLoopback test
- StandbyFabricLoopback test
- Internal PortLoopback test

Note The corrective (recovery) actions are disabled by default.

Corrective (recovery) Action for Supervisor

The corrective action for sup is as follows:

StandbyFabricLoopback test—The system reloads the standby supervisor and after three retries, the system powers off the standby supervisor.

Note After reload, when the standby supervisor comes online, the Health Monitoring Diagnostics starts by default.

Note One retry means a complete cycle of reloading the standby supervisor followed by threshold number of consecutive failures of StandbyFabricLoopback test.

High Availability

A key part of high availability is detecting hardware failures and taking corrective action in a live system. GOLD contributes to the high availability of the system by detecting hardware failures and providing feedback to software components to make switchover decisions.

Cisco MDS 9700 Family switches support stateless restart for GOLD by applying the running configuration after a reboot. After supervisor switchover, GOLD resumes diagnostics from the new active supervisor.

Licensing Requirements for Online Diagnostics

Product	License Requirement
Cisco NX-OS	Online diagnostics require no license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided at no extra charge to you. For a complete explanation of the Cisco NX-OS licensing scheme, see the Cisco MDS 9000 Family NX-OS Licensing Guide.

Default Settings

Table 29: Default Online Diagnostic Parameters , on page 191 lists the default settings for online diagnostic parameters.

Table 29: Default Online Diagnostic Parameters

Parameters	Default
Bootup diagnostics level	complete
Health Monitoring tests	active
Corrective (Recovery) actions	disabled

Configuring Online Diagnostics

Setting the Bootup Diagnostic Level

To configure the bootup diagnostics to run the complete set of tests, or to bypass all bootup diagnostic tests for a faster module bootup time, perform these tasks:

Note It is recommended to set the bootup online diagnostics level to **complete**.

Procedure

Step 1 **configure terminal**

Example:

```
switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)#
```

Places in the global configuration mode.

Step 2 **diagnostic bootup level** {**complete** | **bypass** }

Example:

```
switch(config)# diagnostic bootup level complete
```

Configures the bootup diagnostic level to trigger diagnostics when the device boots:

- **complete**—Performs all bootup diagnostics. The default is complete.
- **bypass**—Does not perform any bootup diagnostics.

Step 3 **show diagnostic bootup level**

Example:

```
switch(config)# show diagnostic bootup level
```

(Optional) Displays the bootup diagnostic level (bypass or complete) that is currently in place on the device.

Step 4 **copy running-config startup-config**

Example:

```
switch(config)# copy running-config startup-config
```

(Optional) Copies the running configuration to the startup configuration.

Displaying the List of Available Tests

Procedure

show diagnostic content module *slot*

Example:

```
switch# show diagnostic content module 1
```

(Optional) Displays the list of information about the diagnostics and their attributes on a given module.

slot—The module number on which the test is activated.

Activating a Health Monitoring Diagnostic Test

Procedure

Step 1 **configure terminal**

Example:

```
switch# configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)#
```

Enters global configuration mode.

Step 2 **diagnostic monitor interval module** *slot* **test** [*test-id* | *name* | **all**] **hour** *hour* **min** *minutes* **second** *sec*

Example:

```
switch(config)# diagnostic monitor interval module 6 test 3 hour 1 min 0 sec 0
```

(Optional) Configures the interval at which the specified test is run. If no interval is set, the test runs at the interval set previously, or the default interval.

The arguments are as follows:

- slot—The module number on which the test is activated.
- test-id—Unique identification number for the test.
- name—Predefined name of the test.
- hour —The range is from 0 to 23 hours.
- minutes—The range is from 0 to 59 minutes.
- seconds—The range is from 0 to 59 seconds.

Step 3 **diagnostic monitor module** *slot* **test** [*test-id* | *name* | **all**]

Example:

```
switch(config)# diagnostic monitor module 6 test 3
switch(config)# diagnostic monitor module 6 test SecondaryBootROM
```

Activates the specified test.

The arguments are as follows:

- slot—The module number on which the test is activated.
- test-id—Unique identification number for the test.
- name—Predefined name of the test.

Step 4 **show diagnostic content module** {*slot* | **all**}

Example:

```
switch(config)# show diagnostic content module 6
```

(Optional) Displays information about the diagnostics and their attributes.

The argument is as follows:

- slot—The module number on which the test is activated.

Deactivating a Health Monitoring Diagnostic Test

Note Inactive tests keep their current configuration but do not run at the scheduled interval.

To deactivate a test, perform this task:

Command	Purpose

no diagnostic monitor module *slot* **test** [*test-id* | *name* | **all**]
Examples:
`switch(config)# no diagnostic monitor interval module 8 test 3`
`switch(config)# no diagnostic monitor interval module 8 test SecondaryBootROM`
	Deactivates the specified test.
	The arguments are as follows:
	• slot—The module number on which the test is activated.
	• test-id—Unique identification number for the test.
	• name—Predefined name of the test.

Starting or Stopping an On-Demand Diagnostic Test

On-demand diagnostic test can be started or stopped, with actions (optional) to modify the number of iterations to repeat the test and determine the action to be taken on test failure.

Note It is recommended to manually start a disruptive diagnostic test during a scheduled network maintenance time.

To start or stop an on-demand diagnostic test, perform these tasks:

Procedure

Step 1 **diagnostic ondemand iteration** *number*

Example:

```
switch# diagnostic ondemand iteration 5
```

(Optional) Configures the number of times that the on-demand test runs. The range is from 1 to 999. The default is 1.

Step 2 **diagnostic ondemand action-on-failure** {**continue failure-count** *num-fails* | **stop**}

Example:

```
switch# diagnostic ondemand action-on-failure stop
```

(Optional) Configures the action to take if the on-demand test fails.

Step 3 **show diagnostic ondemand setting**

Example:

```
switch# show diagnostic ondemand setting
Test iterations = 1
Action on test failure = continue until test failure limit reaches 1
```

(Optional) Displays information about on-demand diagnostics.

Step 4 **diagnostic start module** *slot* **test** [*test-id* | *name* | **all** | **non-disruptive**][**port** *port-number* | **all**]

Example:

```
switch# diagnostic start module 6 test all
```

Starts one or more diagnostic tests on a module.

The arguments are as follows:

- all— All the tests are triggered.

Note The multiple test- id or name can be specified separated by commas.

- non-disruptive—All the non-disruptive tests are triggered.
- port— The tests can be invoked on a single port or range of ports or all ports.

Step 5 **diagnostic run module** *slot* **test** {**PortLoopback** | **RewriteEngineLoopback** | **SnakeLoopback** | **IntPortLoopback** | **ExtPortLoopback**} {**port** *port-id* }

Example:

```
switch# diagnostic run module 3 test PortLoopback port 1
```

Starts the selected test on a module and displays the result on the completion of the test.

Note This command is introduced from the Cisco MDS NX-OS Release 6.2(11c).

For more information, see Starting an On-Demand Diagnostic Test in On-demand Mode, on page 196.

Step 6 **diagnostic stop module** *slot* **test** [*test-id* | *name* | **all**]

Example:

```
switch# diagnostic stop module 6 test all
```

(Optional) Stops one or more diagnostic tests on a module.

Step 7 **show diagnostic status module** *slot*

Example:

```
switch# show diagnostic status module 6
```

(Optional) Displays all the tests which are running and queued up with information about the testing mode for that module.

When the tests are not running or enqueued on the given module, the status is displayed as NA.

Step 8 **show diagnostic result module** *slot* **test** [*test-id* | *name*]

Example:

```
switch# show diagnostic result module 1 test 3 SecondaryBootROM
```

(Optional) Displays the result of the specified test.

Starting an On-Demand Diagnostic Test in On-demand Mode

OHMS (Online Health Management System) supports invoking tests in an “on-demand mode” which displays the results immediately after running the test.

From the Cisco MDS NX-OS Release 6.2(11c), GOLD supports invoking a specific test from a set of tests in “on-demand mode” and displaying the test results immediately after running the test.

GOLD tests can be invoked in an 'on-demand' mode using the **diagnostic start module** command. The **diagnostic run module** command also supports the same action but there are a few key differences between the two. The following are the differences between the two commands:

- In contrast to the **diagnostic start module** command, the **diagnostic run module** command blocks the current CLI session till the completion of test. After the completion of the test the CLI session is unblocked, and the result is displayed on the same console.

Note The CLI session will be blocked till the completion of test or for a maximum of 15 seconds. If the test is not completed within the time frame of 15 seconds, then GOLD unblocks the CLI session and allows the test to run in the background till completion.

Note Only one test can be invoked on a particular module using the **diagnostic run module** command. If the user attempts to invoke another test on the same module, it displays an error and the test is not invoked.

- The **diagnostic start module** command requires the user to execute the **show diagnostic result** command in order to display the test result. As the test runs in the background (the current CLI session is not blocked), the user needs to issue **show diagnostic result** command to view the result, whereas the test result is implicitly displayed on the same console when the **diagnostic run module** command is executed.
- The results displayed through the diagnostic run command are more intuitive than those from the show diagnostic results command.

Note The maximum number of ports recommended for the **diagnostic run module** command is 5.

Clearing Diagnostic Results

To clear the diagnostic test results, use the following command:

Command	Purpose

diagnostic clear result module [*slot* |**all**] **test** {*test-id* | **all** }
Example:
`switch# diagnostic clear result module 2 test all`
`switch# diagnostic clear result module 2 test 3`
	Clears the test result for the specified test.

Simulating Diagnostic Results

To test the behavior of GOLD in case of a diagnostic test failure, GOLD provides a mechanism to simulate the test failure on a port, sup, or fabric.

Note Simulating a failure after enabling corrective actions will result in triggering an action (see Corrective action) on the component where the failure was simulated.

To simulate a diagnostic test result, use the following command:

Command	Purpose
diagnostic test simulation module *slot* **test** *test-id* {**fail** | **random-fail** | **success** } [**port** *number* | **all**] **Example:** `switch# diagnostic test simulation module 2 test 2 fail`	Simulates a test result.

To clear the simulated diagnostic test result, use the following command:

Command	Purpose
diagnostic test simulation module *slot* **test** *test-id* **clear** **Example:** `switch# diagnostic test simulation module 2 test 2 clear`	Clears the simulated test result.

Enabling Corrective (Recovery) Actions

To enable corrective (recovery) actions, use the following command:

Procedure

Step 1 **configure terminal**

Enters global configuration mode.

Step 2 **diagnostic eem action conservative**

Example:

```
switch(config)# diagnostic eem action conservative
```

Enables corrective or recovery actions.

Note This command is applicable to the system as a whole and cannot be specifically configured to any particular module or test.

Step 3 **no diagnostic eem action conservative**

Disables corrective (recovery) actions.

Verifying the Online Diagnostics

To display GOLD test results, status, and configuration information use one of these commands:

Command	Purpose
show diagnostic bootup level	Displays information about bootup diagnostics.
show diagnostic content module {*slot* | **all**}	Displays information about diagnostic test content for a module.
show diagnostic description module *slot* **test** [*test-name* | **all**]	Displays the diagnostic description.
show diagnostic events [**error** | **info**]	Displays diagnostic events by error and information event type.
show diagnostic ondemand setting	Displays information about on-demand diagnostics.
show diagnostic result module *slot* [**test** [*test-name* | **all**]] [**detail**]	Displays information about the results of a diagnostic.
show diagnostic simulation module *slot*	Displays information about a simulated diagnostic.
show diagnostic status module *slot*	Displays the test status for all tests on a module.
show module	Displays module information including the online diagnostic test status.
show diagnostic eem action	Displays the status of the corrective (recovery) action.

Configuration Examples for Online Diagnostics

This example shows how to start all on-demand tests on a module:

diagnostic start module 6 test all

This example shows how to activate a test and set the test interval for a test on a module:

configure terminal

diagnostic monitor module 6 test 2

diagnostic monitor interval module 6 test 2 hour 3 min 30 sec 0

Additional References

For additional information related to implementing online diagnostics, see the following sections:

Related Documents

Related Topic	Document Title
Online diagnostics CLI commands	Cisco MDS 9000 Series Command Reference

Feature History for Online Diagnostics

Table 30: Feature History for Online Diagnostics, on page 199 lists the release history for this feature.

Table 30: Feature History for Online Diagnostics

Feature Name	Releases	Feature Information
Generic Online Diagnostics (GOLD)	6.2	This feature was introduced.

CHAPTER 11

Configuring Interswitch Link Diagnostics

This chapter describes how to configure the ISL Diagnostics on the Cisco MDS switches.

Information About ISL Diagnostics

ISL Diagnostics feature helps in validating the health of interswitch links between the Cisco MDS switches in a network.

The following tests can be performed using ISL diagnostics:

- Latency test
- Single Hop Traffic Test
- Multihop End-to-End Traffic Test

Supported Platforms

ISL diagnostics is supported on the following platforms:

- Cisco MDS 9700 Series Switches
- Cisco MDS 9396S Switch
- Cisco MDS 9396T Switch
- Cisco MDS 9148T Switch
- Cisco MDS 9132T Switch
- Cisco MDS 9124V Switch
- Cisco MDS 9148V Switch

ISL diagnostics is supported on the following FC modules of the Cisco MDS 9700 Switches:

- 16-Gbps modules on Cisco MDS 9700 Series Switches
  - DS-X9448-768K9
  - DS-X9334-K9

Guidelines and Limitations

- Cisco MDS 9124V 64-Gbps 24-Port Switching Module (DS-9124V) and Cisco MDS 9148V 64-Gbps 48-Port Switching Module (DS-9148V) does not support latency test.
- Cisco MDS 9700 48-port 64-Gbps Switching Module (DS-X9748-3072K9) does not support latency test.
- The diagnostic tests can be performed between two supported modules of different switch families on either side.
- The ISL diagnostic support on the modules is limited to the generator and reflector ports only.
- ISL diagnostics are not supported on other non-MDS switches such as Nexus 2000, Nexus 5000, and so on.
- ISL diagnostics are not supported on any of the FCoE and IPS ports of the Cisco MDS Switches.
- ISL diagnostics are not supported on 16 Gbps FEC-enabled links, but it is supported on 16 Gbps links without FEC.
- ISL diagnostics is not supported on Dense Wavelength-Division Multiplexing (DWDM) links.

Latency Test

Latency test measures the latency of an ISL between two Cisco MDS switches. You can run the latency test on ports that are on the same switch. You can use the **diagnostic isl reflector** and **diagnostic isl generator** commands for running the test. For more information, see Configuring Interswitch Link Diagnostics, on page 201.

Figure 7: Latency Test Running on Ports on the Same Switch

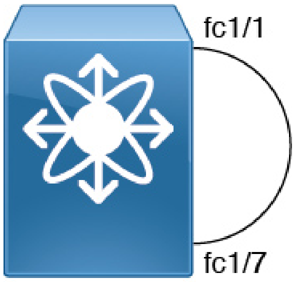

The frame is looped back by the reflector switch port to the generator switch where the timestamps are captured.

Timestamps allow the latency of the link to be measured in both directions as well as the latency of the reflector port. The cable length is calculated using only the link latencies. The accuracy of the reported cable length is +/- 2 meters. On Cisco MDS switches, the cable length (for latency test) is validated up to 50 meters of the cable length.

Note When a latency test is executed, both the generator and reflector ports must be in a admin down ('shutdown') state and not part of any port-channel.

Figure 8: Latency Test, on page 203 shows the details for the latency test.

Figure 8: Latency Test

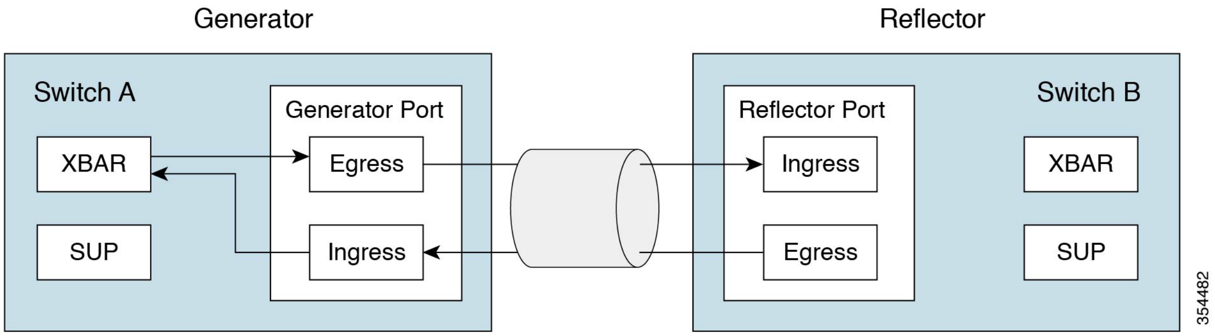

Single Hop Traffic Test

The single hop traffic test validates the health of an ISL by checking the efficiency of the ISL to handle traffic at various frame rates. You can run the single hop traffic test on ports that are on the same switch. You can use the **diagnostic isl reflector** and **diagnostic isl generator** commands for running the test. For more information, see Configuring Interswitch Link Diagnostics, on page 201.

Figure 9: Single Hop Traffic Test Running on Ports on the Same Switch

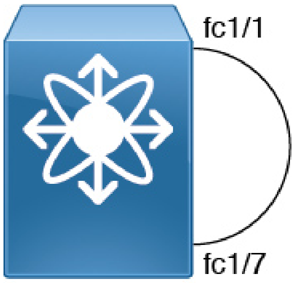

Fibre Channel (FC) frames are generated in the generator switch using the internal traffic generator facility available in the MAC hardware. These frames are transmitted from the generator switch port over the ISL under test. The reflector switch receives the frames, switches them via the normal fabric switching path and transmits the frames back through the received port onto the ISL under test.

The efficiency of the ISL traffic is calculated based on the number of packets received back on the generator switch port.

Note When a single hop test is executed, both the generator and reflector ports must be in a admin down ('shutdown') state and not part of any port-channel.

The traffic test returns with an error in the following scenarios:

- If the ISL is not up.
- If the generator port does not have an internal traffic generator facility.
- When the reflector is not put in a loopback mode.

Figure 10: Single Hop Traffic Test, on page 204 shows the details for Single Hop Traffic Test:

Figure 10: Single Hop Traffic Test

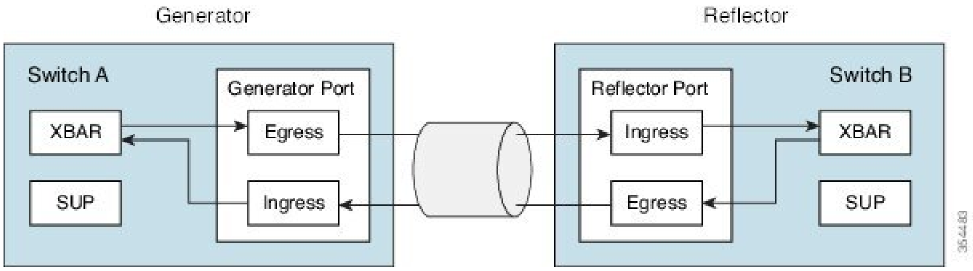

The traffic traverses through all the crossbar links.

Multihop End-to-End Traffic Test

The Multihop test evaluates the health of the ISL between a host switch and a target switch in a fabric.

Before connecting a host to a target in a fabric, test the fabric path between the host port and the target port, using a Multihop test.

Multiple hops can exist between the host switch and the target switch. No specific configuration is required on the intermediate switches.

Note The intermediate switches in the fabric can have any interface or link, for example, FC, FCoE, IPS, and so on, between them, as long as a route exists between the generator and the reflector port.

Fibre channel (FC) frames are generated at the generator switch port and transmitted to the first hop link. These frames traverse intermediate switches until they reach the reflector switch. The reflector switch then switches the frames and returns them to the generator switch. Based on the number of packets received on the generator switch, the efficiency of the ISL is displayed.

The Multihop traffic test is based on the domain IDs of the generator and reflector switches.

Note When a Multihop traffic test is executed , both the generator and reflector ports must be in a admin down ('shutdown') state and not part of any port-channel however, there can be traffic running over the ISLs that is used by the Multihop traffic test.

Figure 11: Multihop End-to-End Traffic Test , on page 205 shows the details for Multihop End-to-End Traffic Test:

Figure 11: Multihop End-to-End Traffic Test

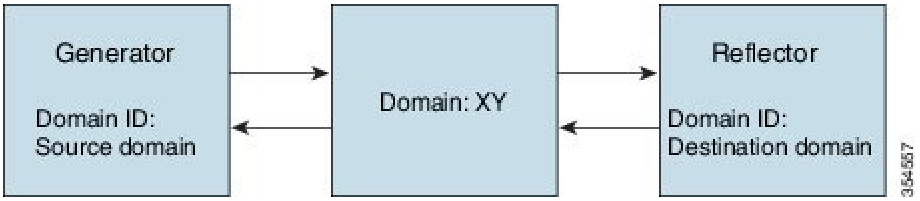

The traffic traverses through all the crossbar links.

Configuring ISL Diagnostics

Configuring Latency Test on Cisco MDS 9700 Series Switches

To configure the latency test between the generator and reflector switches perform the following tasks:

Note This task is supported on the following modules of the Cisco MDS 9700 series switch:

- DS-X9448-768K9
- DS-X9334-K9
- DS-X9648-1536K9

Procedure

Step 1 Enable the test interface on the reflector switch for testing latency by setting it to loopback mode using the following command:

switch B# **diagnostic isl reflector latency_test loop-back interface** *interface id* **enable**

Step 2 Configure the generator switch to run the test and display the results:

switch A# **diagnostic isl latency-test interface** *interface id*

Step 3 To disable the reflector port for latency testing configure the following command on the reflector switch:

switch B# **diagnostic isl reflector latency_test loop-back interface** *interface id* **disable**

Latency Test

Before starting the latency test, shutdown the testing interface ports on both the switches:

```
switch A# shutdown interface fc4/1
switch B# shutdown interface fc1/13
```

This example shows how to enable a port on a reflector switch for latency testing:

```
switch B# diagnostic isl reflector latency_test loop-back interface fc1/13 enable
Reflector Configuration Successful.
```

This example shows how to run the latency test:

```
switch A# diagnostic isl latency-test interface fc4/1
Waiting for sync to be achieved on the link ....
Sync is achieved, Link has been initialized.
Starting the test ....
---------------------------------------------------------
Latency test Result for port:                   fc4/1
Latency in the switch (in ns):                  399
Latency in the cable (in ns):                   39
Length of the cable (accuracy +/- 2m):          4 m
---------------------------------------------------------
```

This example shows how to disable a port on a reflector switch for latency testing:

```
switch B# diagnostic isl reflector latency_test loop-back interface fc1/13 disable
Reflector Configuration Successful.
```

Restart the disabled interface ports:

```
switch A# no shutdown interface fc4/1
switch B# no shutdown interface fc1/13
```

Configuring Latency Test on Other Supported Platforms

To configure the latency test between the generator and reflector switches perform the following tasks:

Note This task is supported on the following Cisco MDS switches:

- Cisco MDS 9396S Switch
- Cisco MDS 9396T Switch
- Cisco MDS 9148T Switch
- Cisco MDS 9132T Switch

Procedure

Step 1 Enable the test interface on the reflector switch for testing latency by setting it to loopback mode using the following command:

switch B# **system health isl reflector latency_test loop-back interface** *interface id* **enable**

Step 2 Configure the generator switch to run the test and display the results:

switch A# **system health isl latency-test interface** *interface id*

Step 3 To disable the reflector port for latency testing configure the following command on the reflector switch:

switch B# **system health isl reflector latency_test loop-back interface** *interface id* **disable**

Latency Test

Before starting the latency test, shutdown the testing interface ports on both the switches:

```
switch A# shutdown interface fc1/13
switch B# shutdown interface fc4/25
```

This example shows how to enable a port on a reflector switch for latency testing:

```
switch B# system health isl reflector latency_test loop-back interface fc4/25 enable
Reflector Configuration Successful.
```

This example shows how to run the latency test:

```
switch A# system health isl latency-test interface fc1/13
Waiting for sync to be achieved on the link ....
Sync is achieved, Link has been initialized.
Starting the test ....


---------------------------------------------------------
Latency test Result for port: fc1/13
Latency in the switch (in ns): 5504
Latency in the cable (in ns): 664
Length of the cable (accuracy +/- 2m): 4.816514 m
---------------------------------------------------------
```

This example shows how to disable a port on a reflector switch for latency testing:

```
switch B# system health isl reflector latency_test loop-back interface fc4/25 disable
Reflector Configuration Successful.
```

Restart the disabled interface ports:

```
switch A# no shutdown interface fc1/13
switch B# no shutdown interface fc4/25
```

Configuring a Single Hop Traffic Test on Cisco MDS 9700 Series Switches

To configure single hop traffic test between a generator switch and a reflector switch, perform the following tasks:

Procedure

Step 1 Enable the test interface on the reflector switch for single hop traffic test by setting it to loopback mode using the following command:

switch B# **diagnostic isl reflector traffic_test loop-back interface** *interface id* **enable**

Step 2 Configure the interface using one of the following options:

- Configure the interface on the generator switch to run the traffic test for a given frame count, frame size, and rate (link speed) parameters:

  switch A# **diagnostic isl generator interface** *interface id* **start frame-count** *number* **rate** *value* **frame_size min** *minimum size* **max** *maximum size* **step** *num*

- Configure the interface on the generator switch to run the traffic test for a given duration, frame size, and rate (link speed) parameters:

  switch A# **diagnostic isl generator interface** *interface id* **start duration** *seconds* **rate** *value* **frame_size min** *minimum size* **max** *maximum size* **step** *num*

Step 3 Disable the reflector port for single hop traffic test:

switch B# **diagnostic isl reflector traffic_test loop-back interface** *interface id* **disable**

Step 4 View the results of the single hop traffic test:

switch B# **show diagnostic isl result interface** *interface id*

Step 5 To stop the single hop traffic test, use the following command:

switch A# **diagnostic isl generator interface** *interface id* **stop**

Single Hop Traffic Test

Before starting the single hop traffic test, shutdown the testing interface ports on both the switches:

```
switch A# shutdown interface fc4/5
switch B# shutdown interface fc9/37
```

This example shows how to enable the test interface on the reflector switch for the single hop traffic test by setting it to loopback mode:

```
switch B# diagnostic isl reflector traffic_test loop-back interface fc9/37 enable
Reflector Configuration Successful.
```

This example shows how to run a traffic test on the generator switch for a particular duration, speed, and frame size parameters:

```
switch A# diagnostic isl generator interface fc4/5 start duration 100 rate 25% frame_size
min 16 max 517 step 1
```

This example shows the results of the single hop traffic test:

```
switch A# show  diagnostic isl result interface fc4/5
---------------------------------------------------------
Single hop  Traffic test Result for port: fc4/5
Packets Transmitted:                           30621868
Packets Recieved:                              30621868
ISL traffic Efficiency (percent):              100.0000
---------------------------------------------------------
```

Restart the disabled interface ports:

```
switch A# no shutdown interface fc4/5
switch B# no shutdown interface fc9/37
```

Configuring a Single Hop Traffic Test on Other Supported Platforms

To configure a single hop traffic test between a generator switch and a reflector switch, perform the following tasks:

Procedure

Step 1 Enable the test interface on the reflector switch for a single hop traffic test by setting it to loopback mode:

switch B# **system health isl reflector traffic_test loop-back interface** *interface id* **enable**

Step 2 Configure the interface using one of the following options:

- Configure the interface on the generator switch to run the traffic test for a given frame count, frame size, and rate (link speed) parameters:

  switch A# **system health isl generator interface** *interface id* **start frame-count** *number* **rate** *value* **frame_size min** *minimum size* **max** *maximum size* **step** *num*

- Configure the interface on the generator switch to run the traffic test for a given duration, frame size, and rate (link speed) parameters:

  switch A# **system health isl generator interface** *interface id* **start duration** *seconds* **rate** *value* **frame_size min** *minimum size* **max** *maximum size* **step** *num*

Step 3 To disable the reflector port for single hop traffic test, configure the following command on the reflector switch:

switch B# **system health isl reflector traffic_test loop-back interface** *interface id* **disable**

Step 4 View the results of the single hop traffic test:

switch B# **show system health isl result interface** *interface id*

Step 5 To stop the single hop traffic test:

switch A# **system health isl generator interface** *interface id* **stop**

Single Hop Traffic Test

Before starting the single hop traffic test, shutdown the testing interface ports on both the switches:

```
switch A# shutdown interface fc12/16
switch B# shutdown interface fc9/37
```

This example shows how to enable the test interface on the reflector switch for the single hop traffic test by setting it to loopback mode:

```
switch B# system health isl reflector traffic_test loop-back interface fc9/37 enable
```

```
Reflector Configuration Successful.
```

This example shows how to run a traffic test for duration parameter on a generator switch:

```
switch A# system health isl generator interface fc12/16 start duration 100
Waiting for sync to be achieved on the link .....
Link initialized successfully. Starting the test.

switch A# system health isl generator interface fc12/16 stop
---------------------------------------------------------
Traffic test Result for port:                    fc12/16
Packets Transmitted:                             5293153
Packets Recieved:                                5293153
ISL traffic Efficiency (percent):                100.0000
---------------------------------------------------------

switch B# system health isl reflector traffic_test loop-back interface fc9/37 disable
Reflector Configuration Successful.
```

This example shows the results of the single hop traffic test:

```
switch A# show system health isl result interface fc12/16
---------------------------------------------------------
Single hop  Traffic test Result for port:      fc12/16
Packets Transmitted:                           1019885186
Packets Recieved:                              1019885186
ISL traffic Efficiency (percent):              100.0000
---------------------------------------------------------
```

Restart the disabled interface ports:

```
switch A# no shutdown interface fc12/16
switch B# no shutdown interface fc9/37
```

Configuring a Multihop Traffic Test on Cisco MDS 9700 Series Switches

To configure a multihop traffic test between a generator switch and a reflector switch, perform the following tasks:

Note For a given VSAN, source domain, and destination domain, there can be only one test running.

Procedure

Step 1 Enable the test interface on the reflector switch by setting it to loopback mode for a given VSAN and domain ID of the generator switch for the multihop traffic test:

From Cisco MDS NX-OS Release 8.4(1) and later, use the following command:

switch B# **diagnostic isl multi_hop reflector loop-back interface** *interface id* **enable vsan** *vsan id* **source-domain** *source id*

Note To obtain the source domain, use the following command on the reflector switch:

switch B# **show fcdomain domain-list vsan** *vsan id*

For Cisco MDS NX-OS Release 8.3(2) and earlier, use the following command:

switch B# **diagnostic isl multi_hop reflector loop-back interface** *interface id* **vsan** *vsan id* **source-domain** *source id* **enable**

Step 2 Configure the interface on the generator switch, to run the multihop traffic test for a given VSAN, destination domain (domain ID of the reflector switch), frame count, link speed, and frame size parameters:

From Cisco MDS NX-OS Release 8.4(1) and later, use the following command:

switch A# **diagnostic isl multi_hop generator interface** *interface id* **start vsan** *vsan id* **dest-domain** *dest id* **frame-count** *number* **rate** *value* **frame_size min** *minimum size* **max** *maximum size* **step** *num*

For Cisco MDS NX-OS Release 8.3(2) and earlier, use the following command:

switch A# **diagnostic isl multi_hop generator interface** *interface id* **vsan** *vsan id* **dest-domain** *dest id* **startframe-count** *number* **rate** *value* **frame_size min** *minimum size* **max** *maximum size* **step** *num*

Step 3 Configure the interface on the generator switch to run the multihop traffic test for a given VSAN, destination domain (domain ID of the reflector switch), duration, rate (link speed), and frame size parameters.

