Send documentation comments to dcnm-san-docfeedback@cisco.com

Security Configuration Guide, Cisco DCNM for SAN

Cisco DCNM for SAN, Release 5.x
July 2011

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

Text Part Number: OL-24948-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Security Configuration Guide, Cisco DCNM for SAN
© 2011 Cisco Systems, Inc. All rights reserved.

Send documentation comments to dcnm-san-docfeedback@cisco.com

CONTENTS

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

New and Changed Information

As of Cisco DCNM Release 5.2, Cisco Fabric Manager and Cisco Data Center Network Manager for LAN are merged into one unified product called Cisco Data Center Network Manager (DCNM) that can manage both LAN and SAN environments. As a part of this product merger, the name Cisco DCNM for SAN replaces the name Cisco Fabric Manager.

The following documentation changes support the merged Cisco DCNM product:

- Cisco DCNM product documentation for Cisco DCNM Release 5.2 is retitled with the name Cisco DCNM for LAN.
- Cisco Fabric Manager product documentation for Cisco DCNM Release 5.2 is retitled with the name Cisco DCNM for SAN.
- Cisco DCNM for SAN product documentation is now published to the Data Center Network Manager listing page on Cisco.com:
  http://www.cisco.com/en/US/products/ps9369/tsd_products_support_configure.html

  This URL is also the listing page for Cisco DCNM for LAN product documentation.
- Cisco Fabric Manager documentation for software releases earlier than Cisco DCNM Release 5.2, retains the name Cisco Fabric Manager and remains available at its current Cisco.com listing page:
  http://www.cisco.com/en/US/products/ps10495/tsd_products_support_configure.html

  You should continue to use the Cisco Fabric Manager documentation if you are using a release of Cisco Fabric Manager software that is earlier than Cisco DCNM Release 5.2.
- The name DCNM-SAN is used in place of Cisco DCNM for SAN in the user interface of Cisco Data Center Network Manager; likewise, the name DCNM-LAN is used in place of Cisco DCNM for LAN in the user interface. To match the user interface, the product documentation also uses the names DCNM-SAN and DCNM-LAN.
- The following new publications support both Cisco DCNM for LAN and DCNM for SAN, and address the new licensing model, the new installation process, and the new features of Cisco DCNM:
  - *Cisco DCNM Installation and Licensing Guide*
  - *Cisco DCNM Release Notes*
- For a complete list of Cisco DCNM documentation, see the “Related Documentation” section in the Preface.

As of MDS NX-OS Rease 4.2(1), software configuration information is available in new feature-specific configuration guides for the following information:

- System management
- Interfaces
- Fabric

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Quality of service
- Security
- IP services
- High availability and redundancy

The information in these new guides previously existed in the *Cisco MDS 9000 Family CLI Configuration Guide* and in the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*. Those configuration guides remain available on Cisco.com and should be used for all software releases prior to Fabric Manager Release 5.0(1a). Each guide addresses the features introduced in or available in a particular release. Select and view the configuration guide that pertains to the software installed in your switch.

Some information from the *Cisco MDS 9000 Family CLI Configuration Guide* and the *Cisco MDS 9000 Family Fabric Manager Configuration Guide* now appears in the following guides that are common among products that run the \/ operating system:

- *Cisco NX-OS Licensing Guide* – Explains the licensing model and describes the feature licenses.
- *Cisco NX-OS Fundamentals Guide* – Describes the switch setup utility and includes general CLI, file system, and configuration information.

For a complete list of document titles, see the list of Related Documentation in the “Preface.”

To find additional information about Cisco DCNM for SAN Release 5.x, see the *Cisco MDS 9000 Family Release Notes* available at the following Cisco Systems website:

http://www.cisco.com/en/US/products/ps5989/prod_release_notes_list.htm

About This Guide

The information in the new *Cisco Fabric Manager Security Configuration Guide* previously existed in Part 5: Security of the *Cisco MDS 9000 Family Fabric Manager Configuration Guide.*

Table 1 lists the New and Changed features for this guide, starting with MDS NX-OS Release 4.2(1).