From Cisco MDS NX-OS Release 8.4(1) and later, use the following command:

switch A# **diagnostic isl multi_hop generator interface** *interface id* **start vsan** *vsan id* **dest-domain** *dest id* **duration** *seconds* **rate** *value* **frame_size min** *minimum size* **max** *maximum size* **step** *num*

Note To obtain the destination domain use the following command on the generator switch:

switch A# **show fcdomain domain-list vsan** *vsan id*

For Cisco MDS NX-OS Release 8.3(2) and earlier, use the following command:

switch A# **diagnostic isl multi_hop generator interface** *interface id* **vsan** *vsan id* **dest-domain** *dest id***start duration** *seconds* **rate** *value* **frame_size min** *minimum size* **max** *maximum size* **step** *num*

Step 4 To disable the reflector port for the multihop traffic test configure the following command on the reflector switch:

From Cisco MDS NX-OS Release 8.4(1) and later, use the following command:

switch B# **diagnostic isl multi_hop reflector loop-back interface** *interface id* **disable**

For Cisco MDS NX-OS Release 8.3(2) and earlier, use the following command:

switch B# **diagnostic isl multi_hop reflector loop-back interface** *interface id* **vsan** *vsan id***source-domain** *source id* **disable**

Step 5 View the results of the multihop traffic test:

switch A# **show diagnostic isl result interface** *interface id*

Step 6 To stop the multihop traffic test:

From Cisco MDS NX-OS Release 8.4(1) and later, use the following command:

switch A# **diagnostic isl multi_hop generator interface** *interface id* **stop**

For Cisco MDS NX-OS Release 8.3(2) and earlier, use the following command:

switch A# **diagnostic isl multi_hop generator interface** *interface id* **vsan** *vsan id* **dest-domain** *dest id* **stop**

MultiHop Traffic test

Before starting the multihop traffic test, shutdown the testing interface ports on both the switches:

```
switch A# shutdown interface fc4/10
switch B# shutdown interface fc9/36
```

This example shows how to display the domain list on both the generator and reflector switches:

```
switch# show fcdomain domain-list vsan 1
Number of domains: 3
Domain ID              WWN
---------    -----------------------
0x85(133)    20:01:00:0d:ec:b7:20:01 [Principal]
0xef(239)    20:01:40:55:39:0c:70:81 [Local]
0x02(2)      20:01:00:0d:ec:b7:28:c1
```

This example shows how to enable the test interface on the reflector switch by setting it to loopback mode for a given VSAN and domain ID of the generator switch for Multihop Traffic Test:

```
switch B# diagnostic isl multi_hop reflector loop-back interface fc9/36 enable vsan 1
source_domain 239
```

This example shows how to run a traffic test on the generator switch for a particular duration, speed, and frame size parameters:

```
switch A# diagnostic isl multi_hop generator interface fc4/10 start vsan 1 dest_domain 133
 duration 100 rate 16G frame_size min 16 max 517 step 1
```

This example shows the results of the multihop traffic test:

```
switchA #show diagnostic isl result interface fc 4/10

---------------------------------------------------------
Multi hop Traffic test Result for port:    fc4/10
Packets Transmitted:                             6131424
Packets Recieved:                                6131424
ISL traffic Efficiency (percent):                100.0000
---------------------------------------------------------
```

Restart the disabled interface ports:

```
switch A# no shutdown interface fc4/10
switch B# no shutdown interface fc9/36
```

Configuring a Multihop Traffic Test on Other Supported Platforms

To configure a multihop traffic test between a generator switch and a reflector switch, perform the following tasks:

Procedure

Step 1 Enable the test interface on the reflector switch by setting it to loopback mode for a given VSAN and domain ID of the generator switch:

From Cisco MDS NX-OS Release 8.4(1) and later, use the following command:

switch B# **system health isl multi_hop reflector loop-back interface** *interface id***enable vsan** *vsan id* **source-domain** *source id*

For Cisco MDS NX-OS Release 8.3(2) and earlier, use the following command:

switch B# **system health isl multi_hop reflector loop-back interface** *interface id* **vsan** *vsan id* **source-domain** *source id* **enable**

To obtain the source domain, use the following command on the reflector switch:

switch B# **show fcdomain domain-list vsan** *vsan id*

Step 2 Configure the interface on the generator switch to run the multihop traffic test for a given VSAN, destination domain (domain ID of the reflector switch), frame count, link speed, and frame size parameters:

From Cisco MDS NX-OS Release 8.4(1) and later, use the following command:

switch A# **system health isl multi_hop generator interface** *interface id* **start vsan** *vsan id* **dest-domain** *dest id* **frame-count** *number* **rate** *value* **frame_size min** *minimum size* **max** *maximum size* **step** *num*

For Cisco MDS NX-OS Release 8.3(2) and earlier, use the following command:

switch A# **system health isl multi_hop generator interface** *interface id* **vsan** *vsan id* **dest-domain** *dest id* **start frame-count** *number* **rate** *value* **frame_size min** *minimum size* **max** *maximum size* **step** *num*

Step 3 Configure the interface on the generator switch to run the multihop traffic test for a given VSAN, destination domain (domain ID of the reflector switch), duration, rate (link speed), and frame size parameters:

From Cisco MDS NX-OS Release 8.4(1) and later, use the following command:

switch A# **system health isl multi_hop generator interface** *interface id* **start vsan** *vsan id* **dest-domain** *dest id* **duration** *seconds* **rate** *value* **frame_size min** *minimum size* **max** *maximum size* **step** *num*

For Cisco MDS NX-OS Release 8.3(2) and earlier, use the following command:

switch A# **system health isl multi_hop generator interface** *interface id* **vsan** *vsan id* **dest-domain** *dest id* **start duration** *seconds* **rate** *value* **frame_size min** *minimum size* **max** *maximum size* **step** *num*

To obtain the destination domain use the following command on the generator switch:

switch A# **show fcdomain domain-list vsan** *vsan id*

Step 4 Disable the reflector port for the multihop traffic test, configure the following command on the reflector switch:

From Cisco MDS NX-OS Release 8.4(1) and later, use the following command:

switch B# **system health isl multi_hop reflector loop-back interface** *interface id* **disable**

For Cisco MDS NX-OS Release 8.3(2) and earlier, use the following command:

switch B# **system health isl multi_hop reflector loop-back interface** *interface id* **vsan** *vsan id* **source-domain** *source id* **disable**

Step 5 View the results of the multihop traffic test:

switch A# **show system health isl result interface** *interface id*

Step 6 To stop the multihop traffic test on the generator switch, use the following command:

From Cisco MDS NX-OS Release 8.4(1) and later, use the following command:

switch A# **system health isl multi_hop generator interface** *interface id* **stop**

For Cisco MDS NX-OS Release 8.3(2) and earlier, use the following command:

switch A# **system health isl multi_hop generator interface** *interface id* **vsan** *vsan id* **dest-domain** *dest id* **stop**

Multihop Traffic test

Before starting the multihop traffic test, shutdown the testing interface ports on both the switches:

```
switch A# shutdown interface fc3/18
switch B# shutdown interface fc9/36
```

This example shows how to display the domain list on both the generator and reflector switches:

```
switch# show fcdomain domain-list vsan 1
Number of domains: 3
Domain ID              WWN
---------    -----------------------
0x85(133)    20:01:00:0d:ec:b7:20:01 [Principal]
0xef(239)    20:01:40:55:39:0c:70:81 [Local]
0x02(2)      20:01:00:0d:ec:b7:28:c1
```

This example shows how to enable loop back to the generator switch interface present in a VSAN from the reflector switch for a multihop traffic test:

```
switch B# system health isl multi_hop reflector loop-back interface fc9/36 enable vsan 1
source_domain 239
```

This example shows how to run a traffic test on the generator switch for a particular duration, speed, and frame size parameters:

```
switch A# system health isl multi_hop generator interface fc3/18 start vsan 1 dest_domain
133 duration 100 rate 16G frame_size min 16 max 517 step 1
```

This example shows the results of the multihop traffic test:

```
switch A# show system health isl result interface fc3/18
---------------------------------------------------------
Multi hop Traffic test Result for port:    fc3/18
Packets Transmitted:                             3065550
Packets Recieved:                                3065550
ISL traffic Efficiency (percent):                100.0000
---------------------------------------------------------
```

Restart the disabled interface ports:

```
switch A# no shutdown interface fc3/18
switch B# no shutdown interface fc9/36
```

Debugging ISL Diagnostics

The following table lists the debug commands for this feature. To display the ISL diagnostic test status use one of the following commands:

Table 31: Debug commands

Command	Reference
Cisco MDS 9700 Switches [Cisco MDS NX-OS Release 8.2(1) and earlier]	
	Displays the status of configured ISL diagnostic tests per port.
Cisco MDS 9700 Switches [Cisco MDS NX-OS Release 8.3(1) and later]	

show diagnostic isl status
`switch#` **`show diagnostic isl status`**
`Status of isl_diag tests in progress:`
`---------------------------------------------------`
`Index  Interface        Mode <Gen/Ref>`
`   Test`
`---------------------------------------------------`
` 0      fc2/41          Reflector        SH`
`Traffic Test`
`---------------------------------------------------`
	Displays the status of configured ISL diagnostic tests per port.

show diagnostic isl result interface *interface id*
`switch#` **`show diagnostic isl result interface fc 5/3`**
`---------------------------------------------------`
`Single hop Traffic test Result for port: fc5/3`
`Packets Transmitted: 30621868`
`Packets Recieved: 30621868`
`ISL traffic Efficiency (percent): 100.0000`
`---------------------------------------------------`
	Displays the results of the Single Hop or Multihop Traffic Tests

Cisco MDS 9396S, Cisco MDS 9396T, Cisco MDS 9148T, Cisco MDS 9132T [Cisco MDS NX-OS Release 8.3(1) and later]	

show system health isl result interface *interface id*
`switch#` **`show system health isl result interface fc 1/18`**
`---------------------------------------------------`
` Single hop  Traffic test Result for port:`
`  fc1/18`
`Packets Transmitted:`
`  1019885186`
`Packets Recieved:`
`  1019885186`
`ISL traffic Efficiency (percent):`
`  100.0000`
`---------------------------------------------------`
	Displays the results of the Single Hop or Multihop Traffic Tests.

<table>
<tr><th>Command</th><th>Reference</th></tr>
<tr><td><b>show system health isl status</b>
<pre>
switch# show system health isl status
Status of isl_diag tests in progress:

Index Interface   Mode <Gen/Ref>
   Test
----

 0   fc1/51   Reflector   SH
Traffic Test

</pre></td><td>Displays the status of ISL Diagnostics tests that are in progress.</td></tr>
</table>

Additional References

For additional information related to implementing online diagnostics, see the following sections:

Related Documents

Related Topic	Document Title
InterSwitch Link Diagnostics CLI commands	Cisco MDS 9000 Family Command Reference

Feature History for Online Diagnostics

Table 32: Feature History for Online Diagnostics, on page 216 lists the release history for this feature.

Table 32: Feature History for Online Diagnostics

Feature Name	Releases	Feature Information
ISL Diagnostics	7.3(0)D1(1)	This feature was introduced.

CHAPTER 12

Using Pathtrace

Pathtrace

The Pathtrace feature builds on the Traceroute feature to provide information about interfaces, such as ingress and egress interface names and the number of transmitted and received frames and errors, at each hop in the path between 2 devices in a fabric. Pathtrace provides an end-to-end view of the shortest path without the need to connect to individual switches and check the Fabric Shortest Path First (FSPF) topology hop by hop.

Pathtrace is used to trace the path from a switch on which the **pathtrace** command is run, to a destination device or all the devices in a destination domain. The Pathtrace feature works with the Fibre Channel, Fibre Channel over Ethernet (FCoE), and Fibre Channel over IP (FCIP) interfaces. Pathtrace collects information about the available paths within the fabric and provides information for devices along the shortest path. Pathtrace displays the source interface, destination interface, cost, speed, and other statistics when used with the **detail** keyword. The **pathtrace** command can also be used to display the reverse path information (from destination back to the source). If the destination cannot be reached, Pathtrace displays the device on which the connectivity terminated.

The statistics displayed for various types of interfaces are:

- Fibre Channel interface—The statistics are displayed for the associated Fibre Channel interfaces.
- Virtual Fibre Channel (VFC) interface—The statistics are displayed for the associated Ethernet interfaces.
- Fibre Channel port channel—The statistics are displayed for port channels.
- VFC port channel—The statistics are displayed for VFC port channels.
- FCIP interface or FCIP port channel—The statistics are displayed for the FCIP interfaces or FCIP port channels.

Guidelines and Limitations for Pathtrace

- Pathtrace is not supported on Cisco MDS switches that are operating in the Cisco NPV mode.
- Pathtrace does not support interop mode.
- Pathtrace is supported only on Cisco MDS switches and not on other vendor switches.

- Pathtrace does not support virtual domains (Inter-VSAN Routing [IVR] for Pathtrace).
- Pathtrace is not manageable via Simple Network Management Protocol (SNMP).
- Pathtrace supports a maximum of 16 hops without the reverse option, and 8 hops with the reverse option.
- Statistics are displayed only for egress interfaces.
- Statistics for FCIP and FCIP port-channel interfaces are not displayed for devices in the path running Cisco MDS NX-OS Release 6.2(5).

Pathtrace Multipath

The Pathtrace Multipath feature builds on the Pathtrace feature to collect and display all Equal-Cost Multi-Path (ECMP) paths and statistics between source and destination switches. This feature provides information for all the links between the 2 endpoints to be displayed, including individual equal-cost links of a port channel. This feature can aid in troubleshooting difficult situations, for example, when a single link in a port channel has errors and the remaining do not.

Guidelines and Limitations for Pathtrace Multipath

- Pathtrace Multipath is not supported on Cisco NPV switches.
- Pathtrace Multipath does not support interop mode.
- Pathtrace Multipath is supported only on Cisco MDS switches and not on other vendor switches.
- Pathtrace Multipath does not support virtual domains (Inter-VSAN Routing [IVR] for Pathtrace Multipath).
- Pathtrace Multipath is not manageable via SNMP.
- Pathtrace Multipath does not have any limitation on the number of hops between the endpoints unlike the Pathtrace feature.
- Pathtrace Multipath is supported on F ports that are connected to Qlogic and Emulex host bus adapters (HBAs).

Using Pathtrace or Pathtrace Multipath

To display per-hop interface information along the paths between 2 devices, run this command:

switch# **pathtrace** {**domain** *id* | **fcid** *id*} **vsan** *id* [[**reverse**] [**detail**] | [**multipath**]]

The following example shows how to trace the path between a switch in which the command is executed and an edge device, using the edge device's FCID:

```
switch# pathtrace fcid 0xca016c vsan 2000
The final destination port type is F_Port
------------------------------------------------------------------------
Hop Domain In-Port          Out-Port         Speed Cost  Switchname
------------------------------------------------------------------------
0   111    embedded         fc1/6            4G    250   switch1
1   202    fc1/6            fc1/1            2G    -     switch2
NOTE: The stats are displayed for the egress interface only
```

The following example shows how to trace both the forward path and the return path between a switch in which the command is executed and an edge device, using the edge device's FCID:

```
switch# pathtrace fcid 0xca016c vsan 2000 reverse
The final destination port type is F_Port
------------------------------------------------------------------------
Hop Domain In-Port          Out-Port         Speed Cost  Switchname
------------------------------------------------------------------------
0   111    embedded         fc1/6            4G    250   switch1
1   202    fc1/6            fc1/1            2G    -     switch2
2   202    embedded         fc1/6            4G    250   switch2
3   111    fc1/6            embedded         -     -     switch1
NOTE: The stats are displayed for the egress interface only
```

The following example shows how to display detailed information about the interfaces (both the forward path and the return path) between a switch in which the command is executed and an edge device, using the edge device's FCID:

```
switch# pathtrace fcid 0xca016c vsan 2000 reverse detail
The final destination port type is F_Port
-------------------------------------------------------------------------------
Hop 0       Domain In-Port          Out-Port         Speed Cost  Switchname
            111    embedded         fc1/6            4G    250   switch1
-------------------------------------------------------------------------------
Stats for egress port: fc1/6
    TxRt(B/s): 2944
    RxRt(B/s): 3632
        TxB_B: 32
        RxB_B: 32
      TxFrame: 137467
      RxFrame: 137475
       Errors: 0
      Discard: 0
          CRC: 0
-------------------------------------------------------------------------------
Hop 1       Domain In-Port          Out-Port         Speed Cost  Switchname
            202    fc1/6            fc1/1            2G    -     switch2
-------------------------------------------------------------------------------
Stats for egress port: fc1/1
    TxRt(B/s): 1424
    RxRt(B/s): 1528
        TxB_B: 0
        RxB_B: 32
      TxFrame: 711
      RxFrame: 649
       Errors: 0
      Discard: 15
          CRC: 0
-------------------------------------------------------------------------------
Hop 2       Domain In-Port          Out-Port         Speed Cost  Switchname
            202    embedded         fc1/6            4G    250   switch2
-------------------------------------------------------------------------------
Stats for egress port: fc1/6
    TxRt(B/s): 3632
    RxRt(B/s): 2952
        TxB_B: 32
        RxB_B: 32
      TxFrame: 137476
      RxFrame: 137467
       Errors: 0
      Discard: 0
```

```
          CRC: 0
-------------------------------------------------------------------------------
Hop 3       Domain In-Port          Out-Port         Speed Cost  Switchname
            111    fc1/6            embedded         -     -     switch1
-------------------------------------------------------------------------------
Stats for egress port: embedded
    TxRt(B/s): -
    RxRt(B/s): -
        TxB_B: -
        RxB_B: -
      TxFrame: -
      RxFrame: -
       Errors: -
      Discard: -
          CRC: -
NOTE: The stats are displayed for the egress interface only
```

The following example shows how to trace all the links (including equal-cost parallel links) in the paths between all the edge devices in a domain and a switch in which the command is executed:

```
switch# pathtrace domain 238 vsan 1 multipath
***NOTE ***
 I - Ingress
 E -  Egress
 M - Member Port-channel
 * - Fport
...................................................
PATH 1   switch1 switch2
Domain   236       235
......................................................
-----------------------------------------------------------------------
HOP 1   switch1(fc1/11)(E)-------(I)(fc1/12)switch2
__________________________________________________________________________________________
Interface  Spd(G)  Tx(B/s)  Rx(B/s)  TxB2B  RxB2B  Errors  Discards   CRC
TxWait(1s/1m/1h/72h) FibDrops   ZoneDrops
__________________________________________________________________________________________
(E)fc1/11  8.0     84       44       64     64     0       2          0     0%/0%/0%/0%
       -          -
(I)fc1/12  8.0     44       84       64     64     0       0          0     0%/0%/0%/0%
       -          -
__________________________________________________________________________________________
HOP 2   switch2(fc1/3)(E)*End Device
__________________________________________________________________________________________
Interface   Spd(G)  Tx(B/s)  Rx(B/s)  TxB2B  RxB2B  Errors  Discards   CRC
TxWait(1s/1m/1h/72h) FibDrops    ZoneDrops
__________________________________________________________________________________________
(E)fc1/3    4.0     0        0        16     64     0       0          0     0%/0%/0%/0%
       -            -
..........................................................................................
PATH 2   switch1 switch2
Domain   236       235
..........................................................................................
__________________________________________________________________________________________
HOP 1   switch1(fc1/12)(E)-------(I)(fc1/11)switch2
__________________________________________________________________________________________
Interface  Spd(G)  Tx(B/s)  Rx(B/s)   TxB2B   RxB2B  Errors  Discards  CRC
TxWait(1s/1m/1h/72h) FibDrops     ZoneDrops
__________________________________________________________________________________________
(E)fc1/12  8.0     64       180       64      64     0       0         0     0%/0%/0%/0%
        -            -
(I)fc1/11  8.0     180      64        64      64     0       0         0     0%/0%/0%/0%
        -            -
__________________________________________________________________________________________
```

```
HOP 2   switch2(fc1/3)(E)*End Device
------------------------------------------------------------------------------------------
Interface  Spd(G)  Tx(B/s)  Rx(B/s) TxB2B  RxB2B   Errors  Discards   CRC
TxWait(1s/1m/1h/72h) FibDrops      ZoneDrops
------------------------------------------------------------------------------------------
(E)fc1/3   4.0     0        0        16     64      0       0          0     0%/0%/0%/0%
        -             -


switch# pathtrace domain 132 vsan 447   multipath
***NOTE ***
I - Ingress
E -  Egress
M - Member Port-channel
* - Fport
...................................................
PATH 1   switch1 switch2
Domain   187       132
......................................................
------------------------------------------------------------------------------------------
HOP 1              switch1(port-channel216)(E)-------(I)(port-channel216)switch2
------------------------------------------------------------------------------------------
Interface          InputRate(B/s)     OutputRate(B/s)    InputFrames(/sec)
OutputFrames(/sec)
------------------------------------------------------------------------------------------
(E)port-channel216 3393959           640827945         161838662680576    1375239938244608

   (M)fcip50       292049             55048436           3239                27507

   (M)fcip51       291539             55052889           3237                27508

   (M)fcip52       291702             55080573           3239                27522

   (M)fcip53       278265             52552382           3090                26258

   (M)fcip54       278291             52561525           3090                26263

   (M)fcip55       278346             52559754           3090                26262

   (M)fcip65       291647             55073072           3238                27518

   (M)fcip66       278491             52584017           3092                26274

   (M)fcip67       278362             52571056           3091                26268

   (M)fcip86       278290             52554341           3090                26259

   (M)fcip87       278426             52587737           3092                26276

   (M)fcip88       278551             52602163           3093                26283

(I)port-channel216 640830213          3394016           1375252823146496    161842957647872

   (M)fcip50       55058685           292105             27512               3240

   (M)fcip51       55080107           291690             27522               3239

   (M)fcip52       55097520           291794             27530               3240

   (M)fcip53       52559881           278311             26262               3090

   (M)fcip54       52570959           278345             26268               3091

   (M)fcip55       52571081           278410             26268               3091
   (M)fcip65       55051714           291539             27507              3237

   (M)fcip66       52564219           278387             26264              3091

   (M)fcip67       52562847           278324             26264              3090

   (M)fcip86       52564931           278345             26265              3091

   (M)fcip87       52571632           278350             26268              3091

   (M)fcip88       52576637           278416             26271              3091


switch# pathtrace domain 83 vsan 70 multipath
***NOTE ***
I - Ingress
E -  Egress
M - Member Port-channel
* - Fport
...................................................
PATH 1   switch1 switch2
Domain   144       83
......................................................
-----------------------------------------------------------------------
HOP 1              switch1(vfc69)(E)-------(I)(vfc69)switch2
---------------------------------------------------------------------------------------
Interface     Spd(G)  FcoeOut(Oct)       FcoeIn(Oct)        FcoeOutPkt         FcoeInPkt

--------------------------------------------------------------------------------------
(E)vfc69      10.0    165604             153648             697                700


(I)vfc69      10.0    153716             166276             701                698
```

Note

- In the output, *embedded* indicates that the respective port is an HBA interface on an edge device.
- Some of the terminologies used in the multipath outputs are defined in the following table:

Table 33: Multipath Terminologies

Term	Description
FCIP	
InputRate(B/s)	The number of bytes received per second on the in port of an FCIP link.
OutputRate(B/s)	The number of bytes received per second on the out port of an FCIP link.
InputFrames(/sec)	The number of frames received per second on the in port of an FCIP link.
OutputFrames(/sec)	The number of frames received per second on the out port of an FCIP link.
vFC	
FcoeOut(Oct)	The number of egress FCoE octets on a vFC interface.
FcoeIn(Oct)	The number of ingress FCoE octets on a vFC interface.
FcoeOutPkt	The number of egress FCoE packets on a vFC interface.
FcoeInPkt	The number of ingress FCoE packets on a vFC interface.

CHAPTER 13

Configuring HBA Link Diagnostics

Overview

The HBA Link Diagnostics feature helps in validating the health of links between Host Bus Adapters (HBAs) and Cisco MDS switches in a network.

Servers connect to Storage Area Networks (SANs) through hardware devices called HBAs. This connectivity comprises of many optical and electrical components that may develop faults during their lifetime. This feature allows identification of faulty cables, transceivers, ASICs, drivers, firmware issues or software issues, thereby eliminating dropped frames and ensuring reliable I/O operations of the server.

Supported Platforms

HBA Link Diagnostics is supported on the following platforms:

- Cisco MDS 48-Port 16-Gbps Fibre Channel Switching Module: DS-X9448-768K9
- Cisco MDS 48-Port 32-Gbps Fibre Channel Switching Module: DS-X9648-1536K9
- Cisco MDS 24/10 SAN Extension Module (FC ports only): DS-X9334-K9
- Cisco MDS 9132T Multilayer Fabric Switch
- Cisco MDS 9148T Multilayer Fabric Switch
- Cisco MDS 9396S Multilayer Fabric Switch
- Cisco MDS 9396T Multilayer Fabric Switch

Guidelines and Limitations

- From Cisco MDS NX-OS Release 8.3(1), the HBA Link Diagnostics feature is supported in N port virtualization (NPV) mode and switch mode. This feature is supported on the following platforms:
  - Cisco MDS 9132T Multilayer Fabric Switch
  - Cisco MDS 9148T Multilayer Fabric Switch
  - Cisco MDS 9396T Multilayer Fabric Switch
  - Cisco MDS 9396S Multilayer Fabric Switch
- In Cisco MDS NX-OS Release 8.2(1), the HBA Link Diagnostics feature is supported only in switch mode on the Cisco MDS 9396S Multilayer Fabric Switch, the HBA Link Diagnostics feature in N port virtualization (NPV) mode is not supported.
- HBA links can be established in F port mode or Auto mode however, HBA link diagnostics tests can be run on interfaces in F port mode only.
- When link diagnostics tests are running, the generator and host bus adapter (HBA) ports cannot be used for regular Fibre Channel (FC) traffic or for other tests, such as, Inter Switch Link (ISL) diagnostics.
- There must be at least one free or unused port in the switch available to be used as the traffic generator port. This port needs to be in admin shutdown status during the HBA link diagnostic tests.
- When a chassis reloads, switch over, or a module hosting the generator or diagnostic ports reloads, the diagnostic tests will be terminated.
- When more than one loopback test fails, only the lowest level failure is reported. We recommend that you first fix the reported failure, and then rerun the tests.
- The port LEDs on the diagnostic port are solid green color even when traffic tests are running.
- The maximum line rate of the diagnostic port that can be tested depends on the generator port's capability and the user specified line rate. For example, if the diagnostic port is running on a 32 Gbps switching module and the generator port is running on a 16 Gbps switching module, and the rate of traffic generation is set to 50%, then the maximum line rate that is supported by the diagnostic port is 8 Gbps.
- HBA link diagnostics tests are not supported on FEC-enabled links at 16 Gbps.

HBA Link Diagnostics Tests

HBA link diagnostics is a useful tool to validate the performance and isolate faulty remote-peer and HBA components. Different types of tests are available to verify the operation of different components in the path to, and the stack of, the target device.

Link diagnostics tests are configured and controlled from the MDS switch. The target HBA and SFP must support the intended types of tests. The link is set to a diagnostic mode which removes it from the SAN fabric. The test traffic can then be run exclusively on the specific link without interfering with the fabric traffic. When the tests are complete, the link can be taken out of the diagnostic mode and returned to service in the SAN fabric.

Two ports, a diagnostic port and a generator port, are required to run the tests. The diagnostic port is the port on which the tests are run. The generator port generates traffic required to run the tests. If the generator port is not explicitly specified by the user while starting the diagnostics tests, then any port that is in admin shutdown status is selected as the generator port.

The following are the different types of link diagnostics tests available on Cisco MDS switches:

- Latency Test
- Loopback Traffic Tests

Both the link diagnostics tests can be run at different supported levels. For more information, see the Levels of HBA Link Diagnostics Tests section.

Latency Test

The Latency test measures the round-trip latency of the link between the HBA and the Cisco MDS switch.

The test frame is looped back by the HBA port to the generator switch port where timestamps are captured. Timestamps allow the latency of the link to be measured in both directions as well as the latency of the HBA port.

The latency test with optical loopback helps in determining the cable length. The cable length calculations are not applicable to any other latency tests. The accuracy of the reported cable length is within +/- 5 meters.

Loopback Traffic Tests

A loopback test sends and receives data from a single port to verify if that port is operational. The loopback traffic tests can be run at different levels. For more information, see the Levels of HBA Link Diagnostics Tests section.

Levels of HBA Link Diagnostics Tests

The HBA Link Diagnostics feature supports the following levels at which link diagnostics tests can be performed:

- Remote Switched
- MAC
- Electrical
- Optical

The following figure denotes the different levels of HBA link diagnostics tests:

Figure 12: Levels of HBA Link Diagnostics Tests

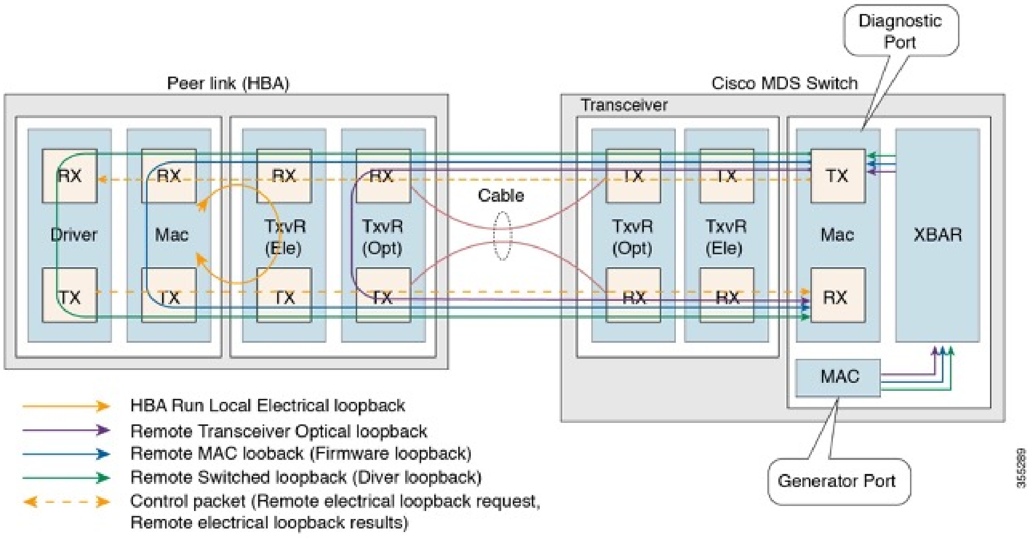
Diagnostic Port
Peer link (HBA)
Cisco MDS Switch
Transceiver
RX
RX
RX
RX
Driver
Mac
TxvR (Ele)
TxvR (Opt)
TX
TX
TX
TX
Cable
TX
TX
TX
TxvR (Opt)
TxvR (Ele)
Mac
XBAR
RX
RX
RX
MAC
Generator Port
HBA Run Local Electrical loopback
Remote Transceiver Optical loopback
Remote MAC looback (Firmware loopback)
Remote Switched loopback (Diver loopback)
Control packet (Remote electrical loopback request, Remote electrical loopback results)
355289

Remote Switched

Frames are looped back by the peer device at the highest layer that is supported by diagnostics in the stack (FC-2 or above). This functionality is implemented in the FC driver on the peer server CPU.

Note Only 100 frames will be transferred in the following scenarios:

- User input of duration as x seconds
- User input of frame-count greater than 100 frames

The requested number of frames is sent when the frame count is less than 100 frames.

MAC

Frames are looped back by the peer device at the MAC (FC-1) layer of the peer HBA. This functionality is implemented in the firmware code on the HBA.

Electrical

Frames are looped back by the peer device at the electrical stage of the transceiver (FC-0) of the peer HBA. This functionality is implemented by the peer HBA firmware programming the local transceiver for electrical loopback.

Note The electrical loopback level does not support Latency test.

Optical

A frame loopback is done in the optical part of the transceiver (FC-0) at the HBA side. The optical loopback is achieved by programming the transceiver from the firmware layer of the HBA.

Configuring HBA Link Diagnostics

To run HBA link diagnostic tests, first set the port connected to the HBA into diagnostic mode and then run the tests from this port.

When link testing is finished, move the port connected to the HBA back to service.

Configuring Link Diagnostic Mode on a Port

To configure link diagnostic mode on a port, perform the following tasks:

Before you begin

- Ensure that supported SFPs are used on the HBA.
- Install supported versions of the driver or firmware on the HBA and configure the diagnostic parameters.

Procedure

Step 1 Enter global configuration mode:

switch# **configure terminal**

Step 2 Specify the interface to be used as the diagnostic port and enter interface configuration mode:

switch(config)# **interface fc** *slot/port*

Step 3 Gracefully shut down the interface and administratively disable the traffic flow

switch(config-if)# **shutdown**

Note Configurations done through ASCII files will not enter link diagnostic mode if the interface is not in admin shut state.

Step 4 Configure the link diagnostic mode on the specified port:

switch(config-if)# **switchport link-diag**

Note To unconfigure the link diagnostic mode on the specified port, use the **no switchport link-diag** command.

Step 5 Enable the interface:

switch(config-if)# **no shutdown**

Step 6 Exit the interface:

switch(config-if)# **end**

Example

Note

- The diagnostic port will be in initializing state when it is configured using the configurations listed in "Configuring Link Diagnostic Mode on a Port" section.
- Unconfigure and reconfigure the link diagnostic mode on the switch whenever there is a change in the HBA such as driver load, unload, HBA port reset, and so on.

The following running configuration shows how to enable link diagnostic mode on an interface. Replace the placeholders with relevant values for your setup.

```
configure terminal
 interface fc <1/1>
  shutdown
  switchport link-diag
  no shutdown
  end
```

Running Link Diagnostic Tests on a Port

To run link diagnostic tests on a port, perform the following task:

Procedure

Run link diagnostic tests on a specified port:

switch# **diagnostic start interface fc** *slot/port* **test link-diag** [**duration** *seconds* | **frame-count** *count*] [**frame-size min** *min_bytes* **max** *max_bytes* **step** *step_size*] [**gen-interface fc** *slot/port*] [**level** { **remote** *levels* |**remote-all**}] [**payload** { **random | fixed** *fixed_payload* }] [**rate** *line_rate*]

Note

- By default, tests are run at all supported levels if it is not explicitly selected using the **level remote** *levels* option.
- The generator port is autoselected if it is not explicitly configured using the **gen-interface fc** *slot/port* option. For more information about this command, see the Cisco MDS 9000 Series Command Reference.
- User specified **frame-count** *count* may not match the actual number of transmitted frames due to in-switch drops.
- Avoid clearing counters or statistics on interfaces where link diagnostic tests are running.
- On interfaces where link diagnostic tests are running, new configurations that are attempted will be successful only after the completion of the link diagnostic tests.

Running Link Diagnostic Tests on a Port

This example shows how to run link diagnostic tests on a diagnostic port. In the following example, link diagnostic mode is configured on fc1/1 interface.

```
switch# diagnostic start interface fc1/1 test link-diag
```

The following command outputs display the result of the tests that are running on a diagnostic port:

```
switch# show diagnostic result interface fc1/1 test link-diag
PWWN of peer port: 21:00:00:24:ff:17:09:ac
Status: Supported (Reflector)
Reflector loopback capabilities:  Xcvr-optical Electrical
Time of Test: Thu Sep 14 00:20:11 2017
Total time taken: 30 seconds
══════════════|═════════════|═════════════|═════════════════════|══════════════|═════════════

                   |                 |                 |     Discards                |
Latency (ns)  |
Loopback Level     |    Tx Frames    |    Rx Frames    |  IN     |  OUT   |BAD
WORDS|In-Switch|External|    Status
══════════════|═════════════|═════════════|══════|══════|═══════|═══════|══════|═════════════
Remote-Switched(R) |                0|                0|       0|       0|         0|
  0|       0|       -NA-
Mac(R)             |                0|                0|       0|       0|         0|
  0|       0|       -NA-
Xcvr-optical(R)    |          1000000|          1000000|       0|       0|         0|
2136|     632|     Success
Electrical(R)      |            20000|            20000|            -NA-              |
-NA-|    -NA-|     Success
════════════════════════════════════════════════════════════════════════════════════════════


Overall Status                          : Success
Cable Length (approx. +/- 5 metres)     : 38.2 metres
```

Note

The annotation (R) denotes the remote-peer or HBA port.

The following command output displays the link diagnostics capabilities of the peer device:

```
switch# show diagnostic result interface fc1/1 test link-diag peer-capability
pWWN of Peer Port: 10:23:34:90:fa:cd:16:6c
Status: Supported (Reflector)
Reflector Loopback Capabilities:   Remote-switched MAC Xcvr-optical
```

The following running configuration shows how to unconfigure the link diagnostic mode on an interface. Replace the placeholders with relevant values for your setup.

```
configure terminal
 interface fc <1/1>
  shutdown
  no switchport link-diag
  no shutdown
  end
```

Terminating Link Diagnostic Tests on a Port

To terminate link diagnostic tests on a port, perform the following task:

Procedure

Terminate link diagnostic tests on a specified port:

switch# **diagnostic stop interface fc** *slot/port* **test link-diag**

Example: Terminating Link Diagnostic Tests on a Port

The following example shows how to terminate link diagnostic tests on a port. In the following example, link diagnostic mode is configured on fc1/1 interface.

```
switch# diagnostic stop interface fc 1/1 test link-diag
```

The following command output displays the result of the terminated tests on a diagnostic port:

```
switch# show diagnostic result interface fc 1/1 test link-diag
PWWN of peer port: 10:00:00:90:fa:c7:e1:e9
Status: Supported (Reflector)
Reflector loopback capabilities:  Remote-switched MAC Xcvr-optical
Time of Test: Wed Sep 20 12:54:59 2017
Total time taken: 10 seconds
==============|=============|============|====================|=============|=============

                   |                 |                 |     Discards              |
Latency (ns)  |
Loopback Level     |    Tx Frames    |    Rx Frames    |  IN    |  OUT   |BAD
WORDS|In-Switch|External|    Status
==============|=============|============|======|======|======|======|======|=============
Remote-Switched(R) |                0|                0|       0|       0|         0|
  0|       0|       -NA-
Mac(R)             |                0|                0|       0|       0|         0|
  0|       0|       -NA-
Xcvr-optical(R)    |              439|              439|           -NA-            |
  0|       0|    Stopped
```

```
Electrical(R)      |                0|               0|       0|       0|        0|
  0|       0|       -NA-
==========================================================================================

Overall Status                         : User Stop/Module Reload/PortDown/ELS error
                                         [DIAG TEST STOPPED]
Cable Length (approx. +/- 5 metres)    : -NA-
```

Note The annotation (R) denotes the remote-peer or HBA port.

Troubleshooting HBA Link Diagnostics

The following commands can be used to troubleshoot common HBA link diagnostics issues:

- To check if link diagnostics is enabled on an interface, use the **show interface fc** *slot/port* command.

```
switch# show interface fc1/1
fc1/1 is down (Initializing)
Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:25:40:55:39:0c:70:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port vsan is 1
    Receive data field Size is 2112
    Beacon is turned off
    Logical type is edge
    Link Diagnostics enabled
    5 minutes input rate 0 bits/sec,0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec,0 bytes/sec, 0 frames/sec
      26654656 frames input,53267399028 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      26654687 frames output,53267399756 bytes
        0 discards,0 errors
      31 input OLS,31  LRR,33 NOS,0 loop inits
      61 output OLS,0 LRR, 27 NOS, 0 loop inits
    Last clearing of "show interface" counters :  never
```

- To check if an interface is being used as the generator port, use the **show interface fc** *slot/port* command.

```
switch# show interface fc 1/1
fc1/2 is down (Administratively down)
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:02:8c:60:4f:0d:20:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port vsan is 1
    Receive data field Size is 2112
    Beacon is turned off
    Logical type is Unknown(0)
    Link Diagnostics generator port
    5 minutes input rate 0 bits/sec,0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec,0 bytes/sec, 0 frames/sec
      0 frames input,0 bytes
        0 discards,0 errors
        0 invalid CRC/FCS,0 unknown class
        0 too long,0 too short
      0 frames output,0 bytes
        0 discards,0 errors
      0 input OLS,0  LRR,0 NOS,0 loop inits
      0 output OLS,0 LRR, 0 NOS, 0 loop inits
    Last clearing of "show interface" counters :  never
```

- To check the link diagnostics tests that are running on the switch, use the **show diagnostic test link-diag status** command.

```
switch# show diagnostic test link-diag status
--------------------------------------------------------------------------------------
Index  Diag-Interface    Gen-Interface    Link-diag Status
--------------------------------------------------------------------------------------

                                          Remote-Switched(R)      MAC(R)
Electrical(R)    Xcvr-optical(R)
--------------------------------------------------------------------------------------
1       fc2/9             fc2/1             NA                    Running
   NA                  NA
```

- To collect information for Cisco technical support for this feature, use the **show tech-support link-diag** command.

CHAPTER 14

Configuring SNMP

The CLI and SNMP use common roles in all switches in the Cisco MDS 9000 Family. You can use SNMP to modify a role that was created using the CLI and vice versa.

Users, passwords, and roles for all CLI and SNMP users are the same. A user configured through the CLI can access the switch using SNMP (for example, the Cisco DCNM-SAN or the Device Manager) and vice versa.

Information About SNMP Security

SNMP is an application layer protocol that facilitates the exchange of management information between network devices. In all Cisco MDS 9000 Family switches, three SNMP versions are available: SNMPv1, SNMPv2c, and SNMPv3 (see #unique_384 unique_384_Connect_42_fig_sjx_bwq_sz).

SNMP Version 1 and Version 2c

SNMP Version 1 (SNMPv1) and SNMP Version 2c (SNMPv2c) use a community string match for user authentication. Community strings provided a weak form of access control in earlier versions of SNMP. SNMPv3 provides much improved access control using strong authentication and should be preferred over SNMPv1 and SNMPv2c wherever it is supported.

SNMP Version 3

Note From Cisco MDS NX-OS Release 8.5(1), AES-128 is the recommended encryption algorithm because it is a strong encryption algorithm. However, DES encryption is also supported.

In-Service System Downgrade (ISSD) using the **install all** command is aborted if users with DES privacy protocol are present in the SNMP database. Users need to be reconfigured using the default AES-128 or deleted. This behavior is seen in Cisco MDS NX-OS Release 8.5(1). DES user support in ISSD cases will be added in future releases. However, in case of a cold reboot, the SNMP users with DES privacy protocol are deleted.

SNMP Version 3 (SNMPv3) is an interoperable standards-based protocol for network management. SNMPv3 provides secure access to devices by a combination of authenticating and encrypting frames over the network. The security features provided in SNMPv3 are:

- Message integrity—Ensures that a packet has not been tampered with in-transit.
- Authentication—Determines the message is from a valid source.
- Encryption—Scrambles the packet contents to prevent it from being seen by unauthorized sources.

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the role in which the user resides. A security level is the permitted level of security within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet.

SNMPv3 CLI User Management and AAA Integration

The Cisco NX-OS software implements RFC 3414 and RFC 3415, including user-based security model (USM) and role-based access control. While SNMP and the CLI have common role management and share the same credentials and access privileges, the local user database was not synchronized in earlier releases.

SNMPv3 user management can be centralized at the AAA server level. This centralized user management allows the SNMP agent running on the Cisco MDS switch to leverage the user authentication service of the AAA server. Once user authentication is verified, the SNMP PDUs are processed further. The AAA server also is used to store user group names. SNMP uses the group names to apply the access/role policy that is locally available in the switch.

CLI and SNMP User Synchronization

Any configuration changes made to the user group, role, or password results in database synchronization for both SNMP and AAA.

To create an SNMP or CLI user, use either the **username** or **snmp-server user** commands.

- The auth passphrase specified in the **snmp-server user** command is synchronized as the password for the CLI user.
- The password specified in the **username** command is synchronized as the auth and priv passphrases for the SNMP user.

Users are synchronized as follows:

- Deleting a user using either command results in the user being deleted for both SNMP and the CLI.
- User-role mapping changes are synchronized in SNMP and the CLI.

Note When the passphrase/password is specified in localized key/encrypted format, the password is not synchronized.

- Existing SNMP users continue to retain the auth and priv passphrases without any changes.
- If the management station creates an SNMP user in the usmUserTable, the corresponding CLI user is created without any password (login is disabled) and will have the network-operator role.

AAA Exclusive Behavior in SNMPv3 Servers

The AAA exclusive behavior feature enables you to authenticate users based on location.

A unique SNMPv3 user is not authenticated if the user is not a local user or a remote AAA user. If the user exists in both the local and remote database, the user will be authenticated or rejected based on whether AAA exclusive behavior is enabled or not.

Table 34: AAA Exclusive Behavior Scenarios

User Location	AAA Server	AAA Exclusive Behavior	User Authentication
Local user database	Disabled	Enabled	User is authenticated.
Local user database	Enabled	Enabled	User is not authenticated.
Local user database	Enabled	Disabled	User is authenticated.
Local user database	Disabled	Disabled	User is authenticated.
Remote and local user databases (same username)	Enabled	Enabled	Remote user is authenticated, but the local user is not authenticated.
Remote and local user databases (same username)	Disabled	Enabled	Local user is authenticated, but the remote user is not authenticated.
Remote and local user databases (same username)	Disabled	Disabled	Local user is authenticated, but the remote user is not authenticated.
Remote and local user databases (same username)	Enabled	Disabled	Local user is authenticated, but the remote user is not authenticated.

Note When AAA servers are unreachable, a fallback option can be configured on the server so that a user is validated against the local user database. The SNMPv3 server returns an error if the user is not available in the local database or in the remote user database. The SNMPv3 server returns an “Unknown user” message without checking the availability of AAA servers when a user is not available in the remote user database.

Restricting Switch Access

You can restrict access to a Cisco MDS 9000 Family switch using IP access control lists (IP-ACLs).

Group-Based SNMP Access

Note Because *group* is a standard SNMP term used industry-wide, we refer to role(s) as group(s) in this SNMP section.

SNMP access rights are organized by groups. Each group in SNMP is similar to a role through the CLI. Each group is defined with three accesses: read access, write access, and notification access. Each access can be enabled or disabled within each group.

You can begin communicating with the agent once your user name is created, your roles are set up by your administrator, and you are added to the roles.

Creating and Modifying Users

You can create users or modify existing users using SNMP, DCNM-SAN, or the CLI.

- SNMP—Create a user as a clone of an existing user in the usmUserTable on the switch. Once you have created the user, change the cloned secret key before activating the user. Refer to RFC 2574.
- DCNM-SAN.
- CLI—Create a user or modify an existing user using the **snmp-server user** command.

A network-operator and network-admin roles are available in a Cisco MDS 9000 Family switch. There is also a default-role if you want to use the GUI (DCNM-SAN and Device Manager). You can also use any role that is configured in the Common Roles database.

Tip All updates to the CLI security database and the SNMP user database are synchronized. You can use the SNMP password to log into either DCNM-SAN or Device Manager. However, after you use the CLI password to log into DCNM-SAN or Device Manager, you must use the CLI password for all future logins. If a user exists in both the SNMP database and the CLI database before upgrading to Cisco MDS SAN-OS Release 2.0(1b), then the set of roles assigned to the user becomes the union of both sets of roles after the upgrade.

AES Encryption-Based Privacy

The Advanced Encryption Standard (AES) is the symmetric cipher algorithm. The Cisco NX-OS software uses AES as one of the privacy protocols for SNMP message encryption and conforms with RFC 3826.

The **priv** option offers a choice of DES or 128-bit AES encryption for SNMP security encryption. Prior to Cisco MDS NX-OS Release 8.5(1), the **priv** option along with the **aes-128** token indicates that this privacy password is for generating a 128-bit AES key. AES-128 has been made the default privacy option from Cisco MDS NX-OS Release 8.5(1). This indicates that any user configured or modified from Cisco MDS NX-OS Release 8.5(1) will use **aes-128** as the privacy option. The AES priv password can have a minimum of eight characters. If the passphrases are specified in clear text, you can specify a maximum of 64 characters. If you use the localized key, you can specify a maximum of 130 characters.

Note For an SNMPv3 operation using the external AAA server, user configurations in the external AAA server require AES to be the privacy protocol to use SNMP PDU encryption.

Traps, Notifications, and Informs

A trap is an unacknowledged message sent from an SNMP agent to SNMP managers in SNMPv1. It is known as a notification in SNMPv2 and SNMPv3. An inform is an acknowledged message sent from an SNMP agent to an SNMP manager. If the response is not received by the agent, it sends the inform request again.

An inform consumes more resources in the agent and in the network. Unlike a trap or notification, which is discarded by the agent as soon as it is sent, an inform request must be held in memory until a response is received, or the request times out. Traps and notifications can be sent only once, while informs can be sent multiple times. Resending informs increases traffic and contributes to a higher overhead on the network. The same traps, notifications, and informs can be sent to multiple host receivers.

Note For SNMPv3 informs to work, you must configure the Network Management Server (NMS) engineID with the SNMP user using the **snmp-server user***name* **engineID** command.

To get a Linux engineID from an NMS, start the **snmptarpd** and look for the **lcd_set_enginetime** string in the output.

```
#snmptrapd -f -D -Le 3162
lcd_set_enginetime: engineID 80 00 1F 88 80 14 D4 89 07 46 D5 74 5A 00 00 00
00 : boots=96, time=0
```

EngineID

An SNMP engineID is used to identify an entity independent of its source address. The entity consists of an SNMP engine and SNMP applications. The engineID is important when protocol data units (PDUs) must traverse proxies or Network Address Translator (NAT), or when the source entity itself has a dynamically assigned transport address or multiple source addresses.

In SNMPv3, engineIDs are also used for encoding and decoding secure PDUs. This is a requirement of the SNMPv3 user-based security model (USM).

There are two types of engineIDs, local and remote. On Cisco MDS 9000 Series switches, only remote engineIDs can be configured. The local engineID is automatically generated by the switch based on the MAC address and does not change.

LinkUp/LinkDown Notifications for Switches

You can configure which LinkUp/LinkDown notifications to enable on switches. You can enable the following types of LinkUp/LinkDown notifications:

- Cisco—Only notifications (cieLinkUp, cieLinkDown) defined in CISCO-IF-EXTENSION-MIB.my are sent for an interface, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface.

- IETF—Only notifications (LinkUp, LinkDown) defined in IF-MIB are sent for an interface, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. Only the varbinds defined in the notification definition are sent with the notifications.
- IEFT extended—Only notifications (LinkUp, LinkDown) defined in IF-MIB are sent for an interface, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. In addition to the varbinds defined in the notification definition, varbinds defined in the IF-MIB specific to the Cisco Systems implementation are sent. This is the default setting.
- IEFT Cisco—Only notifications (LinkUp, LinkDown) defined in IF-MIB and notifications (cieLinkUp, cieLinkDown) defined in CISCO-IF-EXTENSION-MIB.my are sent for an interface, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. Only the varbinds defined in the notification definition are sent with the linkUp and linkDown notifications.
- IEFT extended Cisco—Only notifications (LinkUp, LinkDown) defined in IF-MIB and notifications (cieLinkUp, cieLinkDown) defined in CISCO-IF-EXTENSION-MIB.my are sent for an interface, if ifLinkUpDownTrapEnable (defined in IF-MIB) is enabled for that interface. In addition to the varbinds defined in linkUp and linkDown notification definition, varbinds defined in the IF-MIB specific to the Cisco Systems implementation are sent with the LinkUp and LinkDown notifications.

Note For more information on the varbinds defined in the IF-MIB specific to the Cisco Systems implementation, refer to the *Cisco MDS 9000 Family MIB Quick Reference*.

Scope of LinkUp and LinkDown Trap Settings

The LinkUp and LinkDown trap settings for the interfaces generate traps based on the following scope:

Switch-level Trap Setting	Interface-level Trap Setting	Trap Generated for Interface Links?
Enabled (default)	Enabled (default)	Yes
Enabled	Disabled	No
Disabled	Enabled	No
Disabled	Disabled	No

Default Settings

Table 35: Default SNMP Settings , on page 240 lists the default settings for all SNMP features in any switch.

Table 35: Default SNMP Settings

Parameters	Default
User account	No expiry (unless configured)
Password	None

Configuring SNMP

SNMP is an application layer protocol that facilitates the exchange of management information between network devices.

Assigning SNMP Switch Contact and Location Informations

You can assign the switch contact information, which is limited to 32 characters (without spaces), and the switch location.

To configure contact and location information, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server contact NewUser**

Assigns the contact name for the switch.

Step 3 switch(config)# **no snmp-server contact NewUser**

Deletes the contact name for the switch.

Step 4 switch(config)# **snmp-server location SanJose**

Assigns the switch location.

Step 5 switch(config)# **no snmp-server location SanJose**

Deletes the switch location.

Configuring SNMP Users from the CLI

The passphrase specified in the **snmp-server user** command and the **username** command are synchronized.

Note When the passphrase or password is specified in the **localizedkey** or encrypted format, the password is not synchronized. If a configuration file is copied to the device, the passwords will not be set correctly if the configuration file was generated at a different device. Explicitly configure the desired passwords after copying the configuration into the device.

To create or modify SNMP users from the CLI, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server user joe network-admin auth sha abcd1234**

Creates or modifies the settings for a user (joe) in the network-admin role using the HMAC-SHA-96 authentication password (abcd1234).

Note From Cisco MDS NX-OS Release 8.5(1), AES-128 is the default privacy protocol for SNMPv3.

Step 3 switch(config)# **snmp-server user sam network-admin auth md5 abcdefgh**

Creates or modifies the settings for a user (sam) in the network-admin role using the HMAC-MD5-96 authentication password (abcdefgh).

Step 4 switch(config)# **snmp-server user Bill network-admin auth sha abcd1234 priv abcdefgh**

Creates or modifies the settings for a user (Bill) in the network-admin role using the HMAC-SHA-96 authentication level. AES-128 is used as the privacy encryption parameter from Cisco MDS NX-OS Release 8.5(1). Prior to Cisco MDS NX-OS Release 8.5(1), DES was used as the privacy protocol.

Step 5 switch(config)# **no snmp-server user usernameA**

Deletes the user (usernameA) and all associated parameters.

Step 6 switch(config)# **no snmp-server usam role vsan-admin**

Deletes the specified user (usam) from the vsan-admin role.

Step 7 switch(config)# **snmp-server user user1 network-admin auth md5 0xab0211gh priv 0x45abf342 localizedkey**

Specifies the password to be in localized key format (RFC 2574). The localized key is provided in hexadecimal format (for example, 0xacbdef).

Step 8 switch(config)# **snmp-server user user2 auth md5 asdgfsadf priv aes-128 asgfsgkhkj**

Configures the user2 with the MD5 authentication protocol and AES-128 privacy protocol. This command is supported in releases prior to Cisco MDS NX-OS Release 8.5(1). AES-128 is the default privacy option from Cisco MDS NX-OS Release 8.5(1).

Step 9 switch(config)# **snmp-server user joe sangroup**

Adds the specified user (joe) to the sangroup role.

Step 10 switch(config)# **snmp-server user joe techdocs**

Adds the specified user (joe) to the techdocs role.

Creating or Modifying Passwords

To create or modify passwords for SNMP users from the CLI, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server user user1 role1 auth md5 0xab0211gh priv 0x45abf342 localizedkey**

Specifies the password to be in localized key format using the DES option for security encryption.

Note From Cisco MDS NX-OS Release 8.5(1), AES-128 is the default privacy protocol for SNMPv3.

Step 3 switch(config)# **snmp-server user user1 role2 auth sha 0xab0211gh priv aes-128 0x45abf342 localizedkey**

Specifies the password to be in localized key format using the 128-bit AES option for security encryption

Note This command is supported in releases prior to Cisco MDS NX-OS Release 8.5(1). AES-128 is the default privacy option from Cisco MDS NX-OS Release 8.5(1).

The **snmp-server user** command takes the engineID as an additional parameter. The engineID creates the notification target user (see the Configuring the Notification Target User , on page 252). If the engineID is not specified, the local user is created.

Enforcing SNMPv3 Message Encryption

By default the SNMP agent allows the securityLevel parameters of authNoPriv and authPriv for the SNMPv3 messages that use user-configured SNMPv3 message encryption with auth and priv keys.

To enforce the message encryption for a user, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server user testUser enforcePriv**

Enforces the message encryption for SNMPv3 messages using this user.

Note You can only use this command for previously existing users configured with both auth and priv keys. When the user is configured to enforce privacy, for any SNMPv3 PDU request using securityLevel parameter of either noAuthNoPriv or authNoPriv, the SNMP agent responds with authorizationError.

Step 3 switch(config)# **no snmp-server user testUser enforcePriv**

Disables SNMPv3 message encryption enforcement.

Enforcing SNMPv3 Message Encryption Globally

Alternatively, you can enforce the SNMPv3 message encryption globally on all the users using the following commands:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server globalEnforcePriv**

Enforces the SNMPv3 message encryption for all the users on the switch.

Step 3 switch(config)# **no snmp-server globalEnforcePriv**

Disables global SNMPv3 message encryption enforcement.

Assigning SNMPv3 Users to Multiple Roles

The SNMP server user configuration is enhanced to accommodate multiple roles (groups) for SNMPv3 users. After the initial SNMPv3 user creation, you can map additional roles for the user.

Note Only users belonging to a network-admin role can assign roles to other users.

To configure multiple roles for SNMPv3 users from the CLI, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server user NewUser role1**

Creates or modifies the settings for an SNMPv3 user (NewUser) for the role1 role.

Step 3 switch(config)# **snmp-server user NewUser role2**

Creates or modifies the settings for an SNMPv3 user (NewUser) for the role2 role.

Step 4 switch(config)# **no snmp-server user User5 role2**

Removes role2 for the specified user (User5).

Adding Communities

You can configure read-only or read-write access for SNMPv1 and SNMPv2 users. Refer to RFC 2576.

To create an SNMPv1 or SNMPv2c community, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server community snmp_Community ro**

Adds read-only access for the specified SNMP community.

Step 3 switch(config)# **snmp-server community snmp_Community rw**

Adds read-write access for the specified SNMP community.

Step 4 switch(config)# **no snmp-server community snmp_Community**

Deletes access for the specified SNMP community (default).

Configuring SNMP Trap and Inform Notifications

You can configure the Cisco MDS switch to send notifications to SNMP managers when particular events occur.

Note Switches can forward events (SNMP traps and informs) up to 10 destinations. When you try to configure the eleventh target host for SNMP, the following message is displayed:

```
switch(config)# snmp-server host 10.4.200.173 traps version 2c noauth
reached maximum allowed targets limit
```

- You must enable the RMON traps in the SNMP configuration. For more information, refer to Configuring RMON, on page 179.
- Use the SNMP-TARGET-MIB to obtain more information on the destinations to which notifications are to be sent either as traps or as informs. Refer to the Cisco MDS 9000 Family MIB Quick Reference.

Tip The SNMPv1 option is not available with the **snmp-server host** *ip-address* **informs** command.

Note SNMP hostname using DSN server name starting with 0. or 127. is not supported.

Configuring SNMPv2c Notifications

Configuring SNMPv2c Notifications using IPv4

To configure SNMPv2c notifications using IPv4, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server host 171.71.187.101 traps version 2c private udp-port 1163**

Configures the specified host to receive SNMPv2c traps using SNMPv2c community string (private).

Step 3 switch(config)# **no snmp-server host 171.71.187.101 traps version 2c private udp-port 2162**

Prevents the specified host from receiving SNMPv2c traps on the configured UDP port using SNMPv2c community string (private).

Step 4 switch(config)# **snmp-server host 171.71.187.101 informs version 2c private udp-port 1163**

Configures the specified host to receive SNMPv2c informs using SNMPv2c community string (private).

Step 5 switch(config)# **no snmp-server host 171.71.187.101 informs version 2c private udp-port 2162**

Prevents the specified host from receiving SNMPv2c informs on the configured UDP port using SNMPv2c community string (private).

Configuring SNMPv2c Notifications using IPv6

To configure SNMPv2c notifications using IPv6, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server host 2001:0DB8:800:200C::417A traps version 2c private udp-port 1163**

Configures the specified host to receive SNMPv2c traps using SNMPv2c community string (private).

Step 3 switch(config)# **no snmp-server host 2001:0DB8:800:200C::417A traps version 2c private udp-port 2162**

Prevents the specified host from receiving SNMPv2c traps on the configured UDP port using SNMPv2c community string (private).

Step 4 switch(config)# **snmp-server host 2001:0DB8:800:200C::417A informs version 2c private udp-port 1163**

Configures the specified host to receive SNMPv2c informs using SNMPv2c community string (private).

Step 5 switch(config)# **no snmp-server host 2001:0DB8:800:200C::417A informs version 2c private udp-port 2162**

Prevents the specified host from receiving SNMPv2c informs on the configured UDP port using SNMPv2c community string (private).

Configuring SNMPv2c Notifications using DNS Name

To configure SNMPv2c notifications using the DNS Name of the SNMP notification host myhost.cisco.com, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server host myhost.cisco.com traps version 2c private udp-port 1163**

Configures the specified host to receive SNMPv2c traps using SNMPv2c community string (private).

Step 3 switch(config)# **no snmp-server host myhost.cisco.com traps version 2c private udp-port 2162**

Prevents the specified host from receiving SNMPv2c traps on the configured UDP port using SNMPv2c community string (private).

Step 4 switch(config)# **snmp-server host myhost.cisco.com informs version 2c private udp-port 1163**

Configures the specified host to receive SNMPv2c informs using SNMPv2c community string (private).

Step 5 switch(config)# **no snmp-server host myhost.cisco.com informs version 2c private udp-port 2162**

Prevents the specified host from receiving SNMPv2c informs on the configured UDP port using SNMPv2c community string (private).

Note Switches can forward events (SNMP traps and informs) up to 10 destinations.

Configuring SNMPv3 Notifications

Configuring SNMPv3 Notifications using IPv4

To configure SNMPv3 notifications using IPv4, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server host 16.20.11.14 traps version 3 noauth testuser udp-port 1163**

Configures the specified host to receive SNMPv3 traps using SNMPv3 user (testuser) and securityLevel of noAuthNoPriv.

Step 3 switch(config)# **snmp-server host 16.20.11.14 informs version 3 auth testuser udp-port 1163**

Configures the specified host to receive SNMPv3 informs using SNMPv3 user (testuser) and securityLevel of AuthNoPriv.

Step 4 switch(config)# **snmp-server host 16.20.11.14 informs version 3 priv testuser udp-port 1163**

Configures the specified host to receive SNMPv3 informs using SNMPv3 user (testuser) and securityLevel of AuthPriv.

Step 5 switch(config)# **no snmp-server host 172.18.2.247 informs version 3 testuser noauth udp-port 2162**

Prevents the specified host from receiving SNMPv3 informs.

Configuring SNMPv3 Notifications using IPv6

To configure SNMPv3 notifications using IPv6, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server host 2001:0DB8:800:200C::417A traps version 3 noauth testuser udp-port 1163**

Configures the specified host to receive SNMPv3 traps using SNMPv3 user (testuser) and securityLevel of noAuthNoPriv.

Step 3 switch(config)# **snmp-server host 2001:0DB8:800:200C::417A informs version 3 auth testuser udp-port 1163**

Configures the specified host to receive SNMPv3 informs using SNMPv3 user (testuser) and securityLevel of AuthNoPriv.

Step 4 switch(config)# **snmp-server host 2001:0DB8:800:200C::417A informs version 3 priv testuser udp-port 1163**

Configures the specified host to receive SNMPv3 informs using SNMPv3 user (testuser) and securityLevel of AuthPriv.

Step 5 switch(config)# **no snmp-server host 2001:0DB8:800:200C::417A informs version 3 testuser noauth udp-port 2162**

Prevents the specified host from receiving SNMPv3 informs.

Configuring SNMPv3 Notifications using DNS Name

To configure SNMPv3 notifications using the DNS Name of the SNMP notification host myhost.cisco.com, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server host myhost.cisco.com traps version 3 noauth testuser udp-port 1163**

Configures the specified host to receive SNMPv3 traps using SNMPv3 user (testuser) and securityLevel of noAuthNoPriv.

Step 3 switch(config)# **snmp-server host myhost.cisco.com informs version 3 auth testuser udp-port 1163**

Configures the specified host to receive SNMPv3 informs using SNMPv3 user (testuser) and securityLevel of AuthNoPriv.

Step 4 switch(config)# **snmp-server host myhost.cisco.com informs version 3 priv testuser udp-port 1163**

Configures the specified host to receive SNMPv3 informs using SNMPv3 user (testuser) and securityLevel of AuthPriv.

Step 5 switch(config)# **no snmp-server host myhost.cisco.com informs version 3 testuser noauth udp-port 2162**

Prevents the specified host from receiving SNMPv3 informs.

Authenticating SNMPv3 Users Based on Location

You can authenticate local or remote SNMPv3 users based on their location.

Use the following command in global configuration mode to enable AAA exclusive behavior in SNMPv3 servers:

Command	Purpose
snmp-server aaa exclusive-behavior enable	Enables the AAA exclusive behavior in SNMPv3 servers to authenticate users based on location. Depending on the location of the user and whether the AAA server is enabled, the exclusive behavior is as follows: • If the user is a local user and the AAA server is enabled, queries for the user will fail with an “Unknown user” message. • If the user is a remote AAA user and the AAA server is disabled, queries for the user will fail with an “Unknown user” message. • If the user is both a local user and a remote AAA user and the AAA server is enabled, the queries with remote credentials will succeed, and queries with local credentials will fail with an “Incorrect password” message. If the AAA server is disabled, queries with local remote credentials will succeed, and queries with remote credentials will fail with an “Incorrect password” message.

Enabling SNMP Notifications

Table 36: Enabling SNMP Notifications , on page 250 lists the CLI commands that enable the notifications for Cisco NX-OS MIBs.

Table 36: Enabling SNMP Notifications

MIB	DCNM-SAN Check Boxes
CISCO-ENTITY-FRU-CONTROL-MIB	**Click the Other tab and check FRU Changes.**
CISCO-FCC-MIB	**Click the Other tab and check FCC.**
CISCO-DM-MIB	**Click the FC tab and check Domain Mgr RCF.**
CISCO-NS-MIB	**Click the FC tab and check Name Server.**
CISCO-FCS-MIB	**Click the Other tab and check FCS Rejects.**
CISCO-FDMI-MIB	**Click the Other tab and check FDMI.**
CISCO-FSPF-MIB	**Click the FC tab and check FSPF Neighbor Change.**
CISCO-LICENSE-MGR-MIB	**Click the Other tab and check License Manager.**
CISCO-IPSEC-SIGNALLING-MIB	**Click the Other tab and check IPSEC.**

MIB	DCNM-SAN Check Boxes
CISCO-PSM-MIB	**Click the Other tab and check Port Security.**
CISCO-RSCN-MIB	**Click the FC tab and check RSCN ILS, and RCSN ELS.**
SNMPv2-MIB	**Click the Other tab and check SNMP AuthFailure.**
VRRP-MIB, CISCO-IETF-VRRP-MIB	**Click the Other tab and check VRRP.**
CISCO-ZS-MIB	**Click the FC tab and check Zone Rejects, Zone Merge Failures, Zone Merge Successes, Zone Default Policy Change, and Zone Unsuppd Mode.**

The following notifications are enabled by default:

- entity fru
- license
- link ietf-extended

All other notifications are disabled by default.

You can enable or disable the supported traps at the following levels:

- Switch level—You can use snmp-server enable traps command to enable all the traps in the supported MIBs at the switch level.
- Feature level—You can use snmp-server enable traps command with the feature name to enable traps at the feature level.

```
switch =>snmp-server enable traps callhome ?
event-notify    Callhome External Event Notification
smtp-send-fail  SMTP Message Send Fail notification
```

- Individual traps - You can use snmp-server enable traps command with the feature name to enable traps at the individual level.

```
switch =>snmp-server enable traps callhome event-notify ?
```

Note The snmp-server enable traps CLI command enables both traps and informs, depending on how you configured SNMP. See the notifications displayed with the snmp-server host CLI command.

To enable individual notifications, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server enable traps fcdomain**

Enables the specified SNMP (fcdomain) notification.

Step 3 switch(config)# **no snmp-server enable traps**

Disables the specified SNMP notification. If a notification name is not specified, all notifications are disabled.

Configuring the Notification Target User

You must configure a notification target user on the switch for sending SNMPv3 inform notifications to the SNMP manager.

For authenticating and decrypting the received INFORM PDU, the SNMP manager should have the same user credentials in its local configuration data store of users.

To configure the notification target user, use the following command:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server user testusr auth md5 xyub20gh priv xyub20gh engineID 00:00:00:63:00:01:00:a1:ac:15:10:03**

Configures the notification target user with the specified credentials for the SNMP manager with the specified engine ID.

Note From Cisco MDS NX-OS Release 8.5(1), AES-128 is the default privacy protocol for SNMPv3.

Step 3 switch(config)# **no snmp-server user testusr auth md5 xyub20gh priv xyub20gh engineID 00:00:00:63:00:01:00:a1:ac:15:10:03**

Removes the notification target user.

The credentials of the notification target user are used for encrypting the SNMPv3 inform notification messages to the configured SNMPmanager (as in the **snmp-server host** command).

Configuring LinkUp/LinkDown Notifications for Switches

To configure the LinkUp/LinkDown notification for a switch using NX-OS Release 4.2(1) and later, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **snmp-server enable traps link extended-link**

Enables only IETF extended linkUp notifications.

Step 3 switch(config)# **snmp-server enable traps link extended-linkDown**

Enables only IETF extended linkDown notifications.

Step 4 switch(config)# **snmp-server enable traps link cieLinkDown**

Enables Cisco extended link state down notification.

Step 5 switch(config)# **snmp-server enable traps link cieLinkUp**

Enables Cisco extended link state up notification.

Step 6 switch(config)# **snmp-server enable traps link connUnitPortStatusChange**

Enables FCMGMT The overall status of the connectivity unit Notification.

Step 7 switch(config)# **snmp-server enable traps link delayed-link-state-change**

Enables Delayed link state change.

Disable the delayed link state traps to allow the device to generate port down SNMP alerts immediately.

- Use the **no system delayed-traps enable mode FX** command on NX-OS versions 6.2(5) or lower.
- Use the **no snmp-server enable traps link delayed-link-state-change** command on NX-OS version 6.2(7) and above.

Note For upgrade between specific NX-OS release versions, ensure that delayed link state traps are disabled. When migrating from an earlier release like 5.(x) or 6.1(x) or 6.2(x) to a release 6.2(7) and above, ensure that you explicitly disable the delayed link state traps using **no snmp-server enable traps link delayed-link-state-change** command.

Step 8 switch(config)# **snmp-server enable traps link extended-linkDown**

Enables IETF extended link state down notification.

Step 9 switch(config)# **snmp-server enable traps link extended-linkUp**

Enables IETF extended link state down notification.

Step 10 switch(config)# **snmp-server enable traps link fcTrunkIfDownNotify**

Enables FCFE Link state down notification.

Step 11 switch(config)# **snmp-server enable traps link fcTrunkIfUpNotify**

Enables FCFE Link state up notification.

Step 12 switch(config)# **snmp-server enable traps link fcot-inserted**

Enables FCOT info trap.

Step 13 switch(config)# **snmp-server enable traps link fcot-removed**

Enables FCOT info trap.

Step 14 switch(config)# **snmp-server enable traps link linkDown**

Enables IETF Link state down notification.

Step 15 switch(config)# **snmp-server enable traps link linkUp**

Enables IETF Link state up notification.

Step 16 switch(config)# **no snmp-server enable traps link**

Reverts to the default setting (IETF extended).

Configuring Up/Down SNMP Link-State Traps for Interfaces

By default, SNMP link-state traps are enabled for all interfaces. Whenever a link toggles its state from Up to Down or vice versa, an SNMP trap is generated.

In some instances, you may find that you have numerous switches with hundreds of interfaces, many of which do not require monitoring of the link state. In such cases, you may elect to disable link-state traps.

To disable SNMP link-state traps for specific interfaces, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc** *slot/port*

Specifies the interface on which to disable SNMP link-state traps.

Step 3 switch(config-if)# **no link-state-trap**

Disables SNMP link-state traps for the interface.

Step 4 switch(config-if)# **link-state-trap**

Enables SNMP link-state traps for the interface.

Configuring Entity (FRU) Traps

To enable individual SNMP trap control, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 **switch(config)# snmp-server enable traps entity**

Enables individual SNMP trap control.

Step 3 **switch(config)# snmp-server enable entity_fan_status_change**

Enables entity fan status change.

Step 4 **switch(config)# snmp-server enable entity_mib_change**

Enables entity MIB change.

Step 5 **switch(config)# snmp-server enable entity_module_inserted**

Enables entity module to be inserted.

Step 6 **switch(config)# snmp-server enable entity_module_removed**

Enables entity module to be removed.

Step 7 **switch(config)# snmp-server enable entity_module_status_change**

Enables entity module status change.

Step 8 **switch(config)# snmp-server enable entity_power_out_change**

Enables entity power out change.

Step 9 **switch(config)# snmp-server enable entity_power_status_change**

Enables entity power status change.

Step 10 **switch(config)# snmp-server enable entity_unrecognised_module**

Enables entity unrecognized module.

Note All these traps have to do with legacy FRU traps.

Modifying the AAA Synchronization Time

You can modify how long Cisco NX-OS holds the synchronized user configuration.

Procedure

	Command or Action	Purpose

Step 1	**configure terminal**
	Example:
	`switch# configure terminal`
	`switch(config)#`
		Enters global configuration mode.

Step 2	**snmp-server aaa-user cache-timeout** *seconds* **Example:** `switch(config)# snmp-server aaa-user cache-timeout 1200`	Configures how long the AAA synchronized user configuration stays in the local cache. The range is from 1 to 86400 seconds. The default is 60000.

	Command or Action	Purpose
Step 3	(Optional) **copy running-config startup-config** **Example:** `switch(config)# copy running-config startup-config`	Copies the running configuration to the startup configuration.

Verifying SNMP Configuration

To display the SNMP configuration information, perform one of the following tasks:

Command	Purpose
show running-config	Displays the running configuration **Note** From Cisco MDS NX-OS Release 8.5(1), SNMP users with the configured privacy protocol, AES-128 or DES, are displayed in the running configuration. This is unlike releases prior to Cisco MDS NX-OS Release 8.5(1) where only AES-128 users were displayed the **aes-128** option in the running configuration. From Cisco MDS NX-OS Release 8.5(1), users are configured with AES-128 protocol, by default.
show interface	Displays the SNMP link-state trap configuration for a particular interface
show snmp trap	Displays all the notifications and their status
show snmp	Displays configured SNMP information, counter information for SNMP contact, location, and packet settings.

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference* .

Viewing the Up/Down SNMP Link-State Traps for Interfaces

Whenever you disable an SNMP link-state trap for an interface, the command is also added to the running configuration of the system.

To view the running configuration, use the **show running-config** command for the interface.

```
switch# no link-state-trap
switch# show running-config interface fc2/25

!Command: show running-config interface fc2/25
!Running configuration last done at: Fri Sep 20 11:28:19 2019
!Time: Fri Sep 20 11:28:22 2019

version 8.4(1)

interface fc2/25
  no link-state-trap
  no shutdown
```

To view the SNMP link-state trap configuration for a particular interface, enter the **show interface** command.

```
switch# show interface fc2/25

fc2/25 is trunking
    Hardware is Fibre Channel, SFP is long wave laser cost reduced
    Port WWN is 20:59:54:7f:ee:ea:c0:00
    Peer port WWN is 20:1d:00:de:fb:b1:7b:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Admin Speed is auto max 32 Gbps
    Operating Speed is 32 Gbps
    Rate mode is dedicated
    Port flow-control is ER_RDY
.
.
.
```

Displaying SNMP Traps

You can use the **show snmp trap** command to display all the notifications and their status.

```
switch# show snmp trap
--------------------------------------------------------------------------------
Trap type                                                              Enabled
--------------------------------------------------------------------------------
entity               : entity_mib_change                            Yes
entity               : entity_module_status_change                  Yes
entity               : entity_power_status_change                   Yes
entity               : entity_module_inserted                       Yes
entity               : entity_module_removed                        Yes
entity               : entity_unrecognised_module                   Yes
entity               : entity_fan_status_change                     Yes
entity               : entity_power_out_change                      Yes
link                 : linkDown                                     Yes
link                 : linkUp                                       Yes
link                 : extended-linkDown                            Yes
link                 : extended-linkUp                              Yes
link                 : cieLinkDown                                  Yes
link                 : cieLinkUp                                    Yes
link                 : connUnitPortStatusChange                     Yes
link                 : fcTrunkIfUpNotify                            Yes
link                 : fcTrunkIfDownNotify                          Yes
link                 : delayed-link-state-change                    Yes
link                 : fcot-inserted                                Yes
link                 : fcot-removed                                 Yes
callhome             : event-notify                                 No
callhome             : smtp-send-fail                               No
cfs                  : state-change-notif                           No
cfs                  : merge-failure                                No
fcdomain             : dmNewPrincipalSwitchNotify                   No
fcdomain             : dmDomainIdNotAssignedNotify                  No
fcdomain             : dmFabricChangeNotify                         No
rf                   : redundancy_framework                         Yes
aaa                  : server-state-change                          No
license              : notify-license-expiry                        Yes
license              : notify-no-license-for-feature                Yes
license              : notify-licensefile-missing                   Yes
license              : notify-license-expiry-warning                Yes
```

```
scsi                 : scsi-disc-complete                           No
fcns                 : reject-reg-req                               No
fcns                 : local-entry-change                           No
fcns                 : db-full                                      No
fcns                 : remote-entry-change                          No
rscn                 : rscnElsRejectReqNotify                       No
rscn                 : rscnIlsRejectReqNotify                       No
rscn                 : rscnElsRxRejectReqNotify                     No
rscn                 : rscnIlsRxRejectReqNotify                     No
fcs                  : request-reject                               No
fcs                  : discovery-complete                           No
fctrace              : route                                        No
zone                 : request-reject1                              No
zone                 : merge-success                                No
zone                 : merge-failure                                No
zone                 : default-zone-behavior-change                 No
zone                 : unsupp-mem                                   No
port-security        : fport-violation                              No
port-security        : eport-violation                              No
port-security        : fabric-binding-violation                     No
vni                  : virtual-interface-created                    No
vni                  : virtual-interface-removed                    No
vsan                 : vsanStatusChange                             No
vsan                 : vsanPortMembershipChange                     No
fspf                 : fspfNbrStateChangeNotify                     No
upgrade              : UpgradeOpNotifyOnCompletion                  No
upgrade              : UpgradeJobStatusNotify                       No
feature-control      : FeatureOpStatusChange                        No
vrrp                 : cVrrpNotificationNewMaster                   No
fdmi                 : cfdmiRejectRegNotify                         No
snmp                 : authentication                               No
```

Displaying SNMP Security Information

Use the **show snmp** commands to display configured SNMP information (see the following examples):

SNMP User Details

The following example SNMP user details:

```
switch# show snmp user
______________________________________________________________
                  SNMP USERS
______________________________________________________________
User                          Auth  Priv(enforce) Groups
____                          ____  _____________ ______
admin                         md5   des(no)       network-admin
testusr                       md5   aes-128(no)   role111
                                                  role222
______________________________________________________________
 NOTIFICATION TARGET USERS (configured  for sending V3 Inform)
______________________________________________________________
User                          Auth  Priv
____                          ____  ____
testtargetusr                 md5   des
(EngineID 0:0:0:63:0:1:0:0:0:15:10:3)
```

SNMP Community Information

The following example displays SNMP community information:

```
switch# show snmp community

Community             Group / Access      context
_________             ______________      _______
dcnm_user             network-admin
admin                 network-admin
```

SNMP Host Information

The following example displays SNMP host information:

```
switch# show snmp host
Host                            Port Version  Level  Type   SecName
____                            ____ _______  ______ ____  ______
171.16.126.34                   2162 v2c      noauth trap   public
171.16.75.106                   2162 v2c      noauth trap   public
...
171.31.58.97                    2162 v2c      auth   trap   public
...
```

The **show snmp** command displays counter information for SNMP contact, location, and packet settings. This command provides information that is used entirely by the Cisco MDS 9000 Family DCNM-SAN (refer to the System Management Configuration Guide, Cisco DCNM for SAN). See the following example:

SNMP Information

The following example displays SNMP information:

```
switch# show snmp
sys contact:
sys location:
1631 SNMP packets input
        0 Bad SNMP versions
        0 Unknown community name
        0 Illegal operation for community name supplied
        0 Encoding errors
        64294 Number of requested variables
        1 Number of altered variables
        1628 Get-request PDUs
        0 Get-next PDUs
        1 Set-request PDUs
152725 SNMP packets output
        0 Too big errors
        1 No such name errors
        0 Bad values errors
        0 General errors
Community                      Group / Access
---------                      ------
public                          rw
______________________________________________________________
                  SNMP USERS
```

```
______________________________________________________________
User                          Auth  Priv(enforce) Groups
____                          ____  _____________ ______
admin                         md5   des(no)       network-admin
testusr                       md5   aes-128(no)   role111
                                                  role222
______________________________________________________________
 NOTIFICATION TARGET USERS (configured  for sending V3 Inform)
______________________________________________________________
User                          Auth  Priv
____                          ____  ____
testtargetusr                 md5   des
(EngineID 0:0:0:63:0:1:0:0:0:15:10:3)
```

Displays SNMP Engine IDs

The following example displays SNMP engine IDs:

```
switch# show snmp engineID
Local SNMP engineID: [Hex] 8000000903000DEC2CF180
                     [Dec] 128:000:000:009:003:000:013:236:044:241:128
```

Information on SNMP Security Groups

The following example displays information on SNMP Security groups:

```
switch# show snmp group
groupname: network-admin
security model: any
security level: noAuthNoPriv
readview: network-admin-rd
writeview: network-admin-wr
notifyview: network-admin-rd
storage-type: permanent
row status: active
groupname: network-admin
security model: any
security level: authNoPriv
readview: network-admin-rd
writeview: network-admin-wr
notifyview: network-admin-rd
storage-type: permanent
row status: active
groupname: network-operator
security model: any
security level: noAuthNoPriv
readview: network-operator-rd
writeview: network-operator-wr
notifyview: network-operator-rd
storage-type: permanent
row status: active
groupname: network-operator
security model: any
security level: authNoPriv
readview: network-operator-rd
writeview: network-operator-wr
notifyview: network-operator-rd
storage-type: permanent
row status: active
```

Additional References

For additional information related to implementing SNMP, see the following sections:

MIBs

MIBs	MIBs Link
• CISCO-SNMP-TARGET-EXT-MIB • CISCO-SNMP-VACM-EXT-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

CHAPTER 15

Configuring Domain Parameters

The Fibre Channel domain (fcdomain) feature performs principal switch selection, domain ID distribution, FC ID allocation, and fabric reconfiguration functions as described in the FC-SW-2 standards.

Information About Fibre Channel Domains

The Fibre Channel domain (fcdomain) feature performs principal switch selection, domain ID distribution, FC ID allocation, and fabric reconfiguration functions as described in the FC-SW-2 standards. The domains are configured on a per VSAN basis. If you do not configure a domain ID, the local switch uses a random ID.

This section describes each fcdomain phase:

- Principal switch selection—This phase guarantees the selection of a unique principal switch across the fabric.
- Domain ID distribution—This phase guarantees each switch in the fabric obtains a unique domain ID.
- FC ID allocation—This phase guarantees a unique FC ID assignment to each device attached to the corresponding switch in the fabric.
- Fabric reconfiguration—This phase guarantees a resynchronization of all switches in the fabric to ensure they simultaneously restart a new principal switch selection phase.

Caution Changes to fcdomain parameters should not be performed on a daily basis. These changes should be made by an administrator or individual who is completely familiar with switch operations.

Figure 13: Sample fcdomain Configuration, on page 264 shows a sample fcdomain configuration.

Figure 13: Sample fcdomain Configuration

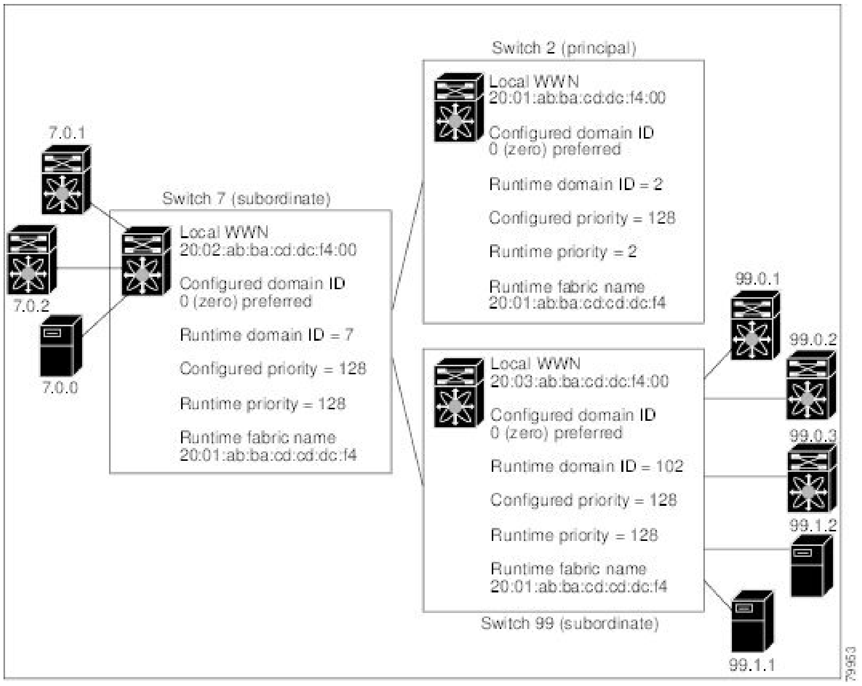

Domain Restart

Fibre Channel domains can be started disruptively or nondisruptively. If you perform a disruptive restart, reconfigure fabric (RCF) frames are sent to other switches in the fabric and data traffic is disrupted on all the switches in the VSAN (including remotely segmented ISLs). If you perform a nondisruptive restart, build fabric (BF) frames are sent to other switches in the fabric and data traffic is disrupted only on the switch.

If you are attempting to resolve a domain ID conflict, you must manually assign domain IDs. A disruptive restart is required to apply most configuration changes, including manually assigned domain IDs. Nondisruptive domain restarts are acceptable only when changing a preferred domain ID into a static one (and the actual domain ID remains the same).

Note It is not recommended to use disruptive restart followed by VSAN suspend/no-suspend, since it is used only for recovery purpose when normal restart does not solve the problem.

Note A static domain is specifically configured by the user and may be different from the runtime domain. If the domain IDs are different, the runtime domain ID changes to take on the static domain ID after the next restart, either disruptive or nondisruptive.

Tip If a VSAN is in interop mode, you cannot restart the fcdomain for that VSAN disruptively.

You can apply most of the configurations to their corresponding runtime values. Each of the following sections provide further details on how the fcdomain parameters are applied to the runtime values.

The **fcdomain restart** command applies your changes to the runtime settings. Use the **disruptive** option to apply most of the configurations to their corresponding runtime values, including preferred domain IDs (see the Domain IDs, on page 267).

Domain Manager All Optimization

Domain Manager All Optimization feature can be used to enable or disable all of the optimization modes.

Note You cannot enable all the optimizations such as Selective Restart, Fast Restart, and Scale Restart in VSANs where Interop mode is enabled (non-native modes). Also you cannot move a VSAN where the optimizations are enabled into Interop mode 1 to 4.

Domain Manager Fast Restart

As of Cisco MDS SAN-OS Release 3.0(2), when a principal link fails, the domain manager must select a new principal link. By default, the domain manager starts a build fabric phase, followed by a principal switch selection phase. Both of these phases involve all the switches in the VSAN and together take at least 15 seconds to complete. To reduce the time required for the domain manager to select a new principal link, you can enable the domain manager fast restart feature.

When fast restart is enabled and a backup link is available, the domain manager needs only a few milliseconds to select a new principal link to replace the one that failed. Also, the reconfiguration required to select the new principal link only affects the two switches that are directly attached to the failed link, not the entire VSAN. When a backup link is not available, the domain manager reverts to the default behavior and starts a build fabric phase, followed by a principal switch selection phase. We recommend using fast restart on most fabrics, especially those with a large number of logical ports (3200 or more), where a logical port is an instance of a physical port in a VSAN.

Domain Manager Scale Restart

During fabric reconfiguration, as and when principal switch assigns a domain ID to a switch (including itself), it transmits an Exchange Fabric Parameter (EFP) request. This request basically carries domain list information of the fabric. So whenever domain list grows there will be a Exchange Fabric Parameter flooded to the fabric. With this feature optimization enabled, a single consolidated Exchange Fabric Parameter request will be flooded by the principal switch once the domain identifier allocation phase is completed. This feature optimization cannot be supported in interop mode.

Scale Restart will be enabled by default in all native VSANs. It will not be enabled in interop VSANs.

Domain Manager Selective Restart

In the Fibre Channel protocol, fabric reconfiguration starts with build fabric frame flooding, which indicates to all the switches in the fabric that the fabric is changing. This process is followed by principal switch selection and domain ID allocation phases. During the build fabric flooding phase, build fabric frames are flooded on all the links. A switch may have more than one link to a peer switch. In such cases, the build fabric frame can be sent to only one of the links to the peer switch. This situation reduces the number of build fabric frames that are to be exchanged during the build fabric phase of fabric reconfiguration. Enabling this feature optimization, sends the build frame to only one of the peer switch links which benefits scaling.

Switch Priority

Any new switch can become the principal switch when it joins a stable fabric. During the principal switch selection phase, the switch with the highest priority becomes the principal switch. If two switches have the same configured priority, the switch with the lower WWN becomes the principal switch.

The priority configuration is applied to runtime when the fcdomain is restarted (see the Domain Restart, on page 264). This configuration is applicable to both disruptive and nondisruptive restarts.

fcdomain Initiation

By default, the fcdomain feature is enabled on each switch. If you disable the fcdomain feature in a switch, that switch can no longer participate with other switches in the fabric. The fcdomain configuration is applied to runtime through a disruptive restart.

Incoming RCFs

You can choose to reject RCF request frames on a per-interface, per-VSAN basis. By default, the RCF reject option is disabled (that is, RCF request frames are not automatically rejected).

The RCF reject option takes immediate effect at runtime through a disruptive restart (see the Domain Restart, on page 264)

You can configure the rcf-reject option on a per-interface, per-VSAN basis. By default, the rcf-reject option is disabled (that is, RCF request frames are not automatically rejected).

The rcf-reject option takes effect immediately. No fcdomain restart is required.

Autoreconfiguring Merged Fabrics

By default, the autoreconfigure option is disabled. When you join two switches belonging to two different stable fabrics that have overlapping domains, the following cases apply:

- If the autoreconfigure option is enabled on both switches, a disruptive reconfiguration phase is started.
- If the autoreconfigure option is disabled on either or both switches, the links between the two switches become isolated.
- RCF is expected only when auto-reconfigure is enabled in entire fabric.

The autoreconfigure option takes immediate effect at runtime. You do not need to restart the fcdomain. If a domain is currently isolated due to domain overlap, and you later enable the autoreconfigure option on both switches, the fabric continues to be isolated. If you enabled the autoreconfigure option on both switches before connecting the fabric, a disruptive reconfiguration (RCF) will occur. A disruptive reconfiguration may affect data traffic. You can nondisruptively reconfigure the fcdomain by changing the configured domains on the overlapping links and eliminating the domain overlap.

Domain IDs

Domain IDs uniquely identify a switch in a VSAN. A switch may have different domain IDs in different VSANs. The domain ID is part of the overall FC ID.

The configured domain ID can be preferred or static. By default, the configured domain ID is 0 (zero) and the configured type is preferred.

Note The 0 (zero) value can be configured only if you use the preferred option.

If you do not configure a domain ID, the local switch sends a random ID in its request. We recommend that you use static domain IDs.

When a subordinate switch requests a domain, the following process takes place (see Figure 14: Configuration Process Using the preferred Option, on page 267):

1. The local switch sends a configured domain ID request to the principal switch.
2. The principal switch assigns the requested domain ID if available. Otherwise, it assigns another available domain ID.

Figure 14: Configuration Process Using the preferred Option

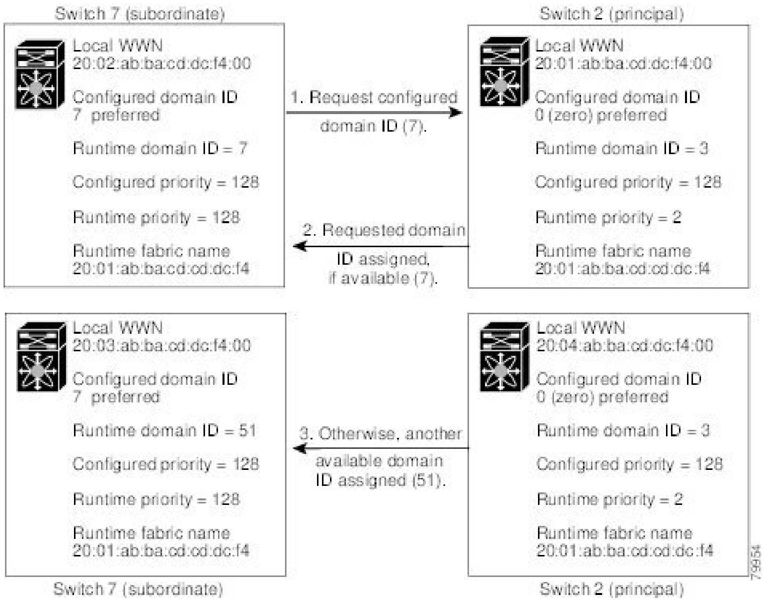

The behavior for a subordinate switch changes based on three factors:

- The allowed domain ID lists.
- The configured domain ID.
- The domain ID that the principal switch has assigned to the requesting switch.

In specific situations, the changes are as follows:

- When the received domain ID is not within the allowed list, the requested domain ID becomes the runtime domain ID and all interfaces on that VSAN are isolated.
- When the assigned and requested domain IDs are the same, the preferred and static options are not relevant, and the assigned domain ID becomes the runtime domain ID.
- When the assigned and requested domain IDs are different, the following cases apply:
  - If the configured type is static, the assigned domain ID is discarded, all local interfaces are isolated, and the local switch assigns itself the configured domain ID, which becomes the runtime domain ID.
  - If the configured type is preferred, the local switch accepts the domain ID assigned by the principal switch and the assigned domain ID becomes the runtime domain ID.

If you change the configured domain ID, the change is only accepted if the new domain ID is included in all the allowed domain ID lists currently configured in the VSAN. Alternatively, you can also configure zero-preferred domain ID.

Tip When the FICON feature is enabled in a given VSAN, the domain ID for that VSAN remains in the static state. You can change the static ID value but you cannot change it to the preferred option.

Note In an IVR without NAT configuration, if one VSAN in the IVR topology is configured with static domain IDs, then the other VSANs (edge or transit) in the topology should also be configured with static domain IDs. In an IVR NAT configuration, if one VSAN in the IVR topology is configured with static domain IDs, then the IVR domains that can be exported to that VSAN must also be assigned static domains.

Caution You must enter the fcdomain restart command if you want to apply the configured domain changes to the runtime domain.

Caution You must restart the **fcdomain** if you want to apply the configured domain changes to the runtime domain.

Note If you have configured an allowed domain ID list, the domain IDs that you add must be in that range for the VSAN. See the Configuring Allowed Domain ID Lists, on page 278.

Specifying Static or Preferred Domain IDs

When you assign a static domain ID type, you are requesting a particular domain ID. If the switch does not get the requested address, it will isolate itself from the fabric. When you specify a preferred domain ID, you are also requesting a particular domain ID; however, if the requested domain ID is unavailable, then the switch will accept another domain ID.

While the static option can be applied at runtime after a disruptive or nondisruptive restart, the preferred option is applied at runtime only after a disruptive restart (see the Domain Restart, on page 264).

Allowed Domain ID Lists

By default, the valid range for an assigned domain ID list is from 1 to 239. You can specify a list of ranges to be in the allowed domain ID list and separate each range with a comma. The principal switch assigns domain IDs that are available in the locally configured allowed domain list.

Use allowed domain ID lists to design your VSANs with non-overlapping domain IDs. This helps you in the future if you need to implement IVR without the NAT feature.

CFS Distribution of Allowed Domain ID Lists

You can enable the distribution of the allowed domain ID lists configuration information to all Cisco MDS switches in the fabric using the Cisco Fabric Services (CFS) infrastructure. This feature allows you to synchronize the configuration across the fabric from the console of a single MDS switch. Since the same configuration is distributed to the entire VSAN, you avoid possible misconfiguration and the likelihood that two switches in the same VSAN have configured incompatible allowed domains.

Use CFS to distribute the allowed domain ID list to ensure consistency in the allowed domain ID lists on all switches in the VSAN.

Note We recommend configuring the allow domain ID list and committing it on the principle switch.

For more information about CFS, see Using the CFS Infrastructure, on page 7.

Contiguous Domain ID Assignments

By default, the contiguous domain assignment is disabled. When a subordinate switch requests the principal switch for two or more domains and the domains are not contiguous, the following cases apply:

- If the contiguous domain assignment is enabled in the principal switch, the principal switch locates contiguous domains and assigns them to the subordinate switches. If contiguous domains are not available, the NX-OS software rejects this request.
- If the contiguous domain assignment is disabled in the principal switch, the principal switch assigns the available domains to the subordinate switch.

Locking the Fabric

The first action that modifies the existing configuration creates the pending configuration and locks the feature in the fabric. Once you lock the fabric, the following conditions apply:

- No other user can make any configuration changes to this feature.
- A pending configuration is created by copying the active configuration. Modifications from this point on are made to the pending configuration and remain there until you commit the changes to the active configuration (and other switches in the fabric) or discard them.

Committing Changes

To apply the pending domain configuration changes to other MDS switches in the VSAN, you must commit the changes. The pending configuration changes are distributed and, on a successful commit, the configuration changes are applied to the active configuration in the MDS switches throughout the VSAN and the fabric lock is released.

Clearing a Fabric Lock

If you have performed a domain configuration task and have not released the lock by either committing or discarding the changes, an administrator can release the lock from any switch in the fabric. If the administrator performs this task, your pending changes are discarded and the fabric lock is released.

The pending changes are only available in the volatile directory and are discarded if the switch is restarted.

FC IDs

When an N or NL port logs into a Cisco MDS 9000 Family switch, it is assigned an FC ID. By default, the persistent FC ID feature is enabled. If this feature is disabled, the following consequences apply:

- An N or NL port logs into a Cisco MDS 9000 Family switch. The WWN of the requesting N or NL port and the assigned FC ID are retained and stored in a volatile cache. The contents of this volatile cache are not saved across reboots.
- The switch is designed to preserve the binding FC ID to the WWN on a best-effort basis. For example, if one N port disconnects from the switch and its FC ID is requested by another device, this request is granted and the WWN with the initial FC ID association is released.
- The volatile cache stores up to 4000 entries of WWN to FC ID binding. If this cache is full, a new (more recent) entry overwrites the oldest entry in the cache. In this case, the corresponding WWN to FC ID association for the oldest entry is lost.
- The switch connection behavior differs between N ports and NL ports:
  - N ports receive the same FC IDs if disconnected and reconnected to any port within the same switch (as long as it belongs to the same VSAN).
  - NL ports receive the same FC IDs only if connected back to the same port on the switch to which they were originally connected.

Persistent FC IDs

When persistent FC IDs are enabled, the following consequences apply:

- The currently *in use* FC IDs in the fcdomain are saved across reboots.
- The fcdomain automatically populates the database with dynamic entries that the switch has learned about after a device (host or disk) is plugged into a port interface.

Persistent FC ID Configuration

When the persistent FC ID feature is enabled, you can enter the persistent FC ID submode and add static or dynamic entries in the FC ID database. By default, all added entries are static. Persistent FC IDs are configured on a per-VSAN basis. Follow these requirements to manually configure a persistent FC ID:

- Ensure that the persistent FC ID feature is enabled in the required VSAN.

- Ensure that the required VSAN is an active VSAN—persistent FC IDs can only be configured on active VSANs.
- Verify that the domain part of the FC ID is the same as the runtime domain ID in the required VSAN. If the software detects a domain mismatch, the command is rejected.
- Verify that the port field of the FC ID is 0 (zero) when configuring an area.

Note FICON uses a different scheme for allocating FC IDs based in the front panel port number. This scheme takes precedence over FC ID persistence in FICON VSANs.

About Unique Area FC IDs for HBAs

Note Read this section only if the HBA port and the storage port are connected to the same switch.

Some HBA ports require a different area ID than storage ports when they are both connected to the same switch. For example, if the storage port FC ID is 0x6f7704, the area for this port is 77. In this case, the HBA port's area can be anything other than 77. The HBA port's FC ID must be manually configured to be different from the storage port's FC ID.

Switches in the Cisco MDS 9000 Family facilitate this requirement with the FC ID persistence feature. You can use this feature to preassign an FC ID with a different area to either the storage port or the HBA port.

Persistent FC ID Selective Purging

Persistent FC IDs can be purged selectively. Static entries and FC IDs currently in use cannot be deleted. Table 37: Purged FC IDs , on page 271 identifies the FC ID entries that are deleted or retained when persistent FC IDs are purged.

Table 37: Purged FC IDs

Persistent FC ID state	Persistent Usage State	Action
Static	In use	Not deleted
Static	Not in use	Not deleted
Dynamic	In use	Not deleted
Dynamic	Not in use	Deleted

Guidelines and Limitations

- When you change the configuration, be sure to save the running configuration. The next time you reboot the switch, the saved configuration is used. If you do not save the configuration, the previously saved startup configuration is used.
- Domain IDs and VSAN values used in all procedures are only provided as examples. Be sure to use IDs and values that apply to your configuration.

Default Settings

Table 38: Default FC domain Parameters, on page 272 lists the default settings for all the FC domain parameters.

Table 38: Default FC domain Parameters

Parameters	Default
fcdomain feature	Enabled.
Configured domain_ID	0 (zero).
Configured domain	Preferred.
auto-reconfigure option	Disabled.
contiguous-allocation option	Disabled.
Priority	128.
Allowed list	1-239.
Fabric name	20:01:00:05:30:00:28:df.
rcf-reject	Disabled.
Persistent FC ID	Enabled.
Allowed domain_ID list configuration distribution	Disabled.

Configuring Fibre Channel Domains

This section describes the fcdomain feature.

Restarting a Domain

Domain Configuration Scenarios:

Switch Configuration

Irrespective of how the switches in VSAN 6 are configured, fcdomain restart disruptive vsan 6 causes all devices of all switches in VSAN 6 to log out, causing data traffic disruption.

Configured domain and the runtime domain are the same

Assuming that the configured domain and the runtime domain are the same on all switches, fcdomain restart vsan 6 does not cause any devices in VSAN 6 to log out.

Configured domain and runtime domain are not the same

Assuming that on some switches in VSAN 6 the configured domain and the runtime domain are not the same, fcdomain restart vsan 6 causes the devices in VSAN 6 attached to the switches whose statically configured and runtime domain differ to log out, causing data traffic disruption.

To restart the fabric disruptively or nondisruptively, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain restart vsan 1**

Is nondisruptive of data traffic over the entire network, but it can be disruptive on a switch if its configured domain is static and numerically not the same as its runtime domain (For example, the configured domain is 11 static and the runtime domain is 99).

Step 3 switch(config)# **fcdomain restart disruptive vsan1**

Distrupts data traffic across all switches in the VSAN.

Enabling Domain Manager All Optimization

To enable the Domain Manager All Optimization feature, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain optimize all vsan 3**

Enables all domain manager optimization (selective-restart, fast-restart, scale-restart) on VSAN 3.

Step 3 switch(config)# **fcdomain optimize all vsan 7 - 10**

Enables domain manager all optimization on the range of VSANs from VSAN 7 to VSAN 10.

Step 4 switch(config)# **no fcdomain optimize all vsan 8**

Disables domain manager all optimization on VSAN 8.

Enabling Domain Manager Fast Restart

To enable the domain manager fast restart feature in Cisco SAN-OS Release 3.0(2) or later, or MDS NX-OS Release 4.1(1a) or later, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain optimize fast-restart vsan 3**

Enables domain manager fast restart on VSAN 3.

Step 3 switch(config)# **fcdomain optimize fast-restart vsan 7 - 10**

Enables domain manager fast restart on the range of VSANs from VSAN 7 to VSAN 10.

Step 4 switch(config)# **no fcdomain optimize fast-restart vsan 8**

Disables (default) domain manager fast restart on VSAN 8.

Enabling Domain Manager Scale Restart

To enable the domain manager scale restart feature, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain optimize scale-restart vsan 3**

Enables domain manager scale restart on VSAN 3.

Step 3 switch(config)# **fcdomain optimize scale-restart vsan 7 - 10**

Enables (default) domain manager scale restart on the range of VSANs from VSAN 7 to VSAN 10.

Step 4 switch(config)# **no fcdomain optimize scale-restart vsan 8**

Disables domain manager scale restart on VSAN 8.

Enabling Domain Manager Selective Restart

To enable the domain manager selective restart feature in Cisco SAN-OS Release 3.0(2) or later, or MDS NX-OS Release 4.1(1a) or later, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain optimize selective-restart vsan 3**

Enables domain manager selective restart on VSAN 3.

Step 3 switch(config)# **fcdomain optimize selective-restart vsan 7 - 10**

Enables domain manager selective restart on the range of VSANs from VSAN 7 to VSAN 10.

Step 4 switch(config)# **no fcdomain optimize selective-restart vsan 8**

Disables (default) domain manager selective restart on VSAN 8.

Configuring Switch Priority

Note By default, the configured priority is 128. The valid range to set the priority is between 1 and 254. Priority 1 has the highest priority. Value 255 is accepted from other switches, but cannot be locally configured.

To configure the priority for the principal switch, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain priority 25 VSAN 99**

Configures a priority of 25 for the local switch in VSAN 99.

Step 3 switch(config)# **no fcdomain priority 25 VSAN 99**

Reverts the priority to the factory default (128) in VSAN 99.

Configuring Fabric Names

To set the fabric name value for a disabled fcdomain, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan 3**

Assigns the configured fabric name value in VSAN 3.

Step 3 switch(config)# **no fcdomain fabric-name 20:1:ac:16:5e:0:21:01 vsan 3010**

Changes the fabric name value to the factory default (20:01:00:05:30:00:28:df) in VSAN 3010.

Rejecting Incoming RCFs

To reject incoming RCF request frames, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc1/1**

switch(config-if)#

Configures the specified interface.

Step 3 switch(config-if)# **fcdomain rcf-reject vsan 1**

Enables the RCF filter on the specified interface in VSAN 1.

Step 4 switch(config-if)# **no fcdomain rcf-reject vsan 1**

Disables (default) the RCF filter on the specified interface in VSAN 1.

Enabling Autoreconfiguration

To enable automatic reconfiguration in a specific VSAN (or range of VSANs), follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain auto-reconfigure vsan 10**

Enables the automatic reconfiguration option in VSAN 10.

Step 3 switch(config)# **no fcdomain auto-reconfigure 69**

Disables the automatic reconfiguration option and reverts it to the factory default in VSAN 69.

Configuring Domain IDs

Domain IDs uniquely identify a switch in a VSAN. A switch may have different domain IDs in different VSANs. The domain ID is part of the overall FC ID.

The configured domain ID can be preferred or static. By default, the configured domain ID is 0 (zero) and the configured type is preferred.

Specifying Static or Preferred Domain IDs

Note Within a VSAN all switches should have the same domain ID type (either static or preferred). If a configuration is mixed (some switches with static domain types and others with preferred), then you may experience link isolation.

When a new domain ID is configured, the new configuration has to be applied by manually restarting the domain using the fcdomain restart command; if a discrepancy is detected between the configured domain ID and the runtime domain ID during the subsequent fabric merge, the link will be isolated.

To specify a static or preferred domain ID, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain domain 3 preferred vsan 8**

Configures the switch in VSAN 8 to request a preferred domain ID 3 and accepts any value assigned by the principal switch. The domain is range is 1 to 239.

Step 3 switch(config)# **no fcdomain domain 3 preferred vsan 8**

Resets the configured domain ID to 0 (default) in VSAN 8. The configured domain ID becomes 0 preferred.

Step 4 switch(config)# **fcdomain domain 2 static vsan 237**

Configures the switch in VSAN 237 to accept only a specific value and moves the local interfaces in VSAN 237 to an isolated state if the requested domain ID is not granted.

Step 5 switch(config)# **no fcdomain domain 18 static vsan 237**

Resets the configured domain ID to factory defaults in VSAN 237. The configured domain ID becomes 0 preferred.

Configuring Allowed Domain ID Lists

If you configure an allowed list on one switch in the fabric, we recommend that you configure the same list in all other switches in the fabric to ensure consistency or use CFS to distribute the configuration.

To configure the allowed domain ID list, perform these steps:

Before you begin

An allowed domain ID list must satisfy the following conditions:

- If this switch is a principal switch, all the currently assigned domain IDs must be in the allowed list.
- If this switch is a subordinate switch, the local runtime domain ID must be in the allowed list.
- The locally configured domain ID of the switch must be in the allowed list.
- The intersection of the assigned domain IDs with other already configured domain ID lists must not be empty.

Procedure

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters configuration mode.
Step 2	switch# **fcdomain allowed 50-110 vsan 4**	Configures the list to allow switches with the domain ID 50 through 110 in VSAN 4. switch# **no fcdomain allowed 50-110 vsan 4** Reverts to the factory default of allowing domain IDs from 1 through 239 in VSAN 5.

Enabling Allowed Domain ID Distribution

CFS distribution of allowed domain ID lists is disabled by default. You must enable distribution on all switches to which you want to distribute the allowed domain ID lists.

To enable (or disable) allowed domain ID list configuration distribution, follow these steps:

Before you begin

All switches in the fabric must be running Cisco SAN-OS Release 3.0(1) or later to distribute the allowed domain ID list using CFS.

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain distribute**

Enables domain configuration distribution.

Step 3 switch(config)# **no fcdomain distribute**

Disables (default) domain configuration distribution.

Committing Changes

To commit pending domain configuration changes and release the lock, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain commit vsan 10**

Commits the pending domain configuration changes.

Discarding Changes

At any time, you can discard the pending changes to the domain configuration and release the fabric lock. If you discard (terminate) the pending changes, the configuration remains unaffected and the lock is released.

To discard pending domain configuration changes and release the lock, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain abort vsan 10**

Discards the pending domain configuration changes.

Enabling Contiguous Domain ID Assignments

To enable contiguous domains in a specific VSAN (or a range of VSANs), follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain contiguous-allocation vsan 81-83**

Enables the contiguous allocation option in VSAN 81 through 83.

Note The **contiguous-allocation** option takes immediate effect at runtime. You do not need to restart the fcdomain.

Step 3 switch(config)# **no fcdomain contiguous-allocation vsan 1030**

Disables the contiguous allocation option and reverts it to the factory default in VSAN 1030.

Configuring FC IDs

When an N or NL port logs into a Cisco MDS 9000 Family switch, it is assigned an FC ID.

Enabling the Persistent FC ID Feature

If you connect to the switch from an AIX or HP-UX host, be sure to enable the persistent FC ID feature in the VSAN that connects these hosts.

A persistent FC ID assigned to an F port can be moved across interfaces and can continue to maintain the same persistent FC ID.

Note

- FC IDs are enabled by default. This change of default behavior from releases prior to Cisco MDS SAN-OS Release 2.0(1b) prevents FC IDs from being changed after a reboot. You can disable this option for each VSAN.
- Persistent FC IDs with loop-attached devices (FL ports) must remain connected to the same port in which they were configured.
- Due to differences in Arbitrated Loop Physical Address (ALPA) support on devices, FC ID persistency for loop-attached devices is not guaranteed.
- For Cisco MDS 9124, 9134, 9148, 9148S, and 9250i switches, ensure that you allocate a complete FCID area per interface and that the last byte to the right of the FCID (port_id) is always zero for these platforms (except for an MDS 9148 running in the NPIV mode connected to an NPV switch). Hence, you cannot configure static FCIDs with non-zero port_ids. For example, the following will not work on MDS 9124, 9134, 9148, 9148S, and 9250i:

  vsan 1000 **wwn** 33:e8:00:05:30:00:16:df **fcid** 0x070128

  It should be changed to **vsan** 1000 **wwn** 33:e8:00:05:30:00:16:df **fcid** 0x070100.

To enable the persistent FC ID feature, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain fcid persistent vsan** *1000*

```
FCID(s) persistent feature is enabled.
```

Activates (default) persistency of FC IDs in VSAN 1000.

Step 3 switch(config)# **no fcdomain fcid persistent vsan** *20*

Disables the FC ID persistency feature in VSAN 20.

Configuring Persistent FC IDs

To configure persistent FC IDs, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcdomain fcid database**

switch(config-fcid-db)#

Enters FC ID database configuration submode.

Step 3 switch(config-fcid-db)# **vsan 1000 wwn 33:e8:00:05:30:00:16:df fcid 0x070128**

Configures a device WWN (33:e8:00:05:30:00:16:df) with the FC ID 0x070128 in VSAN 1000.

Note To avoid assigning a duplicate FC ID, use the **show fcdomain address-allocation vsan** command to display the FC IDs in use.

Step 4 switch(config-fcid-db)# **vsan 1000 wwn 11:22:11:22:33:44:33:44 fcid 0x070123 dynamic**

Configures a device WWN (11:22:11:22:33:44:33:44) with the FC ID 0x070123 in VSAN 1000 in dynamic mode.

Step 5 switch(config-fcid-db)# **vsan 1000 wwn 11:22:11:22:33:44:33:44 fcid 0x070100 area**

Configures a device WWN (11:22:11:22:33:44:33:44) with the FC IDs 0x070100 through 0x701FF in VSAN 1000.

Note To secure the entire area for this fcdomain, assign 00 as the last two characters of the FC ID.

Configuring Unique Area FC IDs for an HBA

To configure a different area ID for the HBA port, follow these steps:

Note The procedure in this example uses a switch domain of 111(6f hex). The HBA port connects to interface fc1/9 and the storage port connects to interface fc 1/10 in the same switch.

Procedure

Step 1 Obtain the port WWN (Port Name field) ID of the HBA using the **show flogi database** command.

```
switch# show flogi database
-------------------------------------------------------------------------------------------
INTERFACE   VSAN    FCID        PORT NAME                   NODE NAME
-------------------------------------------------------------------------------------------
 fc1/9      3       0x6f7703    50:05:08:b2:00:71:c8:c2     50:05:08:b2:00:71:c8:c0
 fc1/10     3       0x6f7704    50:06:0e:80:03:29:61:0f     50:06:0e:80:03:29:61:0f
```

Note Both FC IDs in this setup have the same area 77 assignment.

Step 2 Shut down the HBA interface in the MDS switch.

```
switch# configure terminal
switch(config)# interface fc1/9
switch(config-if)# shutdown
switch(config-if)# end
switch#
```

Example:

Step 3 Verify that the FC ID feature is enabled using the **show fcdomain vsan** command.

```
switch# show fcdomain vsan 1
Local switch run time information:
        State: Stable
        Local switch WWN:    20:01:54:7f:ee:de:b3:01
        Running fabric name: 20:01:00:05:9b:2c:1c:71
        Running priority: 128
        Current domain ID: 0xee(238)
Local switch configuration information:
        State: Enabled
        FCID persistence: Disabled
        Auto-reconfiguration: Disabled
        Contiguous-allocation: Disabled
        Configured fabric name: 20:01:00:05:30:00:28:df
        Optimize Mode: Disabled
        Configured priority: 128
        Configured domain ID: 0x00(0) (preferred)
Principal switch run time information:
        Running priority: 2
Interface               Role          RCF-reject
----------------    -------------    ------------
fc1/1               Non-principal    Disabled
fc1/2               Upstream         Disabled
fc1/11              Non-principal    Disabled
fc1/37              Non-principal    Disabled
port-channel 1      Downstream       Disabled
```

```
----------------    -------------    ------------
```

If this feature is disabled, continue with this procedure to enable the persistent FC ID.

If this feature is already enabled, skip to step 7.

Step 4 Enable the persistent FC ID feature in the Cisco MDS switch.

```
switch# configure terminal
switch(config)# fcdomain fcid persistent vsan 1
switch(config)# end
switch#
```

Step 5 Assign a new FC ID with a different area allocation. In this example, we replace *77* with *ee*.

```
switch# configure terminal
switch(config)# fcdomain fcid database
switch(config-fcid-db)# vsan 3 wwn 50:05:08:b2:00:71:c8:c2 fcid 0x6fee00 area
```

Step 6 Enable the HBA interface in the Cisco MDS switch.

```
switch# configure terminal
switch(config)# interface fc1/9
switch(config-if)# no shutdown
switch(config-if)# end
switch#
```

Step 7 Verify the pWWN ID of the HBA using the **show flogi database** command.

```
switch# show flogi database
------------------------------------------------------------------------------------------
INTERFACE   VSAN    FCID        PORT NAME                  NODE NAME
------------------------------------------------------------------------------------------
fc1/9       3       0x6fee00    50:05:08:b2:00:71:c8:c2    50:05:08:b2:00:71:c8:c0
fc1/10      3       0x6f7704    50:06:0e:80:03:29:61:0f    50:06:0e:80:03:29:61:0f
```

Note Both FC IDs now have different area assignments.

Purging Persistent FC IDs

To purge persistent FC IDs, follow this step:

Procedure

Step 1 switch# **purge fcdomain fcid vsan 4**

Purges all dynamic and unused FC IDs in VSAN 4.

Step 2 switch# **purge fcdomain fcid vsan 3-5**

Purges dynamic and unused FC IDs in VSAN 3, 4, and 5.

Clearing a Fabric Lock

To release a fabric lock, issue the **clear fcdomain session vsan** command in EXEC mode using a login ID that has administrative privileges.

```
switch# clear fcdomain session vsan 10
```

Verifying FC Domain Configuration

To display the domain ID configuration information, perform the following tasks:

Command	Purpose
show fcdomain status	Displays the status of CFS distribution for allowed domain ID lists.
show fcdomain pending	Displays the pending configuration changes.
show fcdomain session-status vsan	Displays the status of the distribution session.
show fcdomain	Displays global information about fcdomain configurations.
show fcdomain domain-list	Displays the list of domain IDs of all switches.
show fcdomain allowed vsan	Displays the list of allowed domain IDs configured on this switch.
show fcdomain fcid persistent	Displays all existing, persistent FC IDs for a specified VSAN.
show fcdomain statistics	Displays frame and other fcdomain statistics for a specified VSAN or PortChannel.
show fcdomain address-allocation	Displays FC ID allocation statistics including a list of assigned and free FC IDs.
show fcdomain address-allocation cache	Displays the valid address allocation cache.

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference* .

Displaying CFS Distribution Status

You can display the status of CFS distribution for allowed domain ID lists using the **show fcdomain status** command.

```
switch# show fcdomain status
CFS distribution is enabled
```

Displaying Pending Changes

You can display the pending configuration changes using the show fcdomain pending command:

```
switch# show fcdomain pending vsan 10
Pending Configured Allowed Domains
----------------------------------
VSAN 10
Assigned or unallowed domain IDs: 1-9,24,100,231-239.
[User] configured allowed domain IDs: 10-230.
```

You can display the differences between the pending configuration and the current configuration using the show fcdomain pending-diff command.

```
switch#show fcdomain pending-diff vsan 10
Current Configured Allowed Domains
----------------------------------
VSAN 10
Assigned or unallowed domain IDs: 24,100.
[User] configured allowed domain IDs: 1-239.
Pending Configured Allowed Domains
----------------------------------
VSAN 10
Assigned or unallowed domain IDs: 1-9,24,100,231-239.
[User] configured allowed domain IDs: 10-230.
```

Displaying Session Status

You can display the status of the distribution session using the **show fcdomain session-status vsan** command.

```
switch# show fcdomain session-status vsan 1
Last Action: Distribution Enable
Result: Success
```

Displaying fcdomain Information

Global fcdoman Information

Use the **show fcdomain** command to display global information about fcdomain configurations. See the following example:

Note In the following example, the fcdomain feature is disabled. Consequently, the runtime fabric name is the same as the configured fabric name.

```
switch# show fcdomain vsan 2
The local switch is the Principal Switch.
Local switch run time information:
        State: Stable
        Local switch WWN:    20:01:00:0b:46:79:ef:41
        Running fabric name: 20:01:00:0b:46:79:ef:41
        Running priority: 128
        Current domain ID: 0xed(237)
Local switch configuration information:
        State: Enabled
        FCID persistence: Disabled
        Auto-reconfiguration: Disabled
        Contiguous-allocation: Disabled
        Configured fabric name: 20:01:00:05:30:00:28:df
        Optimize Mode: Disabled
        Configured priority: 128
        Configured domain ID: 0x00(0) (preferred)
Principal switch run time information:
        Running priority: 128
No interfaces available.
switch# show fcdomain vsan 1
The local switch is the Principal Switch.
Local switch run time information:
State: Stable
Local switch WWN: 20:01:54:7f:ee:46:5b:41
Running fabric name: 20:01:54:7f:ee:46:5b:41
Running priority: 128
Current domain ID: 0xe9(233)
Local switch configuration information:
State: Enabled
FCID persistence: Enabled
Auto-reconfiguration: Disabled
Contiguous-allocation: Disabled
Configured fabric name: 20:01:00:05:30:00:28:df
Optimize Mode: Enabled (Fast Restart, Selective Restart, Scale Restart)
Configured priority: 128
Configured domain ID: 0xe9(233) (static)
Principal switch run time information:
Running priority: 128
No interfaces available.
switch#
```

Note If a scale-restart feature was enabled and the other optimization modes were disabled when downgrading from Cisco MDS 6.2(9) release or later to 6.2(7) or older releases, the optimize mode will be a blank instead of disabled

fcdomain Lists

Use the **show fcdomain domain-list** command to display the list of domain IDs of all switches belonging to a specified VSAN. This list provides the WWN of the switches owning each domain ID. The following example shows the following:

- A switch with WWN of 20:01:00:05:30:00:47:df is the principal switch and has domain 200.
- A switch with WWN of 20:01:00:0d:ec:08:60:c1 is the local switch (the one where you typed the CLI command to show the domain-list) and has domain 99.
- The IVR manager obtained virtual domain 97 using 20:01:00:05:30:00:47:df as the WWN for a virtual switch.

```
switch# show fcdomain domain-list vsan 76
Number of domains: 3
Domain ID              WWN
---------    -----------------------
0xc8(200)    20:01:00:05:30:00:47:df [Principal]
 0x63(99)    20:01:00:0d:ec:08:60:c1 [Local]
 0x61(97)    50:00:53:0f:ff:f0:10:06 [Virtual (IVR)]
```

Allowed Domain ID Lists

Use the **show fcdomain allowed vsan** command to display the list of allowed domain IDs configured on this switch. See the followng example:

```
switch# show fcdomain allowed vsan 1

Assigned or unallowed domain IDs: 1-96,100,111-239.
[Interoperability Mode 1] allowed domain IDs: 97-127.
[User] configured allowed domain IDs: 50-110.
```

Tip Ensure that the requested domain ID passes the Cisco NX-OS software checks, if **interop 1** mode is required in this switch.

Persistent FC IDs in a Specified VSAN

Use the **show fcdomain fcid persistent** command to display all existing, persistent FC IDs for a specified VSAN. You can also specify the **unused** option to view only persistent FC IDs that are still not in use. See the following examples:

```
switch# show fcdomain fcid persistent vsan 1000
Total entries 2.
Persistent FCIDs table contents:
VSAN              WWN                FCID         Mask        Used    Assignment
----    -----------------------  --------      -----------    ----    ----------
1000    11:11:22:22:11:11:12:23    0x700101    SINGLE FCID      NO    STATIC
1000    44:44:33:33:22:22:11:11    0x701000    ENTIRE AREA      NO    DYNAMIC
```

All Persisent FC IDs in fcdomain

The following example displays all persistent FC IDs in the fcdomain:

```
switch# show fcdomain fcid persistent
Total entries 2.
Persistent FCIDs table contents:
VSAN              WWN                FCID         Mask        Used    Assignment
----    -----------------------  --------      -----------    ----    ----------
1000    11:11:22:22:11:11:22:22    0x700501    SINGLE FCID      NO    STATIC
1003    44:44:33:33:22:22:11:11    0x781000    ENTIRE AREA     YES    DYNAMIC
```

fcdomain Statistics for a Specified VSAN

Use the **show fcdomain statistics** command to display frame and other fcdomain statistics for a specified VSAN or PortChannel. See the following example and Domain Manager Selective Restart, on page 265.

```
switch# show fcdomain statistics vsan1

VSAN Statistics
```

```
        Number of Principal Switch Selections: 5
        Number of times Local Switch was Principal: 0
        Number of 'Build Fabric's: 3
        Number of 'Fabric Reconfigurations': 0
```

fcdomain Statistics for a Specified PortChannel

The following example displays fcdomain statistics for a specified PortChannel:

```
switch# show fcdomain statistics interface port-channel 10 vsan 1

        Interface Statistics:
                Transmitted      Received
                -----------      --------
                EFPs      13       9
                DIAs      7        7
                RDIs      0        0
                ACCs      21       25
                RJTs      1        1
                BFs       2        2
                RCFs      4        4
                Error     0        0
                Total     48       48
        Total Retries: 0
        Total Frames: 96
                -----------      --------
```

FC ID Information

Use the **show fcdomain address-allocation** command to display FC ID allocation statistics including a list of assigned and free FC IDs. See the following example:

```
switch# show fcdomain address-allocation vsan 1
Free FCIDs: 0x020000 to 0x02fdff
            0x02ff00 to 0x02fffe

Assigned FCIDs: 0x02fe00 to 0x02feff
                0x02ffff

Reserved FCIDs: 0x020100 to 0x02f0ff
                0x02fe00 to 0x02feff
                0x02ffff


Number free FCIDs: 65279
Number assigned FCIDs: 257
Number reserved FCIDs: 61697
```

Address Allocation Information

Use the **show fcdomain address-allocation cache** command to display the valid address allocation cache. The cache is used by the principal switch to reassign the FC IDs for a device (disk or host) that exited and reentered the fabric. In the cache content, VSAN refers to the VSAN that contains the device, WWN refers to the device that owned the FC IDs, and mask refers to a single or entire area of FC IDs. See the following example:

```
switch# show fcdomain address-allocation cache
Cache content:
line#    VSAN              WWN                FCID         mask
-----    ----    -----------------------    --------    -----------
   1.      12    21:00:00:e0:8b:08:a2:21    0xef0400    ENTIRE AREA
   2.       6    50:06:04:82:c3:a1:2f:5c    0xef0002    SINGLE FCID
   3.       8    20:4e:00:05:30:00:24:5e    0xef0300    ENTIRE AREA
   4.       8    50:06:04:82:c3:a1:2f:52    0xef0001    SINGLE FCID
```

CHAPTER 16

Monitoring Network Traffic Using SPAN

This chapter describes the Switched Port Analyzer (SPAN) features provided in switches in the Cisco MDS 9000 Family.

Information About SPAN

The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic through a Fibre Channel interface. Traffic through any Fibre Channel interface can be replicated to a special port called the SPAN destination port (SD port). Any Fibre Channel port in a switch can be configured as an SD port. Once an interface is in SD port mode, it cannot be used for normal data traffic. You can attach a Fibre Channel Analyzer to the SD port to monitor SPAN traffic.

SD ports do not receive frames, they only transmit a copy of the SPAN source traffic. The SPAN feature is nonintrusive and does not affect switching of network traffic for any SPAN source ports (see Figure 15: SPAN Transmission, on page 291).

Figure 15: SPAN Transmission

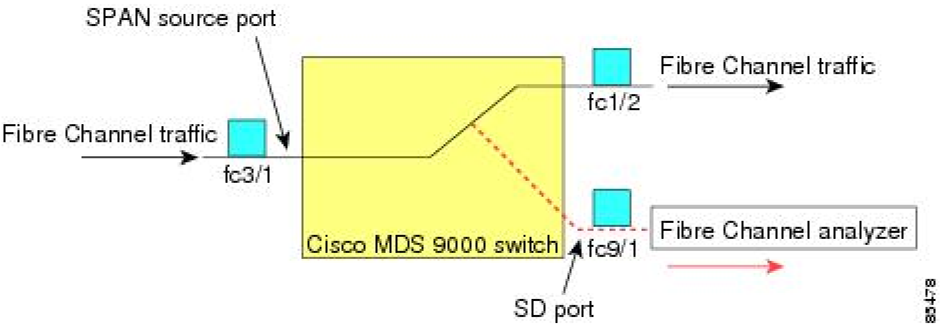

SPAN Sources

SPAN sources refer to the interfaces from which traffic can be monitored. You can also specify VSAN as a SPAN source, in which case, all supported interfaces in the specified VSAN are included as SPAN sources. When a VSAN as a source is specified, then all physical ports and PortChannels in that VSAN are included as SPAN sources. You can choose the SPAN traffic in the ingress direction, the egress direction, or both directions for any source interface:

- Ingress source (Rx)—Traffic entering the switch fabric through this source interface is *spanned* or copied to the SD port (see Figure 16: SPAN Traffic from the Ingress Direction, on page 292).

Figure 16: SPAN Traffic from the Ingress Direction

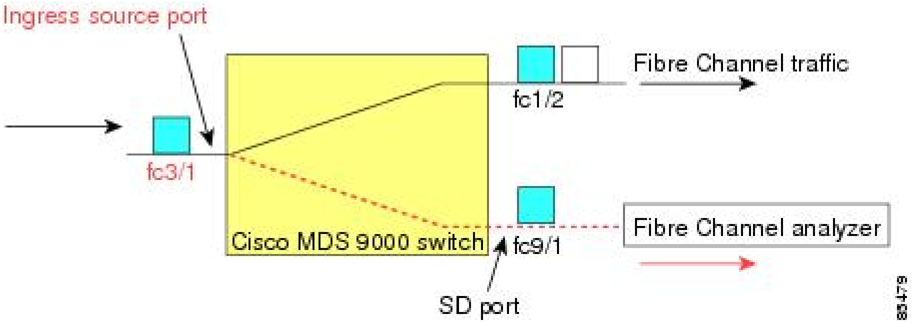

- Egress source (Tx)—Traffic exiting the switch fabric through this source interface is spanned or copied to the SD port (see Figure 17: SPAN Traffic from Egress Direction, on page 292).

Figure 17: SPAN Traffic from Egress Direction

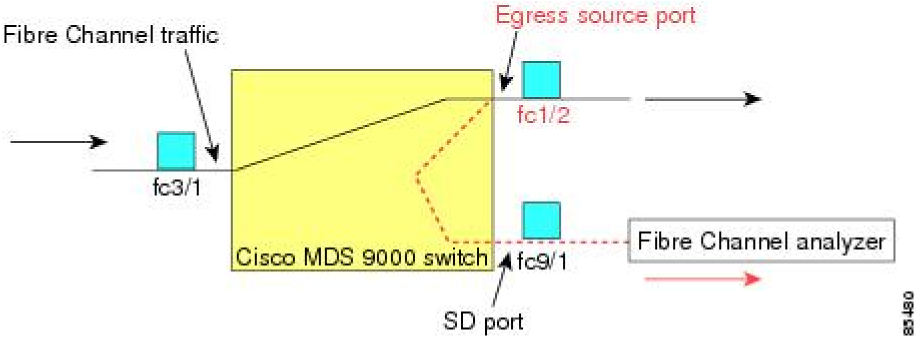

IPS Source Ports

SPAN capabilities are available on the FCIP and iSCSI interfaces on the IP Storage Services (IPS) ports. The SPAN feature is only implemented on the FCIP and iSCSI virtual Fibre Channel interfaces, not IPS ports themselves. You can configure SPAN for ingress traffic, egress traffic, or traffic in both directions for any one of the 24 FCIP interfaces that are available in the IPS module.

Note

- You can configure SPAN for Ethernet traffic using Cisco switches or routers that are connected to the Cisco MDS 9000 Series IPS modules.
- Cisco MDS 9200i switch does not support iSCSI.

From Cisco MDS NX-OS Release 8.5(1), traffic that is sent to a Fibre Channel port that is configured as an SD port can be spanned from FCIP interfaces.

The following are the restrictions for using SPAN for ingress or egress traffic that is sent to a Fibre Channel port that is configured as an SD port can be spanned from FCIP interfaces:

- Only one FCIP interface can be added as the ingress SPAN source.
- FCIP port channel cannot be added as the ingress SPAN source. However, individual FCIP member links can be added as the ingress SPAN source.
- Either the ingress or egress SPAN source can be added in a SPAN session, but not bidirectional. To do a bidirectional SPAN, configure two SPAN sessions, one for ingress and the other for egress, to the same destination SD ports.
- You cannot configure Fibre Channel and FCIP interfaces together as ingress or egress source.

Allowed Source Interface Types

The SPAN feature is available for the following interface types:

- Physical ports such as F ports, FL ports, TE ports, E ports, and TL ports.
- Interface sup-fc0 (traffic to and from the supervisor):
  - The Fibre Channel traffic from the supervisor module to the switch fabric through the sup-fc0 interface is called ingress traffic. It is spanned when sup-fc0 is chosen as an ingress source port.
  - The Fibre Channel traffic from the switch fabric to the supervisor module through the sup-fc0 interface is called egress traffic. It is spanned when sup-fc0 is chosen as an egress source port.
- PortChannels
  - All ports in the PortChannel are included and spanned as sources.
  - You cannot specify individual ports in a PortChannel as SPAN sources. Previously configured SPAN-specific interface information is discarded.
- IPS module specific Fibre Channel interfaces:
  - iSCSI interfaces
  - FCIP interfaces

Note In Cisco MDS 9700 Series Switches, iSCSI ports are not applicable for the Allowed Source Interface Types.

VSAN as a Source

SPAN sources refer to the interfaces from which traffic can be monitored. When a VSAN as a source is specified, then all physical ports and PortChannels in that VSAN are included as SPAN sources. A TE port is included only when the port VSAN of the TE port matches the source VSAN. A TE port is excluded even if the configured allowed VSAN list may have the source VSAN, but the port VSAN is different.

You cannot configure source interfaces (physical interfaces, PortChannels, or sup-fc interfaces) and source VSANs in the same SPAN session.

SPAN Sessions

Each SPAN session represents an association of one destination with a set of source(s) along with various other parameters that you specify to monitor the network traffic. One destination can be used by one or more SPAN sessions. You can configure up to 16 SPAN sessions in a switch. Each session can have several source ports and one destination port.

To activate any SPAN session, at least one source and the SD port must be up and functioning. Otherwise, traffic is not directed to the SD port.

Tip A source can be shared by two sessions, however, each session must be in a different direction—one ingress and one egress.

You can temporarily deactivate (suspend) any SPAN session. The traffic monitoring is stopped during this time.

Note On a Cisco MDS 9250i Multiservice Fabric switch, packet drops will occur if the SPAN port cannot keep up with incoming frame bursts. To avoid these packet drops, the speed of the SPAN destination port should be equal to the maximum speed of the source ports. However, when the source is an FCIP interface, the speed of the SPAN destination port should be more than 10G because the FCIP interface is running over a 10G Ethernet physical interface.

Specifying Filters

You can perform VSAN-based filtering to selectively monitor network traffic on specified VSANs. You can apply this VSAN filter to all sources in a session (see). Only VSANs present in the filter are spanned.

You can specify session VSAN filters that are applied to all sources in the specified session. These filters are bidirectional and apply to all sources configured in the session. Each SPAN session represents an association of one destination with a set of source(s) along with various other parameters that you specify to monitor the network traffic.

SD Port Characteristics

An SD port has the following characteristics:

- Ignores BB_credits.

- Allows data traffic only in the egress (Tx) direction.
- Does not require a device or an analyzer to be physically connected.
- Supports only 1 Gbps or 2 Gbps speeds. The auto speed option is not allowed.
- Multiple sessions can share the same destination ports.
- If the SD port is shut down, all shared sessions stop generating SPAN traffic.
- The outgoing frames can be encapsulated in Extended Inter-Switch Link (EISL) format.
- The SD port does not have a port VSAN.
- SD ports cannot be configured using Storage Services Modules (SSMs).
- The port mode cannot be changed if it is being used for a SPAN session.

Note

- If you need to change an SD port mode to another port mode, first remove the SD port from all sessions and then change the port mode using the **switchport mode** command.
- In Cisco MDS 9700 Series Switches, the SD Port supports only 2 Gbps, 4 Gbps, 8 Gbps, and 16 Gbps speeds. The auto speed option is not allowed

SPAN Conversion Behavior

SPAN features (configured in any prior release) are converted as follows:

- If source interfaces and source VSANs are configured in a given session, then all the source VSANs are removed from that session.

For example, before Cisco MDS SAN-OS Release 1.0(4):

```
Session 1 (active)
   Destination is fc1/9
   No session filters configured
   Ingress (rx) sources are
      vsans 10-11
     fc1/3,
   Egress (tx) sources are
     fc1/3,
```

Once upgraded to Cisco MDS SAN-OS Release 1.1(1):

```
Session 1 (active)
   Destination is fc1/9
   No session filters configured
   Ingress (rx) sources are
     fc1/3,
   Egress (tx) sources are
     fc1/3,
```

For Cisco MDS 9700 Series Switches:

```
switch(config-if)# monitor session 1
switch(config-monitor)# source interface fc5/1
switch(config-monitor)# destination interface fc2/9
switch(config-monitor)# no shut
switch(config-monitor)# show monitor session all
session 1
---------------
ssn direction : both
state : up
source intf :
rx : fc5/1
tx : fc5/1
both : fc5/1
source VLANs :
rx :
tx :
both :
source exception :
rate-limit : Auto
filter VLANs : filter not specified
destination ports : fc2/9
```

Session 1 had both source interfaces and source VSANs before the upgrade. After the upgrade, the source VSANs were removed (rule 1).

- If interface level VSAN filters are configured in source interfaces, then the source interfaces are also removed from the session. If this interface is configured in both directions, it is removed from both directions.

For example, before Cisco MDS SAN-OS Release 1.0(4):

```
Session 2 (active)
   Destination is fc1/9
   No session filters configured
   Ingress (rx) sources are
      vsans 12
     fc1/6 (vsan 1-20),
   Egress (tx) sources are
     fc1/6 (vsan 1-20),
```

Once upgraded to Cisco MDS SAN-OS Release 1.1(1):

```
Session 2 (inactive as no active sources)
   Destination is fc1/9
   No session filters configured
   No ingress (rx) sources
   No egress (tx) sources
```

Note The deprecated configurations are removed from persistent memory once a switchover or a new startup configuration is implemented.

Session 2 had a source VSAN 12 and a source interface fc1/6 with VSAN filters specified in Cisco MDS SAN-OS Release 1.0(4). When upgraded to Cisco MDS SAN-OS Release 1.1(1) the following changes are made:

- • The source VSAN (VSAN 12) is removed (rule 1).
  • The source interface fc1/6 had VSAN filters specified—it is also removed (rule 2).

Monitoring Traffic Using Fibre Channel Analyzers

You can use SPAN to monitor traffic on an interface without any traffic disruption. This feature is especially useful in troubleshooting scenarios in which traffic disruption changes the problem environment and makes it difficult to reproduce the problem. You can monitor traffic in either of the following two ways:

- Without SPAN
- With SPAN

Monitoring Without SPAN

You can monitor traffic using interface fc1/1 in a Cisco MDS 9000 Family switch that is connected to another switch or host. You need to physically connect a Fibre Channel analyzer between the switch and the storage device to analyze the traffic through interface fc1/1 (see Figure 18: Fibre Channel Analyzer Usage Without SPAN, on page 297).

Figure 18: Fibre Channel Analyzer Usage Without SPAN

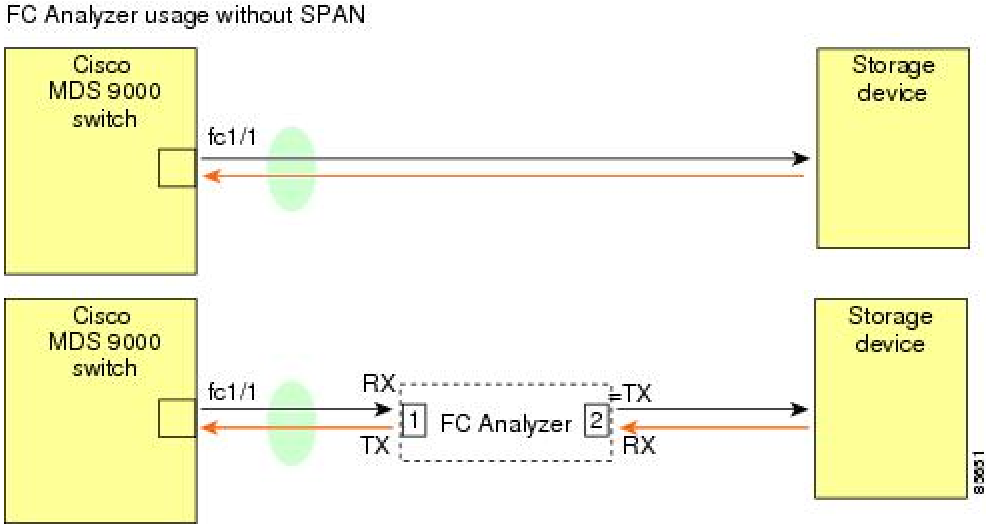

This type of connection has the following limitations:

- It requires you to physically insert the FC analyzer between the two network devices.
- It disrupts traffic when the Fibre Channel analyzer is physically connected.
- The analyzer captures data only on the Rx links in both port 1 and port 2. Port 1 captures traffic exiting interface fc1/1 and port 2 captures ingress traffic into interface fc1/1.

Monitoring with SPAN

Using SPAN you can capture the same traffic scenario (see Figure 18: Fibre Channel Analyzer Usage Without SPAN, on page 297) without any traffic disruption. The Fibre Channel analyzer uses the ingress (Rx) link at port 1 to capture all the frames going out of the interface fc1/1. It uses the ingress link at port 2 to capture all the ingress traffic on interface fc1/1.

Using SPAN you can monitor ingress traffic on fc1/1 at SD port fc2/2 and egress traffic on SD port fc2/1. This traffic is seamlessly captured by the FC analyzer (see Figure 19: Fibre Channel Analyzer Using SPAN, on page 298).

Figure 19: Fibre Channel Analyzer Using SPAN

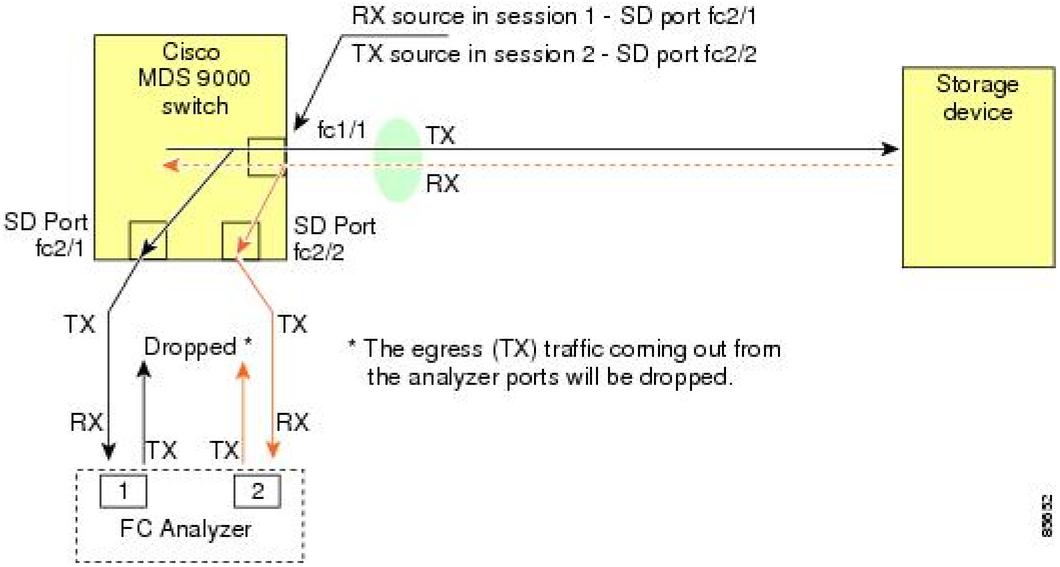
RX source in session 1 - SD port fc2/1
TX source in session 2 - SD port fc2/2
Cisco MDS 9000 switch
fc1/1
TX
RX
Storage device
SD Port fc2/1
SD Port fc2/2
TX
Dropped *
TX
* The egress (TX) traffic coming out from the analyzer ports will be dropped.
RX
RX
TX
TX
1
2
FC Analyzer
85652

Single SD Port to Monitor Traffic

You do not need to use two SD ports to monitor bidirectional traffic on any interface (see Figure 19: Fibre Channel Analyzer Using SPAN, on page 298). You can use one SD port and one FC analyzer port by monitoring traffic on the interface at the same SD port fc2/1.

Figure 20: Fibre Channel Analyzer Using a Single SD Port, on page 298 shows a SPAN setup where one session with destination port fc2/1 and source interface fc1/1 is used to capture traffic in both ingress and egress directions. This setup is more advantageous and cost effective than the setup shown in Figure 19: Fibre Channel Analyzer Using SPAN, on page 298. It uses one SD port and one port on the analyzer, instead of using a full, two-port analyzer.

Figure 20: Fibre Channel Analyzer Using a Single SD Port

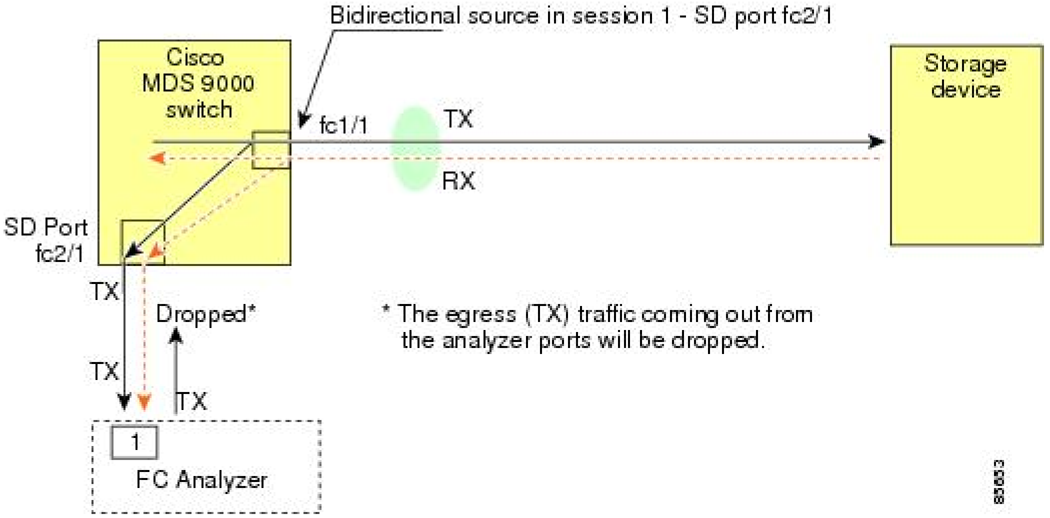
Bidirectional source in session 1 - SD port fc2/1
Cisco MDS 9000 switch
fc1/1
TX
RX
Storage device
SD Port fc2/1
TX
Dropped*
TX
TX
* The egress (TX) traffic coming out from the analyzer ports will be dropped.
1
FC Analyzer
85653

To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

SD Port Configuration

The SD port in the destination switch enables the FC analyzer to receive the RSPAN traffic from the Fibre Channel tunnel. Figure 21: RSPAN Tunnel Configuration, on page 299 depicts an RSPAN tunnel configuration, now that tunnel destination is also configured.

Figure 21: RSPAN Tunnel Configuration

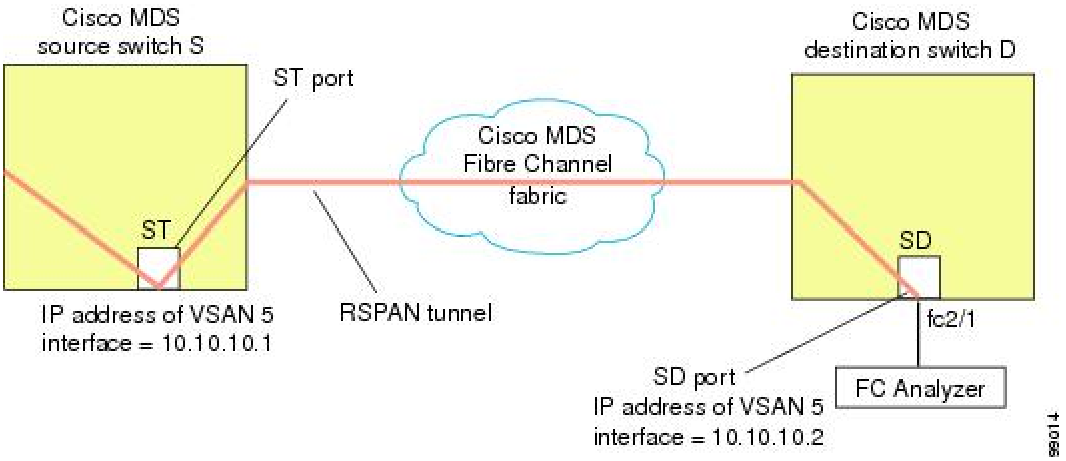

Note SD ports cannot be configured using Storage Services Modules (SSMs).

Mapping the FC Tunnel

The **tunnel-id-map** option specifies the egress interface of the tunnel at the destination switch (see Figure 22: FC Tunnel Configuration, on page 299).

Figure 22: FC Tunnel Configuration

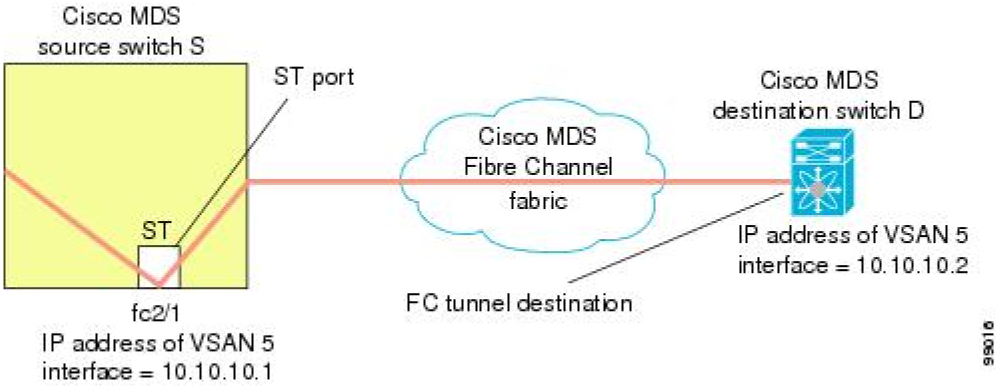

Creating VSAN Interfaces

Figure 23: FC Tunnel Configuration, on page 300 depicts a basic FC tunnel configuration.

Figure 23: FC Tunnel Configuration

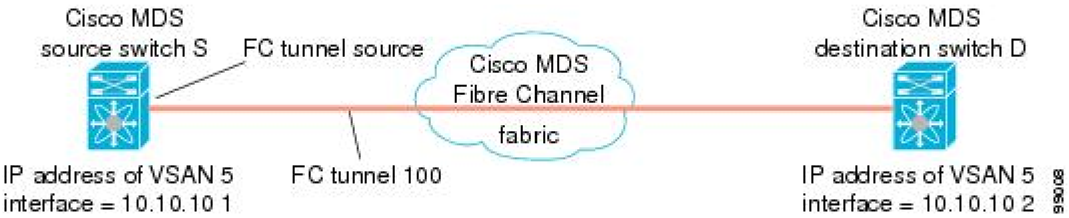

Note This example assumes that VSAN 5 is already configured in the VSAN database.

Remote SPAN

Note Remote SPAN is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem, Cisco Fabric Switch for IBM BladeSystem, Cisco Fabric Switch 9250i, and Cisco Fabric Switch 9100S.

The Remote SPAN (RSPAN) feature enables you to remotely monitor traffic for one or more SPAN sources distributed in one or more source switches in a Fibre Channel fabric. The SPAN destination (SD) port is used for remote monitoring in a destination switch. A destination switch is usually different from the source switch(es) but is attached to the same Fibre Channel fabric. You can replicate and monitor traffic in any remote Cisco MDS 9000 Family switch or director, just as you would monitor traffic in a Cisco MDS source switch.

The RSPAN feature is nonintrusive and does not affect network traffic switching for those SPAN source ports. Traffic captured on the remote switch is tunneled across a Fibre Channel fabric which has trunking enabled on all switches in the path from the source switch to the destination switch. The Fibre Channel tunnel is structured using trunked ISL (TE) ports. In addition to TE ports, the RSPAN feature uses two other interface types (see Figure 24: RSPAN Transmission, on page 301):

- SD ports—A passive port from which remote SPAN traffic can be obtained by the FC analyzer.
- ST ports—A SPAN tunnel (ST) port is an entry point port in the source switch for the RSPAN Fibre Channel tunnel. ST ports are special RSPAN ports and cannot be used for normal Fibre Channel traffic.

Figure 24: RSPAN Transmission

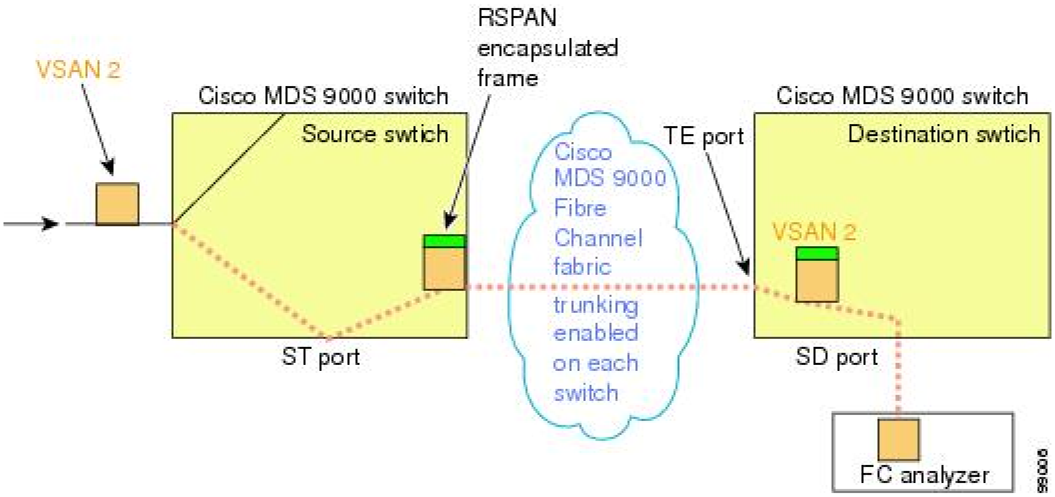

Advantages of Using RSPAN

The RSPAN features has the following advantages:

- Enables nondisruptive traffic monitoring at a remote location.
- Provides a cost effective solution by using one SD port to monitor remote traffic on multiple switches.
- Works with any Fibre Channel analyzer.
- Is compatible with the Cisco MDS 9000 Port Analyzer adapters.
- Does not affect traffic in the source switch, but shares the ISL bandwidth with other ports in the fabric.

FC and RSPAN Tunnels

An FC tunnel is a logical data path between a source switch and a destination switch. The FC tunnel originates from the source switch and terminates at the remotely located destination switch.

RSPAN uses a special Fibre Channel tunnel (FC tunnel) that originates at the ST port in the source switch and terminates at the SD port in the destination switch. You must bind the FC tunnel to an ST port in the source switch and map the same FC tunnel to an SD port in the destination switch. Once the mapping and binding is configured, the FC tunnel is referred to as an RSPAN tunnel (see Figure 25: FC and RSPAN Tunnel, on page 302).

Figure 25: FC and RSPAN Tunnel

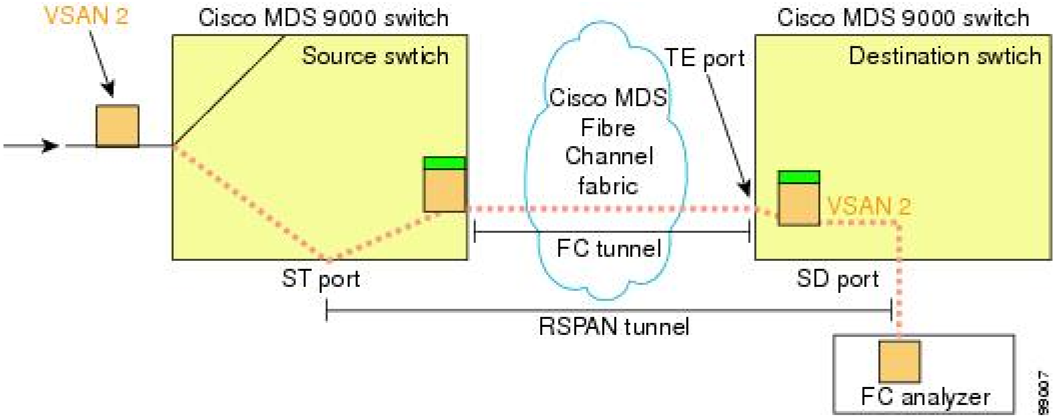

ST Port Configuration

Note In Cisco MDS 9700 Series Switches, SPAN tunnel port (ST port) is not supported.

Once the FC tunnel is created, be sure to configure the ST port to bind it to the FC tunnel at the source switch. The FC tunnel becomes an RSPAN tunnel once the binding and mapping is complete.

Figure 26: Binding the FC Tunnel, on page 302 depicts a basic FC tunnel configuration.

Figure 26: Binding the FC Tunnel

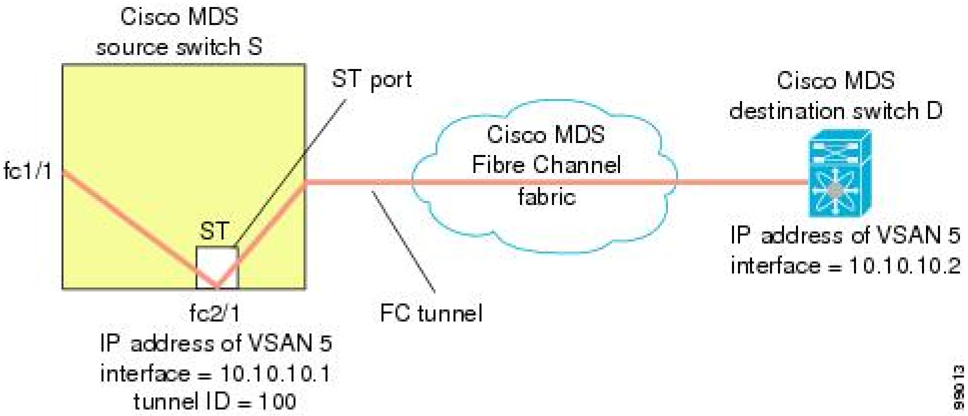

ST Port Characteristics

ST ports have the following characteristics:

- ST ports perform the RSPAN encapsulation of the FC frame.
- ST ports do not use BB_credits.
- One ST port can only be bound to one FC tunnel.
- ST ports cannot be used for any purpose other than to carry RSPAN traffic.

- ST ports cannot be configured using Storage Services Modules (SSMs).

Creating Explicit Paths

You can specify an explicit path through the Cisco MDS Fibre Channel fabric (source-based routing), using the **explicit-path** option. For example, if you have multiple paths to a tunnel destination, you can use this option to specify the FC tunnel to always take one path to the destination switch. The software then uses this specified path even if other paths are available.

This option is especially useful if you prefer to direct the traffic through a certain path although other paths are available. In an RSPAN situation, you can specify the explicit path so the RSPAN traffic does not interfere with the existing user traffic. You can create any number of explicit paths in a switch (see Figure 27: Explicit Path Configuration, on page 303).

Figure 27: Explicit Path Configuration

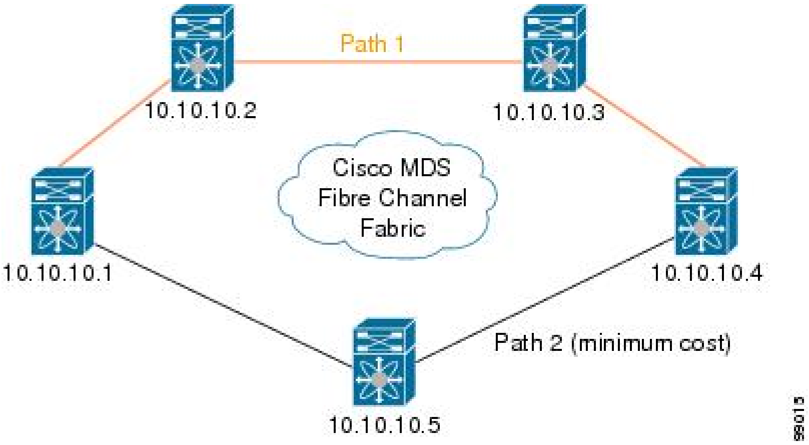

Guidelines and Limitations

Cisco MDS 9700 Series Switches Guidelines

The following guidelines and limitations apply for Cisco MDS 9700 Series Switches:

- In Cisco MDS 9700 Series Switches, SPAN is replaced by Monitor.
- In Cisco MDS 9700 Series Switches, SPAN tunnel port (ST port) is not supported.
- In Cisco MDS 9700 Series Switches, RSPAN is replaced by Remote Monitor.
- For Cisco MDS 9700 Series Switches, Generation 2 Fabric Switches is not supported

SPAN Configuration Guidelines

The following guidelines and limitations apply for SPAN configurations:

- You can configure up to 16 SPAN sessions with multiple ingress (Rx) sources.

- The number of source ports must be less than or equal to 16. However, we recommend that you configure a maximum of only two source ports per SPAN or monitor session.
- You can configure a maximum of three SPAN sessions with one egress (Tx) port.
- In a 32-port switching module, you must configure the same session in all four ports in one port group (unit). If you wish, you can also configure only two or three ports in this unit.

Note This is not applicable for Cisco MDS 9700 Series Switches.

- SPAN frames are dropped if the sum of the bandwidth of the sources exceeds the speed of the destination port.
- Frames dropped by a source port are not spanned.
- SPAN does not capture pause frames in a Fibre Channel over Ethernet (FCoE) network because pause frames sent from the virtual expansion (VE) port are generated and terminated by the outermost MAC layer. For more information on FCoE, see the Cisco NX-OS FCoE Configuration Guide for Cisco Nexus 7000 and Cisco MDS 9500.
- In case of an IVR configuration and topology, SPAN cannot capture the egress (Tx) of the source port. To span the complete traffic flow, add the source ports taking part in the flow in ingress (Rx) direction.

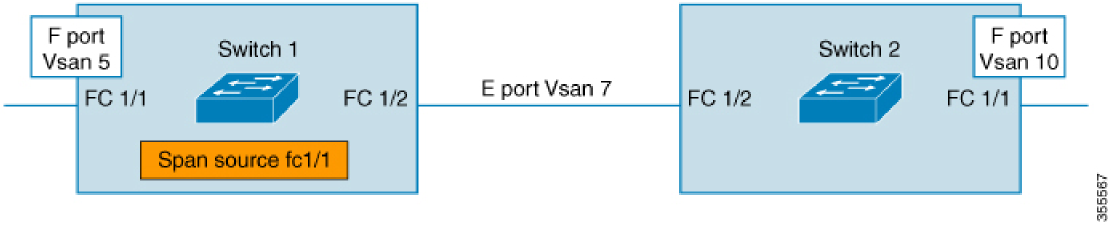

Consider FC1/1, in the above illustration, as the SPAN source port. In this case, traffic egressing (Tx) from FC1/1 will not be spanned. Only packets entering (Rx) FC1/1 will be spanned. To capture the complete flow, span FC1/1 (Rx) and FC1/2 (Rx) in a single session going to a single destination.

Guidelines to Configure VSANs as a Source

The following guidelines apply when configuring VSANs as a source:

- Traffic on all interfaces included in a source VSAN is spanned only in the ingress direction.
- If a VSAN is specified as a source, you cannot perform interface-level SPAN configuration on the interfaces that are included in the VSAN. Previously configured SPAN-specific interface information is discarded.
- If an interface in a VSAN is configured as a source, you cannot configure that VSAN as a source. You must first remove the existing SPAN configurations on such interfaces before configuring VSAN as a source.
- Interfaces are only included as sources when the port VSAN matches the source VSAN. Figure 28: VSAN as a Source, on page 305 displays a configuration using VSAN 2 as a source:

- All ports in the switch are in VSAN 1 except fc1/1.
- Interface fc1/1 is the TE port with port VSAN 2. VSANs 1, 2, and 3 are configured in the allowed list.
- VSAN 1 and VSAN 2 are configured as SPAN sources.

Figure 28: VSAN as a Source

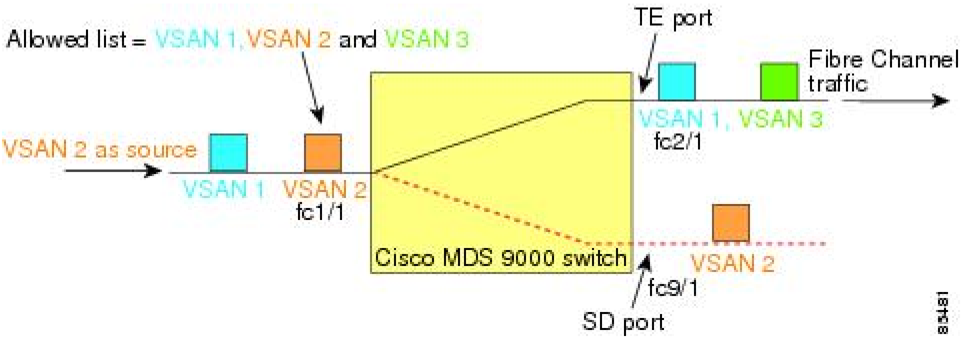

For this configuration, the following apply:

- VSAN 2 as a source includes only the TE port fc1/1 that has port VSAN 2.
- VSAN 1 as a source does not include the TE port fc1/1 because the port VSAN does not match VSAN 1.

Guidelines to Specifying Filters

The following guidelines apply to SPAN filters:

- PortChannel configurations are applied to all ports in the PortChannel.
- If no filters are specified, the traffic from all active VSANs for that interface is spanned by default.
- While you can specify arbitrary VSAN filters in a session, traffic can only be monitored on the port VSAN or on allowed-active VSANs in that interface.

RSPAN Configuration Guidelines

The following guidelines apply for a SPAN configuration:

- All switches in the end-to-end path of the RSPAN tunnel must belong to the Cisco MDS 9000 Family.
- All VSANs with RSPAN traffic must be enabled. If a VSAN containing RSPAN traffic is not enabled, it is dropped.
- The following configurations must be performed on *each* switch in the end-to-end path of the Fibre Channel tunnel in which RSPAN is to be implemented:
  - Trunking must be enabled (enabled by default) and the trunk enabled link must be the lowest cost link in the path.
  - VSAN interface must be configured.
  - The Fibre Channel tunnel feature must be enabled (disabled by default).
  - IP routing must be enabled (disabled by default).

Note If the IP address is in the same subnet as the VSAN, the VSAN interface does not have to be configured for all VSANs on which the traffic is spanned.

- A single Fibre Channel switch port must be dedicated for the ST port functionality.
- Do not configure the port to be monitored as the ST port.
- The FC tunnel's IP address must reside in the same subnet as the VSAN interface.

Default SPAN and RSPAN Settings

Table 39: Default SPAN Configuration Parameters , on page 306 lists the default settings for SPAN parameters.

Table 39: Default SPAN Configuration Parameters

Parameters	Default
SPAN session	Active. **Note** For Cisco MDS 9700 Series Switches, the default value for Monitor session is Shut.
If filters are not specified	SPAN traffic includes traffic through a specific interface from all active VSANs.
Encapsulation	Disabled.
SD port	Output frame format is Fibre Channel.

Table 40: Default RSPAN Configuration Parameters , on page 306 lists the default settings for RSPAN parameters.

Table 40: Default RSPAN Configuration Parameters

Parameters	Default
FC tunnel	Disabled
Explicit path	Not configured
Minimum cost path	Used if explicit path is not configured

Configuring SPAN

The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic through a Fibre Channel interface.

Configuring SD Ports for SPAN

Configuring SD Port for SPAN Monitoring

To configure an SD port for SPAN monitoring, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc9/1**

Configures the specified interface.

Step 3 switch(config-if)# **switchport mode SD**

Configures the SD port mode for interface fc9/1.

Step 4 switch(config-if)# **switchport speed 1000**

Configures the SD port speed to 1000 Mbps.

Note In Cisco MDS 9700 Series Switches, the switch port speed is 8000 Mbps.

Step 5 switch(config-if)# **no shutdown**

Enables traffic flow through this interface.

Configuring SPAN Session

To configure a SPAN session, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **span session 1**

switch(config-span)#

Configures the specified SPAN session (1). If the session does not exist, it is created.

Note In Cisco MDS 9700 Series Switches, SPAN is replaced by Monitor.

Step 3 switch(config)# **no span session 1**

Deletes the specified SPAN session (1).

Step 4 switch(config-span)# **destination interface fc9/1**

Configures the specified destination interface (fc 9/1) in a session.

Step 5 switch(config-span)# **no destination interface fc9/1**

Removes the specified destination interface (fc 9/1).

Step 6 switch(config-span)# **source interface fc7/1**

Configures the source (fc7/1) interface in both directions.

Note While configuring SPAN sources on the Cisco MDS 9124 Fabric Switch, the direction (Rx and Tx) needs to be explicitly mentioned.

Step 7 switch(config-span)# **no source interface fc7/1**

Removes the specified destination interface (fc 7/1) from this session.

Step 8 switch(config-span)# **source interface sup-fc0**

Configures the source interface (sup-fc0) in the session.

Step 9 switch(config-span)# **source interface fc1/5 - 6, fc2/1 -3**

Configures the specified interface ranges in the session.

Step 10 switch(config-span)# **source vsan 1-2**

Configures source VSANs 1 and 2 in the session.

Step 11 switch(config-span)# **source interface port-channel 1**

Configures the source PortChannel (port-channel 1).

Step 12 switch(config-span)# **source interface fcip 51**

Configures the source FCIP interface in the session.

Step 13 switch(config-span)# **source interface iscsi 4/1**

Configures the source iSCSI interface in the session.

Note This is not applicable for MDS 9700 Series Switches.

Step 14 switch(config-span)# **source interface svc1/1 tx traffic-type initiator**

Configures the source SVC interface in the Tx direction for an initiator traffic type.

Note This is not applicable for MDS 9700 Series Switches.

Step 15 switch(config-span)# **no source interface port-channel 1**

Deletes the specified source interface (port-channel 1).

Step 16 switch(config-span)# **shutdown**

Temporarily suspends the session.

Note This is applicable for MDS 9700 Series Switches.

Configuring SPAN Filter

To configure a SPAN filter, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **span session 1**

switch(config-span)#

Configures the specified session (1).

Note In Cisco MDS 9700 Series Switches, SPAN is replaced by monitor session 1.

Step 3 switch(config-span)# **source interface fc9/1 tx**

Configures the source fc9/1 interface in the egress (Tx) direction.

Step 4 switch(config-span)# **source filter vsan 1-2**

Configures VSANs 1 and 2 as session filters.

Step 5 switch(config-span)# **source interface fc7/1 rx**

Configures the source fc7/1 interface in the ingress (Rx) direction.

Configuring SPAN for Generation 2 Fabric Switches

Cisco Generation 2 fabric switches (such as MDS 9124) support SPAN sessions in both directions, Rx and Tx.

Note While using Generation 2 fabric switches, you cannot create an additional active SPAN session when you already have one.

You can specify multiple SPAN source interfaces in Rx and Tx directions. However, the direction should be explicitly mentioned at the end of the command. The SPAN will reject any source interface configuration that fails to mention the direction.

Configuring Ingress SPAN Sessions

To configure for ingress SPAN sessions, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **span session 1**

switch(config-span)#

Configures the specified session (1).

Step 3 switch(config-span)# **destination interface fc1/1**

Configures interface fc1/1 as the destination.

Step 4 switch(config-span)# **source interface fc1/2 rx**

Configures the source interface fc1/2 in the ingress direction.

Configuring Egress SPAN Session

To configure for egress SPAN sessions, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **span session 1**

switch(config-span)#

Configures the specified session (1).

Step 3 switch(config-span)# **destination interface fc1/1**

Configures interface fc1/1 as the destination.

Step 4 switch(config-span)# **source interface fc1/2 tx**

Configures the source interface fc1/2 in the egress direction.

Examples

This example shows how to configure Cisco MDS 9124 for Multiple SPAN Interfaces

```
switch(config-span)# span session 1
switch(config-span)# destination interface fc1/1
switch(config-span)# source interface fc1/2 rx
switch(config-span)# source interface fc1/2 tx
```

Generation 2 Fabric Switches support VSAN filters for one VSAN only in the egress direction; this restriction does not apply to the ingress direction. For example, if you have an interface that is a TE port, with an active VSAN of 1 to 5, and you specify a VSAN filter for VSAN 2, then only the traffic on VSAN 2 will be filtered.

```
switch(config-span)# span session 1
switch(config-span)# source filter vsan 2
switch(config-span)# destination interface fc1/1
switch(config-span)# source interface fc1/2 tx
```

However, if you specify the VSAN filter for VSANs 1 to 2, then traffic from all VSANs (1 to 5) is filtered, which makes the filter useless.

```
switch(config-span)# span session 1
switch(config-span)# source filter vsan 1-2
switch(config-span)# destination interface fc1/1
switch(config-span)# source interface fc1/2 tx
```

Suspending and Reactivating SPAN Sessions

You can temporarily deactivate (suspend) any SPAN session. The traffic monitoring is stopped during this time.

To temporarily suspend or reactivate a SPAN session filter, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **span session 1**

switch(config-span)#

Configures the specified session (1).

Step 3 switch(config-span)# **suspend**

Temporarily suspends the session.

Step 4 switch(config-span)# **no suspend**

Reactivates the session.

Encapsulating Frames

The frame encapsulation feature is disabled by default. If you enable the encapsulation feature, all outgoing frames are encapsulated.

The **switchport encap eisl** command only applies to SD port interfaces. If encapsulation is enabled, you see a new line (Encapsulation is eisl) in the **show interface** *SD_port_interface* command output.

To encapsulate outgoing frames (optional), follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **interface fc9/32**

Configures the specified interface.

Step 3 switch(config-if)# **switchport mode SD**

Configures the SD port mode for interface fc9/32.

Step 4 switch(config-if)# **switchport encap eisl**

Enables the encapsulation option for this SD port.

Step 5 switch(config-if)# **no switchport encap eisl**

Disables (default) the encapsulation option.

Configuring Fibre Channel Analyzers Using SPAN

To configure SPAN on the source and destination interfaces, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **span session 1**

switch(config-span)#

Creates the SPAN session 1.

Step 3 switch(config-span)## **destination interface fc2/1**

Configures the destination interface fc2/1.

Step 4 switch(config-span)# **source interface fc1/1 rx**

Configures the source interface fc1/1 in the ingress direction.

Step 5 switch(config)# **span session 2**

switch(config-span)#

Creates the SPAN session 2.

Step 6 switch(config-span)## **destination interface fc2/2**

Configures the destination interface fc2/2.

Step 7 switch(config-span)# **source interface fc1/1 tx**

Configures the source interface fc1/1 in the egress direction.

To configure Fibre Channel Analyzers using SPAN for the example in , follow these steps:

Procedure

Step 1 Configure SPAN on interface fc1/1 in the ingress (Rx) direction to send traffic on SD port fc2/1 using session 1.

Step 2 Configure SPAN on interface fc1/1in the egress (Tx) direction to send traffic on SD port fc2/2 using session 2.

Step 3 Physically connect fc2/1 to port 1 on the Fibre Channel analyzer.

Step 4 Physically connect fc2/2 to port 2 on the Fibre Channel analyzer.

Configuring Single SD Port to Monitor Traffic

To configure SPAN on a single SD port, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **span session 1**

switch(config-span)#

Creates the SPAN session 1.

Step 3 switch(config-span)## **destination interface fc2/1**

Configures the destination interface fc2/1.

Step 4 switch(config-span)# **source interface fc1/1**

Configures the source interface fc1/1 on the same SD port.

Configuring the Source Switch

This section identifies the tasks that must be performed in the source switch (Switch S):

Creating VSAN Interfaces

To create a VSAN interface in the source switch for the scenario in , follow these steps:

Procedure

Step 1 switchS# **configure terminal**

Enters configuration mode.

Step 2 switchS(config)# **interface vsan 5**

switchS(config-if)#

Configures the specified VSAN interface (VSAN 5) in the source switch (switch S).

Step 3 switchS(config-if)# **ip address 10.10.10.1 255.255.255.0**

Configures the IPv4 address and subnet for the VSAN interface 5 in the source switch (switch S).

Step 4 switchS(config-if)# **no shutdown**

Enables traffic flow through this interface.

Enabling FC Tunnels

Note

- FC tunnels do not work over nontrunking ISLs.
- The interface cannot be operationally up until the FC tunnel mapping is configured in the destination switch.

To enable the FC tunnel feature, follow these steps:

Procedure

Step 1 switchS# **configure terminal**

Enters configuration mode.

Step 2 switchS(config)# **fc-tunnel enable**

Enables the FC tunnel feature (disabled by default).

Note Be sure to enable this feature in each switch in the end-to-end path in the fabric.

Initiating the FC Tunnel

To initiate the FC tunnel in the source switch for the scenario in , follow these steps:

Procedure

Step 1 switchS# **configure terminal**

Enters configuration mode.

Step 2 switchS(config)# **interface fc-tunnel 100**

switchS(config-if)#

Initiates the FC tunnel (100) in the source switch (switch S). The tunnel IDs range from 1 to 255.

Step 3 switchS(config-if)# **source 10.10.10.1**

Maps the IPv4 address of the source switch (switch S) to the FC tunnel (100).

Step 4 switchS(config-if)# **destination 10.10.10.2**

Maps the IPv4 address of the destination switch (switch D) to the FC tunnel (100).

Step 5 switchS(config-if)# **no shutdown**

Enables traffic flow through this interface.

Configuring the ST Port

Note ST ports cannot be configured using Storage Services Modules (SSMs).

To configure an ST port, follow these steps:

Procedure

Step 1 switchS# **configure terminal**

Enters configuration mode.

Step 2 switchS(config)# **interface fc2/1**

Configures the specified interface.

Step 3 switchS(config-if)# **switchport mode ST**

Configures the ST port mode for interface fc2/1.

Step 4 switchS(config-if)# **switchport speed 2000**

Configures the ST port speed to 2000 Mbps.

Step 5 switchS(config-if)# **rspan-tunnel interface fc-tunnel 100**

Associates and binds the ST port with the RSPAN tunnel (100).

Step 6 switchS(config-if)# **no shutdown**

Enables traffic flow through this interface.

Configuring an RSPAN Session

A RSPAN session is similar to a SPAN session, with the destination interface being an RSPAN tunnel.

To configure an RSPAN session in the source switch for the scenario in , follow these steps:

Procedure

Step 1 switchS# **configure terminal**

Enters configuration mode.

Step 2 switchS(config)# **span session 2**

switchS(config-span)#

Configures the specified SPAN session (2). If the session does not exist, it is created. The session ID ranges from 1 to 16.

Step 3 switchS(config-span)# **destination interface fc-tunnel 100**

Configures the specified RSPAN tunnel (100) in a session.

Step 4 switchS(config-span)# **source interface fc1/1**

Configures the source interface (fc1/1) for this session and spans the traffic from interface fc1/1 to RSPAN tunnel 100.

Configuring All Intermediate Switches

This section identifies the tasks that must be performed in all intermediate switches in the end-to-end path of the RSPAN tunnel:

Configuring VSAN Interfaces

depicts an RSPAN tunnel configuration terminating in the destination switch (Switch D).

Note This example assumes that VSAN 5 is already configured in the VSAN database.

To create a VSAN interface in the destination switch for the scenario in , follow these steps:

Procedure

Step 1 switchD# **configure terminal**

Enters configuration mode.

Step 2 switchD(config)# **interface vsan 5**

switchD(config-if)#

Configures the specified VSAN interface (VSAN 5) in the destination switch (Switch D).

Step 3 switchD(config-if)# **ip address 10.10.10.2 255.255.255.0**

Configures the IPv4 address and subnet for the VSAN interface in the destination switch (Switch D).

Step 4 switchD(config-if)# **no shutdown**

Enables traffic flow to administratively allow traffic (provided the operational state is up).

Enabling IP Routing

The IP routing feature is disabled by default. Be sure to enable IP routing in each switch (including the source and destination switches) in the end-to-end path in the fabric. This procedure is required to set up the FC tunnel.

Configuring the Destination Switch

This section identifies the tasks that must be performed in the destination switch (Switch D):

Configuring VSAN Interfaces

depicts an RSPAN tunnel configuration terminating in the destination switch (Switch D).

Note This example assumes that VSAN 5 is already configured in the VSAN database.

Configuring the SD Port

Note SD ports cannot be configured using Storage Services Modules (SSMs).

To configure an SD port for the scenario in , follow these steps:

Procedure

Step 1 switchD# **configure terminal**

Enters configuration mode.

Step 2 switchD(config)# **interface fc2/1**

Configures the specified interface.

Step 3 switchD(config-if)# **switchport mode SD**

Configures the SD port mode for interface fc2/1.

Step 4 switchD(config-if)# **switchport speed 2000**

Configures the SD port speed to 2000 Mbps.

Step 5 switchD(config-if)# **no shutdown**

Enables traffic flow through this interface.

Mapping the FC Tunnel

To terminate the FC tunnel in the destination switch for the scenario in , follow these steps:

Procedure

Step 1 switchD# **configure terminal**

Enters configuration mode.

Step 2 switchD(config)# **fc-tunnel tunnel-id-map 100 interface fc2/1**

Terminates the FC tunnel (100) in the destination switch (switch D). The tunnel ID range is from 1 to 255.

Creating Explicit Paths

To create an explicit path for the scenario in , follow these steps:

Before you begin

The explicit path must be created in the source switch. To configure an explicit path, you must first create the path and then configure the use of any one path. If an explicit path is not configured, the minimum cost path is used by default. If an explicit path is configured and is functioning, the specified path is used.

Procedure

Step 1 switchS# **configure terminal**

Enters configuration mode.

Step 2 switchS(config)# **fc-tunnel explicit-path Path1**

switch(config-explicit-path)#

Places you at the explicit path prompt for the path named Path 1.

Step 3 switchS(config-explicit-path)# **next-address 10.10.10.2 strict**

switchS(config-explicit-path)# **next-address 10.10.10.3 strict**

switchS(config-explicit-path)# **next-address 10.10.10.4 strict**

Specifies that the next hop VSAN interface IPv4 addresses and the previous hops specified in the explicit path do not require direct connection.

Step 4 switchS(config)# **fc-tunnel explicit-path Path2**

switch(config-explicit-path)#

Places you at the explicit path prompt for Path2.

Step 5 switchS(config-explicit-path)# **next-address 10.10.10.5 strict**

Example:

```
switchS(config-explicit-path)# next-address 10.10.10.4 strict
```

Specifies that the next hop VSAN interface IPv4 addresses and the previous hops specified in the explicit path do not require direct connection.

Step 6 switchS(config)# **fc-tunnel explicit-path Path3**

switch(config-explicit-path)#

Places you at the explicit path prompt for Path3.

Step 7 switchS(config-explicit-path)# **next-address 10.10.10.3 loose**

Configures a minimum cost path in which the 10.10.10.3 IPv4 address exists.

Note In , Path 3 is the same as Path 1—10.10.10.3 exists in Path 1. Using the **loose** option, you can achieve the same effect with one command instead of issuing three commands (using the **strict** option) in Step 3.

Referencing the Explicit Path

To reference the explicit path, follow these steps:

Procedure

Step 1 switchS# **configure terminal**

Enters configuration mode.

Step 2 switchS(config)# **interface fc-tunnel 100**

References the tunnel ID for Path1.

Step 3 switchS(config)# **explicit-path Path1**

Links Path1 to the tunnel ID.

This configuration explicitly specifies Path 1 to be used for the RSPAN traffic. Refer to RFC 3209 for further details on explicit paths and source-based routing.

Monitoring RSPAN Traffic

Once the session is configured, other SPAN sources for this session can also be configured as required. Figure 29: Fibre Channel Analyzer Using a Single SD Port to Monitor RSPAN Traffic, on page 320 shows an RSPAN setup where one session with destination port fc2/1 and source interface fc1/1 is used to capture traffic in both ingress and egress directions.

Figure 29: Fibre Channel Analyzer Using a Single SD Port to Monitor RSPAN Traffic

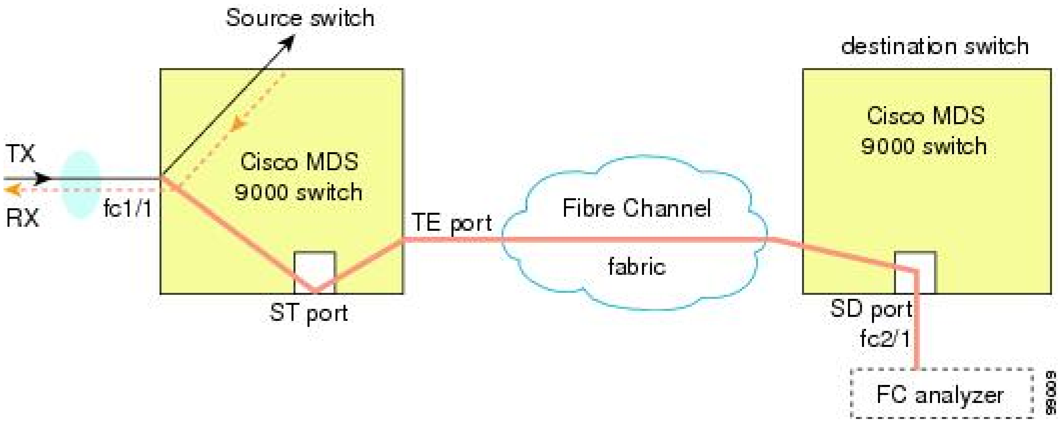

To use this setup, the analyzer should have the capability of distinguishing ingress and egress traffic for all captured frames.

Verifying SPAN Configuration

To display the SPAN configuration information, perform one of the following tasks:

Command	Purpose
show span	Displays SPAN Sessions in a Brief Format **Note** In Cisco MDS 9700 Series Switches, **show span** command is replaced by **show monitor** command.
show span session 7	Displays a Specific SPAN Session in Detail **Note** In Cisco MDS 9700 Series Switches, **show span session 7** command is replaced by **show monitor session 7** command.
show span session	Displays ALL SPAN Sessions **Note** In Cisco MDS 9700 Series Switches, **show span session** command is replaced by **show monitor session all** command.
show int fc9/32	Displays an SD Port Interface with Encapsulation Enabled
show interface brief	Displays ST Port Interface Information
show interface fc1/11	Displays Detailed Information for the ST Port Interface
show fc-tunnel	Displays the FC Tunnel Status
show fc-tunnel tunnel-id-map	Displays FC Tunnel Egress Mapping Information
show fc-tunnel explicit-path	Displays FC Tunnel Explicit Mapping Information
show interface fc-tunnel 200	Displays the FC Tunnel Interface

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference* .

Displaying SPAN Information

Use the **show span** command to display configured SPAN information. See the following examples:

SPAN Sessions in a Brief Format

The following example displays SPAN sessions in a brief format:

```
switch# show span session brief
--------------------------------------------------------
 Session  Admin          Oper           Destination
          State          State          Interface
--------------------------------------------------------
 7        no suspend     active         fc2/7
 1        suspend        inactive       not configured
 2        no suspend     inactive       fc3/1
```

Specified SPAN Session in Detail

The following example displays a specific SPAN session in detail:

```
switch# show span session 7
Session 7 (active)
   Destination is fc2/7
   No session filters configured
   No ingress (rx) sources
   Egress (tx) sources are
     port-channel 7,
```

The following example configures two SPAN sessions to the same destination SD port. This will send bidirectional traffic for the FCIP interface to the destination port.

```
switch# configure
swtch(config)# span session 1
swtch(config-span)# source interface fcip 104 rx
switch(config-span)# destination interface fc1/5

switch# configure
switch(config)# span session 2
switch(config-span)# source interface fcip 104 tx
switch(config-span)# destination interface fc1/5

switch# show span session 1
Session 1 (active)
Destination is fc1/5
   No session filters configured
   Ingress (rx) sources are
     fcip104,
   No egress (tx) sources


switch# show span session 2
Session 2 (active)
Destination is fc1/5
   No session filters configured
   No ingress (rx) sources
   Egress (tx) sources are
     fcip104,
```

ALL SPAN Sessions

The following example displays ALL SPAN Sessions:

```
switch# show span session
Session 1 (inactive as no destination)
Destination is not specified
   Session filter vsans are 1
   No ingress (rx) sources
   No egress (tx) sources
Session 2 (active)
   Destination is fc9/5
   No session filters configured
   Ingress (rx) sources are
      vsans 1
   No egress (tx) sources
Session 3 (admin suspended)
   Destination is not configured
   Session filter vsans are 1-20
   Ingress (rx) sources are
     fc3/2, fc3/3, fc3/4, fcip 51,
     port-channel 2, sup-fc0,
   Egress (tx) sources are
     fc3/2, fc3/3, fc3/4, sup-fc0,
```

SD Port Interface with Encapsulation Enabled

The following example displays SD Port Interface with Encapsulation Enabled

```
switch# show int fc9/32
fc9/32 is up
    Hardware is Fibre Channel
    Port WWN is 22:20:00:05:30:00:49:5e
    Admin port mode is SD
    Port mode is SD
    Port vsan is 1
    Speed is 1 Gbps
    Receive Buffer Size is 2112
    Encapsulation is eisl
<----------------
Displays the enabled encapsulation status
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      0 frames input, 0 bytes, 0 discards
        0 CRC,  0 unknown class
        0 too long, 0 too short
      0 frames output, 0 bytes, 0 discards
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
```

0 output OLS, 0 LRR, 0 NOS, 0 loop inits

Displaying RSPAN Information

Use the **show** commands to display configured RSPAN information. See the following examples:

ST Port Interface Information

The following example displays ST Port Interface information:

```
switch# show interface brief
-------------------------------------------------------------------------------
Interface   Vsan     Admin   Admin   Status      Oper    Oper    Port-channel
                     Mode    Trunk               Mode    Speed
                             Mode                        (Gbps)
-------------------------------------------------------------------------------
fc1/1       1        auto    on      trunking    TE          2       --
...
fc1/14      1        auto    on      trunking    TE          2       --
fc1/15      1        ST      on      up          ST          2       --
...
```

```
fc2/9       1       auto    on      trunking    TE          2       port-channel 21
fc2/10      1       auto    on      trunking    TE          2       port-channel 21
...
fc2/13      999     auto    on      up          F           1       --
fc2/14      999     auto    on      up          FL          1       --
fc2/15      1       SD      --      up          SD          2       --
fc2/16      1       auto    on      trunking    TE          2       --
--------------------------------------------------------------------------------------
Interface           Status      Speed
                                (Gbps)
--------------------------------------------------------------------------------------
sup-fc0             up          1
--------------------------------------------------------------------------------------
Interface           Status      IP Address              Speed       MTU
--------------------------------------------------------------------------------------
mgmt0               up        172.22.36.175/22          100 Mbps    1500
--------------------------------------------------------------------------------------
Interface           Status      IP Address              Speed       MTU--
--------------------------------------------------------------------------------------
vsan5               up          10.10.10.1/24           1 Gbps      1500
--------------------------------------------------------------------------------------
Interface           Vsan        Admin           Status      Oper    Oper
                                Trunk                       Mode    Speed
                                Mode                                (Gbps)
--------------------------------------------------------------------------------------
port-channel 21     1           on              trunking    TE      4
--------------------------------------------------------------------------------------
Interface           Status      Dest IP Addr    Src IP Addr     TID    Explicit Path
--------------------------------------------------------------------------------------
fc-tunnel 100       up          10.10.10.2      10.10.10.1      100
```

Detailed Information for the ST Port Interface

The following example displays detailed information for the ST port interface:

```
switch# show interface fc1/11
fc1/11 is up
    Hardware is Fibre Channel
    Port WWN is 20:0b:00:05:30:00:59:de
   Admin port mode is ST
    Port mode is ST
    Port vsan is 1
    Speed is 1 Gbps
   Rspan tunnel is fc-tunnel 100
    Beacon is turned off
    5 minutes input rate 248 bits/sec, 31 bytes/sec, 0 frames/sec
    5 minutes output rate 176 bits/sec, 22 bytes/sec, 0 frames/sec
      6862 frames input, 444232 bytes
        0 discards, 0 errors
        0 CRC, 0 unknown class
        0 too long, 0 too short
      6862 frames output, 307072 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      0 output OLS, 0 LRR, 0 NOS, 0 loop inits
```

FC Tunnel Status

The following example displays FC tunnel status:

```
switch# show fc-tunnel
fc-tunnel is enabled
```

FC Tunnel Egress Mapping Information

The following example displays FC tunnel egress mapping information:

```
switch# show fc-tunnel tunnel-id-map
tunnel id egress interface
    150     fc3/1
    100     fc3/1
```

Note Multiple tunnel IDs can terminate at the same interface.

FC Tunnel Explicit Mapping Information

The following example displays FC tunnel mapping information:

```
switch# show fc-tunnel explicit-path
Explicit path name: Alternate1
       10.20.1.2 loose
       10.20.1.3 strict
Explicit path name: User2
       10.20.50.1 strict
       10.20.50.4 loose
```

SPAN Mapping Information

The following example displays the SPAN mapping information

```
switch# show span session
Session 2 (active)
   Destination is fc-tunnel 100
   No session filters configured
   Ingress (rx) sources are
     fc2/16,
   Egress (tx) sources are
     fc2/16,
```

FC Tunnel Interface

The following example displays the FC Tunnel Interface:

```
switch# show interface fc-tunnel 200
fc-tunnel 200 is up
Dest   IP Addr: 200.200.200.7   Tunnel ID: 200
Source IP Addr: 200.200.200.4   LSP ID: 1
Explicit Path Name:
```

Configuration Examples for RSPAN

Note RSPAN can be combined with the local SPAN feature so SD ports forward local SPAN traffic along with remote SPAN traffic. Various SPAN source and tunnel scenarios are described in this section.

Single Source with One RSPAN Tunnel

The source Switch S and the destination Switch D are interconnected through a Fibre Channel fabric. An RSPAN tunnel is configured as a destination interface for the SPAN session and the ST port forwards SPAN traffic through the RSPAN tunnel(see Figure 30: RSPAN Scenario with One Source Switch, One Destination Switch, and One Tunnel, on page 326).

Figure 30: RSPAN Scenario with One Source Switch, One Destination Switch, and One Tunnel

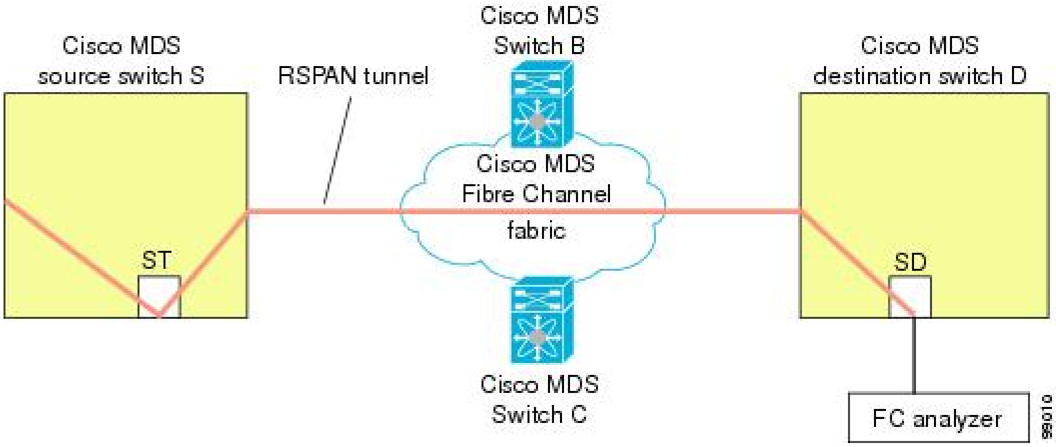

Single Source with Multiple RSPAN Tunnels

Single Source with Multiple RSPAN Tunnels , on page 326 displays two separate RSPAN tunnels configured between Switches S and N. Each tunnel has an associated ST port in the source switch and a separate SD port in the destination switch. This configuration is useful for troubleshooting purposes.

Figure 31: RSPAN Scenario with One Source Switch, One Destination Switch, and Multiple Tunnels

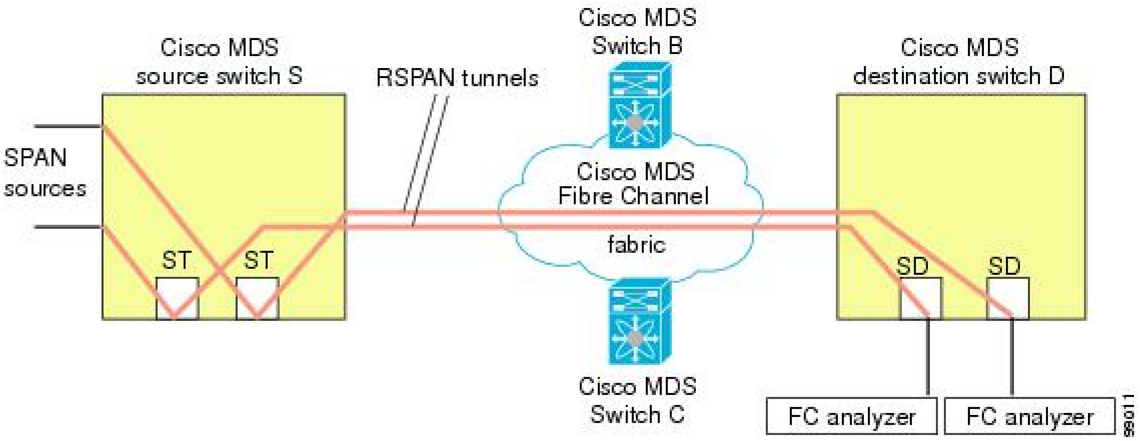

Multiple Sources with Multiple RSPAN Tunnels

Figure 32: RSPAN Scenario with Two Source Switches, a Destination Switch, and Multiple Tunnels, on page 327 displays two separate RSPAN tunnels configured between Switches S1 and S2. Both tunnels have an associated ST port in their respective source switch and terminate in the same SD port in the destination switch.

Figure 32: RSPAN Scenario with Two Source Switches, a Destination Switch, and Multiple Tunnels

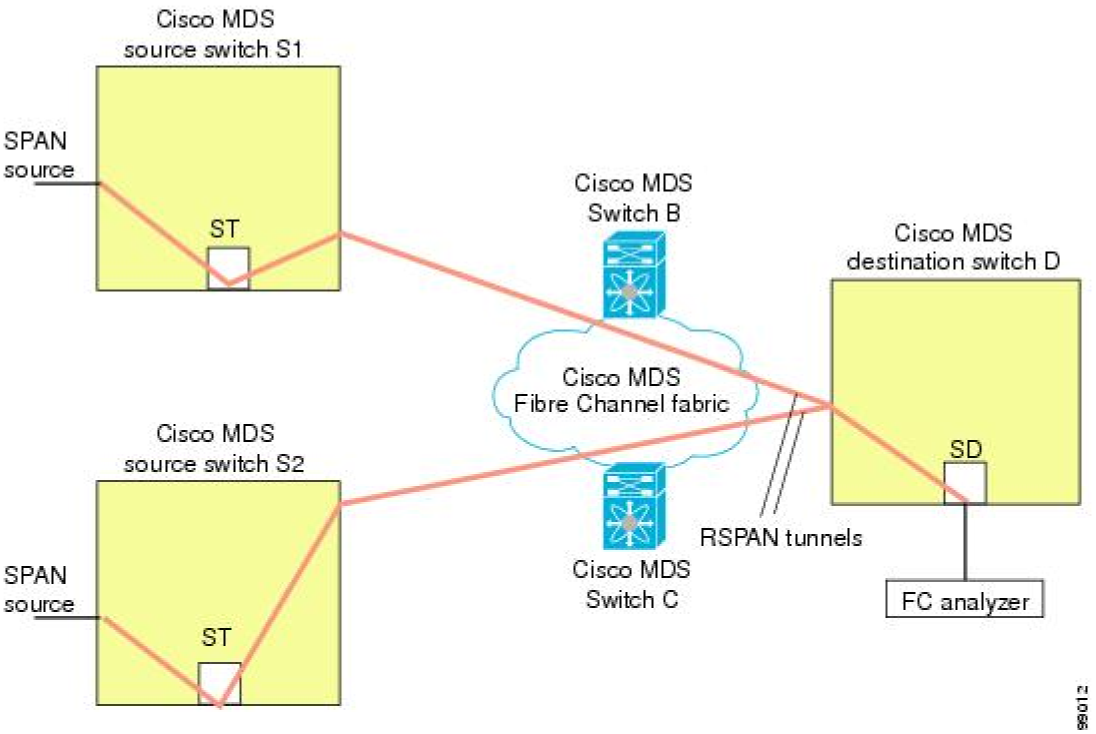

This configuration is useful for remote monitoring purposes. For example, the administrator may be at the destination switch and can remotely monitor the two source switches.

CHAPTER 17

Configuring Fabric Configuration Server

This chapter describes the Fabric Configuration Server (FCS) feature provided in the Cisco MDS 9000 Family of directors and switches.

Information About FCS

The Fabric Configuration Server (FCS) provides discovery of topology attributes and maintains a repository of configuration information of fabric elements. A management application is usually connected to the FCS on the switch through an N port. The FCS views the entire fabric based on the following objects:

- Interconnect element (IE) object—Each switch in the fabric corresponds to an IE object. One or more IE objects form a fabric.
- Port object—Each physical port in an IE corresponds to a port object. This includes the switch ports (xE, Fx, and TL ports) and their attached Nx ports.
- Platform object—A set of nodes may be defined as a platform object to make it a single manageable entity. These nodes are end-devices (host systems, storage subsystems) attached to the fabric. Platform objects reside at the edge switches of the fabric.

Each object has its own set of attributes and values. A null value may also be defined for some attributes.

In the Cisco MDS 9000 Family switch environment, multiple VSANs constitute a fabric, where one instance of the FCS is present per VSAN.

As of Cisco NX-OS Release 4.1(1), FCS supports the discovery of virtual devices. The **fcs virtual-device-add** command, issued in FCS configuration submode, allows you to discover virtual devices in a particular VSAN or in all VSANs. The devices that are zoned for IVR must be discovered with this command and have request domain_ID (RDI) enabled, before activating the IVR zone set.

If you have attached a management application to a switch, all the frames directed towards the FCS in the switch are part of the port VSAN in the switch port (Fx port). Your view of the management application is limited only to this VSAN. However, information about other VSANs that this switch is part of can be obtained either through the SNMP or CLI.

In the following image, Management Application 1 (M1) is connected through an F port with port VSAN ID 1, and Management Application 2 (M2) is connected through an F port with port VSAN ID 2. M1 can query the FCS information of switches S1 and S3, and M2 can query switches S3 and S4. Switch S2 information is not known to both of them. FCS operations can be done only on those switches that are visible in the VSAN. Note that M2 can send FCS requests only for VSAN 2 even though S3 is also a part of VSAN 1.

Figure 33: FCSs in a VSAN Environment

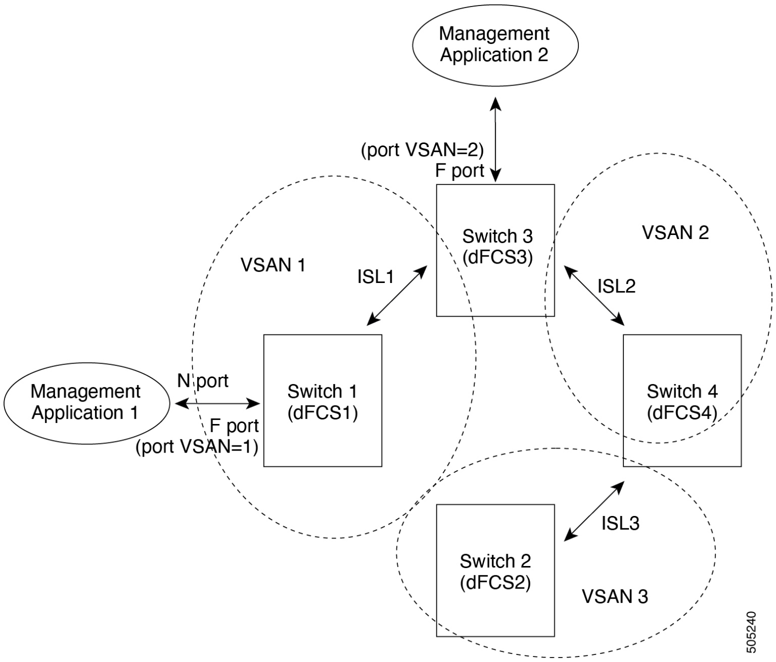

Significance of FCS

This section lists the significance of FCSs.

- FCSs support network management including the following:
    - N port management application can query and obtain information about fabric elements.
    - SNMP manager can use the FCS management information base (MIB) to start discovery and obtain information about the fabric topology.
- FCSs support TE and TL ports in addition to the standard F and E ports.
- FCS can maintain a group of modes with a logical name and management address when a platform registers with it. FCSs maintain a backup of all registrations in secondary storage and update it with every change. When a restart or switchover happens, FCSs retrieve the secondary storage information and rebuild its database.
- SNMP manager can query FCSs for all IEs, ports, and platforms in the fabric.

Default Settings

Table 41: Default FCS Settings , on page 331 lists the default FCS settings.

Table 41: Default FCS Settings

Parameters	Default
Global checking of the platform name	Disabled.
Platform node type	Unknown.

Configuring FCS

The Fabric Configuration Server (FCS) provides discovery of topology attributes and maintains a repository of configuration information of fabric elements.

Specifying a FCS Name

You can specify if the unique name verification is for the entire fabric (globally) or only for locally (default) registered platforms.

Note Set this command globally only if all switches in the fabric belong to the Cisco MDS 9000 Family.

To enable global checking of the platform name, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcs plat-check-global vsan 1**

Enables global checking of the platform name.

Step 3 switch(config)# **no fcs plat-check-global vsan 1**

Disables (default) global checking of the platform name.

Registering Platform Attributes

To register platform attributes, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch(config)# **fcs register**

switch(config-fcs-register)#

Enters the FCS registration submode.

Step 3 switch(config-fcs-register)# **platform name SamplePlatform vsan 1**

switch(config-fcs-register-attrib)#

Enters the FCS registration attributes submode.

Step 4 switch(config-fcs-register)# **no platform name SamplePlatform vsan 1**

switch(config-fcs-register)#

Deletes a registered platform.

Step 5 switch(config-fcs-register-attrib)# **mgmt-addr 1.1.1.1**

Configures the platform management IPv4 address.

Step 6 switch(config-fcs-register-attrib)# **no mgmt-addr 1.1.1.1**

Deletes the platform management IPv4 address.

Step 7 switch(config-fcs-register-attrib)# **mgmt-addr 2001:0DB8:800:200C::417A**

Configures the platform management IPv6 address.

Step 8 switch(config-fcs-register-attrib)# **no mgmt-addr 2001:0DB8:800:200C::417A**

Deletes the platform management IPv6 address.

Step 9 switch(config-fcs-register-attrib)# **nwwn 11:22:33:44:55:66:77:88**

Configures the platform node name.

Step 10 switch(config-fcs-register-attrib)# **no nwwn 11:22:33:44:55:66:77:88**

Deletes the platform node name.

Step 11 switch(config-fcs-register-attrib)# **type 5**

Configures the fc-gs-3 defined platform type.

Step 12 switch(config-fcs-register-attrib)# **no type 5**

Deletes the configured type and reverts the switch to its factory default of unknown type.

Step 13 switch(config-fcs-register-attrib)# **exit**

Exits the FCS registration attributes submode.

Step 14 switch(config-fcs-register)# **exit**

Exits the FCS registration submode.

Verifying FCS Configuration

To display the FCS configuration information, perform one of the following tasks:

Command	Purpose
show fcs database	Displays FCS Local Database Information.
show fcs ie vsan 1	Displays a List of All IEs for a Specific VSAN.
show fcs ie nwwn 20:01:00:05:30:00:16:df vsan 1	Displays Interconnect Element Object Information for a Specific nWWN
show fcs platform name SamplePlatform vsan 1	Displays Information for a Specific Platform
show fcs platform vsan 1	Displays a List of Platforms for a Specified VSAN
show fcs port vsan 24	Displays a List of Switch Ports in a Specified VSAN
show fcs port pwwn 20:51:00:05:30:00:16:de vsan 24	Displays Port Information for a Specified pWWN
show fcs statistics	Displays FCS Statistics
show fcs vsan	Displays Platform Settings for Each VSAN

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS 9000 Family Command Reference* .

Displaying FCS Elements

Use the **show fcs** commands to display the status of the WWN configuration (see Example FCS Local Database Information, on page 333 to Platform Settings for Each VSAN, on page 336).

FCS Local Database Information

The following example displays FCS local database information:

```
switch# show fcs database
FCS Local Database in VSAN: 1
------------------------------
Switch WWN               : 20:01:00:05:30:00:16:df
Switch Domain Id         : 0x7f(127)
Switch Mgmt-Addresses    : snmp://172.22.92.58/eth-ip
                           http://172.22.92.58/eth-ip
Fabric-Name              : 20:01:00:05:30:00:16:df
Switch Logical-Name      : 172.22.92.58
Switch Information List  : [Cisco Systems*DS-C9509*0*20:00:00:05:30:00
Switch Ports:
-------------------------------------------------------------------
Interface  pWWN                     Type     Attached-pWWNs
-------------------------------------------------------------------
```

```
fc2/1      20:41:00:05:30:00:16:de  TE       20:01:00:05:30:00:20:de
fc2/2      20:42:00:05:30:00:16:de  Unknown  None
fc2/17     20:51:00:05:30:00:16:de  TE       20:0a:00:05:30:00:20:de
FCS Local Database in VSAN: 5
------------------------------
Switch WWN                : 20:05:00:05:30:00:12:5f
Switch Domain Id          : 0xef(239)
Switch Mgmt-Addresses     : http://172.22.90.171/eth-ip
                            snmp://172.22.90.171/eth-ip
                            http://10.10.15.10/vsan-ip
                            snmp://10.10.15.10/vsan-ip
Fabric-Name               : 20:05:00:05:30:00:12:5f
Switch Logical-Name       : 172.22.90.171
Switch Information List   : [Cisco Systems*DS-C9509**20:00:00:05:30:00:12:5e]
Switch Ports:
-------------------------------------------------------------------
Interface  pWWN                     Type     Attached-pWWNs
-------------------------------------------------------------------
fc3/1      20:81:00:05:30:00:12:5e  TE       22:01:00:05:30:00:12:9e
fc3/2      20:82:00:05:30:00:12:5e  TE       22:02:00:05:30:00:12:9e
fc3/3      20:83:00:05:30:00:12:5e  TE       22:03:00:05:30:00:12:9e
```

List of All IEs for a Specific VSAN

The following example displays list of all IEs for a specific VSAN:

```
switch# show fcs ie vsan 1
IE List for VSAN: 1
-------------------------------------------------------------------
IE-WWN                   IE-Type                           Mgmt-Id
-------------------------------------------------------------------
20:01:00:05:30:00:16:df  Switch (Local)                    0xfffc7f
20:01:00:05:30:00:20:df  Switch (Adjacent)                 0xfffc64
[Total 2 IEs in Fabric]
```

Interconnect Element Object Information for a Specific nWWN

The following example displays interconnect element object information for a specific nWWN:

```
switch# show fcs ie nwwn 20:01:00:05:30:00:16:df vsan 1
IE Attributes
-------------
Domain-Id = 0x7f(127)
Management-Id = 0xfffc7f
Fabric-Name = 20:01:00:05:30:00:16:df
Logical-Name = 172.22.92.58
Management Address List =
        snmp://172.22.92.58/eth-ip
        http://172.22.92.58/eth-ip
Information List:
        Vendor-Name = Cisco Systems
        Model Name/Number = DS-C9509
        Release-Code = 0
```

Information for a Specific Platform

The following example displays information for a specific platform:

```
switch# show fcs platform name SamplePlatform vsan 1
Platform Attributes
-------------------
Platform Node Names:
        11:22:33:44:55:66:77:88
Platform Type = Gateway
Platform Management Addresses:
        1.1.1.1
```

List of Platforms for a Specified VSAN

The following example displays list of platforms for a specified VSAN:

```
switch# show fcs platform vsan 1
Platform List for VSAN: 1
Platform-Names
--------------
SamplePlatform
[Total 1 Platforms in Fabric]
```

List of Switch Ports in a Specified VSAN

The following example displays a list of switch ports in a specified VSAN:

```
switch# show fcs port vsan 24
Port List in VSAN: 24
                 -- IE WWN: 20:18:00:05:30:00:16:df  --
-------------------------------------------------------------------------
Port-WWN                 Type      Module-Type            Tx-Type
-------------------------------------------------------------------------
20:41:00:05:30:00:16:de  TE_Port   SFP with Serial Id      Shortwave Laser
20:51:00:05:30:00:16:de  TE_Port   SFP with Serial Id      Shortwave Laser
[Total 2 switch-ports in IE]
                 -- IE WWN: 20:18:00:05:30:00:20:df  --
-------------------------------------------------------------------------
Port-WWN                 Type      Module-Type            Tx-Type
-------------------------------------------------------------------------
20:01:00:05:30:00:20:de  TE_Port   SFP with Serial Id      Shortwave Laser
20:0a:00:05:30:00:20:de  TE_Port   SFP with Serial Id      Shortwave Laser
[Total 2 switch-ports in IE]
```

Port Information for a Specified pWWN

The following example displays port information for a specified pWWN:

```
switch# show fcs port pwwn 20:51:00:05:30:00:16:de vsan 24
Port Attributes
---------------
Port Type = TE_Port
Port Number = 0x1090000
Attached-Port-WWNs:
        20:0a:00:05:30:00:20:de
Port State = Online
```

FCS Statistics

The following example displays FCS statistics:

```
switch# show fcs statistics
FCS Statistics for VSAN: 1
---------------------------
FCS Rx Get Reqs   :2
FCS Tx Get Reqs   :7
FCS Rx Reg Reqs   :0
FCS Tx Reg Reqs   :0
FCS Rx Dereg Reqs :0
FCS Tx Dereg Reqs :0
FCS Rx RSCNs      :0
...
FCS Statistics for VSAN: 30
---------------------------
FCS Rx Get Reqs   :2
FCS Tx Get Reqs   :2
FCS Rx Reg Reqs   :0
FCS Tx Reg Reqs   :0
FCS Rx Dereg Reqs :0
FCS Tx Dereg Reqs :0
FCS Rx RSCNs      :0
FCS Tx RSCNs      :0
...
```

Platform Settings for Each VSAN

The following example displays platform settings for each VSAN:

```
switch# show fcs vsan
------------------------------
VSAN    Plat Check fabric-wide
------------------------------
0001    Yes
0010    No
0020    No
0021    No
0030    No
```

Additional References

For additional information related to implementing FCS, see the following section:

Table 42: MIBs

MIBs	MIBs Link
CISCO-FCS-MIB	To locate and download MIBs, go to the following URL: http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html

CHAPTER 18

Fabric Module Error Monitoring

This chapter provides information about Fabric Module Error Monitoring (XbarErrorMonitor) and how to configure it.

Feature History for Fabric Module Error Monitoring

Feature Name	Releases	Feature Information
Fabric Module Error Monitoring (XbarErrorMonitor)	9.3(1)	This feature is introduced.

About Fabric Module Error Monitoring

The fabric modules of Cisco MDS are commonly called Xbars. There are two versions of these fabric modules: Fabric 1 and Fabric 3. Frames that are received by an FC port with CRC error are dropped and not forwarded further. As frames move from component to component and module to module errors may possibly occur. Frames are CRC checked at several places along the switching path. Once a frame is detected as an error it is dropped as soon as possible.

The existing ' Internal CRC Detection and Isolation' feature can detect and take corrective action when these internal CRC errors occur. However, fabric modules can experience other errors that are not, strictly speaking, internal CRC errors. The Fabric Module Error Monitoring (XbarErrorMonitor) feature, introduced in Cisco MDS Release 9.3(1), complements the 'Internal CRC Detection and Isolation' feature, and is designed to detect and take corrective action in the presence of these errors. This feature enables you to monitor certain specific hardware counters in Fabric 1 and Fabric 3 modules that might cause I/O problems in the network setup.

XbarErrorMonitor is a python script that utilizes the MDS scheduler feature to check for these internal errors. It works by having the scheduler execute it periodically (default 120 seconds). Each time it executes it issues a 'show hardware internal errors' command and records the specific counters monitored for the specific fabric module type present in the switch. It then sleeps (suspends) for an amount of time (default 30 seconds), issues another 'show hardware internal errors' command and compares each specific counter to the previous value.

If one or more of the monitored counters is greater than or equal to the threshold value (default 50) then the action specified (default 'log-only') is performed.

Guidelines and Limitations of Fabric Module Error Monitoring

- This feature supports only Cisco MDS 9700 series switches.
- This feature gets automatically enabled when you upgrade to Cisco MDS Release 9.3(1). This feature has the default values - Scheduling interval is 120 seconds, Sleep time is 30 seconds, counter threshold is 50, and default action is log-only.
- This feature monitors the following error counters in a switch:
    - Fabric 1 module counters

      - INTERNAL_ERROR_CNT

      - HIGH_XT_DROP_CNT

      - SAC_XTIMEOUT_INTR_HI
    - Fabric 3 module counters

      - packets dropped destined to port

      - packets drop on receive port

      - double bit ecc error

Note These counters can be seen (if they are non-zero) using the **show hardware internal errors** command.

- By default, this feature monitors the Fabric module 1 and Fabric module 3 counters every two minutes. If a counter crosses the default threshold value of 50, it displays a syslog saying that there is a fault in the respective spine. For example:

```
2022 Jun 28 14:10:38 sw9706-89 %USER-2-SYSTEM_MSG:
xbarErrorMonitor: counter threshold exceeded for xbar 3 for
counter packets dropped destined to port. (Before: 0, After: 128, Delta 128).
```

- If XbarErrorMonitor was started with a certain set of parameters then ensure that when changing any parameter all non-default parameters are specified. For example:

```
xbarErrorMonitor -si 180 enable
xbarErrorMonitor -a log-and-out-of-service enable
```

- When you use the xbarErrorMonitor log-and-out-ofservice enable command, the scheduling interval will be changed back to default 120 second as the si parameter don’t get passed.
- Enabling xbarErrorMonitor creates a scheduler job named xbarErrorMonitor_job and a scheduler schedule named XbarErrorMonitor_Schedule. These should not be deleted or xbarErrorMonitor will cease to function.

Configuring Fabric Module Error Monitoring

Procedure

	Command or Action	Purpose
Step 1	switch# xbarErrorMonitor enable	Enables the XbarErrorMonitor feature on the switch.
Step 2	switch# xbarErrorMonitor disable	(Optional) Disables the XbarErrorMonitor feature.
Step 3	switch# xbarErrorMonitor -h	Modifies the error monitoring parameters. **Note** The help option provides a list of parameters that you can modify so that the monitoring can be done based on the selected parameter. See the 'Configuration Examples' section for a sample usage of this option.
Step 4	switch# xbarErrorMonitor show	Verifies the status of xbar error monitoring.

Configuration Examples

The following example shows how to check the status of the XbarErrorMonitor feature:

```
switch# xbarErrorMonitor show
xbarErrorMonitor 1.0

Status: Enabled
Schedular Interval: 120
Sleep Time: 30
Counter Threshold: 50
Action: log-only
Counters Monitored:
  packets dropped destined to port
  packets drop on receive port
  double bit ecc error
```

The following example shows how to modify the error monitoring parameters:

```
Switch(config)# xbarErrorMonitor --help
usage: xbarErrorMonitor [-h] [-v] [-si] [-st] [-t] [-a]
                        {enable,disable,show,forScheduler} ...

Enable/Disable xbar error monitor on the switch

positional arguments:
  {enable,disable,show,forScheduler}
    enable              Enable xbarErrorMonitor feature
    disable             Disable xbarErrorMonitor feature
    show                Show current status of xbarErrorMonitor feature
    forScheduler        This option is for schedular only, DO NOT USE THIS
                        MANUALLY

optional arguments:
  -h, --help            show this help message and exit
  -v, --version         show program's version number and exit
  -si , --schedular-interval
                        scheduler interval time, value should be between 120s
                        to 3600s. Default value is 120s.
  -st , --sleep-time    sleep time between getting error counters, value
                        should be between 30s to 90s. Default value is 30s.
  -t , --counter-threshold
                        counter threshold value beyond which action will be
                        taken, value should be between 50 to 500. Default
                        value is 50.
  -a , --action         action that needs to be taken when counter breaches
                  the threshold value. 'log-only': Shows only a syslog,
                        'log-and-out-of-service': Shows a syslog as well as
                        puts the xbar out-of-service. Default action is log-
                        only.
```

CHAPTER 19

Configuring Port Pacing

This chapter describes how to configure Port Pacer.

Information About Port Pacing

The Fibre Channel Port Pacer is supported only on Cisco MDS 9700 Director Series switches. The Port Pacer is designed to pace the number of mode F ports that come up simultaneously so that ports are brought up in a phased manner.

During an F port start up, the Port Pacer informs the F Port server that a port is starting up. The Port Pacer waits for the F port server to receive FLOGIs and FDISCs on that port. The Port Pacer attempts to bring up concurrent-ports number of ports simultaneously. However, after the F port server informs the Port Pacer that it has received FLOGI and FDISC for that port, and then the Port Pacer completes the port bring up and updates the port status as up. Subsequently, the next port is attempted for starting up.

By default, F port pacing is disabled. After enabling port pacing, number of FLOGI or FDISC received on the port are being tracked. In the case of all FLOGI or FDISC successfully logged in, which would take few seconds, another set of concurrent ports are brought up. At any given time, FLOGI is processed only for configured concurrent ports. This feature is useful in case of zero FLOGI retries in the hosts.

Guidelines and Limitations

Following are the recommended guidelines and requirements for enabling the Port Pacer:

- Port pacing configurations are supported only for admin port mode F.
- Concurrent-ports port-number needs to be set depending upon the topology and set this value on how many F ports can be brought up simultaneously.

Configuring Port Pacer

Enabling Port Pacing

Note Port pacing configuration is supported only for admin port mode F.

Port pacing command is a system wide command applicable for all admin port mode F ports.

To enable the port pacer, perform these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch# (config)# **system port pacer mode F interface-login-threshold 10 concurrent-ports 1**

Enables the pacer mode for F port with concurrency of 1 and threshold set as 10.

interface-login-threshold specifies the number of FLOGI or FDISC expected on a port.

concurrent-ports specifies the number of admin port mode F ports that can be brought up simultaneously.

Disabling the Port Pacing Configuration

To disable the port pacing configuration, follow these steps:

Procedure

Step 1 switch# **configure terminal**

Enters configuration mode.

Step 2 switch# (config)# **no system port pacer mode F interface-login-threshold 10 concurrent-ports 1**

Disables the pacer mode for F port.