Table 1 New and Changed Features

Feature	New or Changed Topics	Changed in Release	Where Documented
AAA/LDAP	LDAP Server is supported	5.2(1)	Chapter 4, “Configuring Security Features on an External AAA Server”
Changes to SSH	Boot Mode SSH, Passwordfree File copy and SSH	5.0(1a)	Chapter 3, “Configuring Users and Common Role.”
Role Distributions	Enabling Role-Based Configuration Distribution	5.0(1a)	Chapter 3, “Role Distributions.”
Switch AAA Functionalities	Configuring Fallback Mechanism for Authentication, Configuring AAA Server Monitoring Parameters Globally	5.0(1a)	Chapter 4, “Configuring Security Features on an External AAA Server.”
Configuring TACACS+ Server Monitoring Parameters	CHAP Authentication	5.0(1a)	Chapter 4, “Configuring Security Features on an External AAA Server.”
OTP Authentication	One Time Password Support	5.0(1a)	Chapter 4, “Configuring Security Features on an External AAA Server.”

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 1 New and Changed Features

Feature	New or Changed Topics	Changed in Release	Where Documented
Creating Users Guidelines	Caution has been changed	5.0(1a)	Chapter 3, "Configuring Users and Common Role."
Merge Guidelines for RADIUS and TACACS+ Configurations	TACACS test parameters have to be distributed via CFS a note has been changed	5.0(1a)	Chapter 4, "Configuring Security Features on an External AAA Server."
AAA Service Configuration Options	Note has been changed	5.0(1a)	Chapter 4, "AAA Service Configuration Options."

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

Preface

This preface describes the audience, organization, and conventions of the *Security Configuration Guide, Cisoc DCNM for SAN*. It also provides information on how to obtain related documentation.

Audience

This guide is for experienced network administrators who are responsible for configuring and maintaining the Cisco MDS 9000 Family of multilayer directors and fabric switches.

Document Organization

This document is organized as folows:

Chapter	Title	Description
Chapter 1	Security Overview	Provides an overview of the security features supported by the Cisco MDS 9000 Family NX-OS software.
Chapter 2	Configuring FIPS	Describes the configuration guidelines for FIPS and also how to enable FIPS mode and how to conduct FIPS self-tests.
Chapter 3	Configuring Users and Common Role	Describes how to configure users and common roles.
Chapter 4	Configuring Security Features on an External AAA Server	Describes the AAA parameters, user profiles, and RADIUS authentication security options provided in all switches in the Cisco MDS 9000 Family and provides configuration information for these options.
Chapter 5	Configuring IPv4 and IPv6 Access Control Lists	Describes the IPv4 static routing feature and its use to route traffic between VSANs.
Chapter 6	Configuring Certificate Authorities and Digital Certificates	Describes how to interoperate with Certificate Authorities (CAs) and use digital certificates for secure, scalable communication.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Chapter	Title	Description
Chapter 7	Configuring IPsec Network Security	Provides details on the digital certificates, IP Security Protocol (IPsec) open standards, and the Internet Key Exchange (IKE) protocol that it uses to handle protocol and algorithm negotiation.
Chapter 8	Configuring FC-SP and DHCHAP	Describes the DHCHAP protocol, an FC-SP protocol, that provides authentication between Cisco MDS 9000 Family switches and other devices.
Chapter 9	Configuring Port Security	Provides details on port security features that can prevent unauthorized access to a switch port in the Cisco MDS 9000 Family.
Chapter 10	Configuring Fabric Binding	Describes the fabric binding security feature for VSANs, which ensures that ISLs are only enabled between specific switches.
Chapter 11	Configuring Cisco TrustSec Fibre Channel Link Encryption	Describes how the switch allows IP hosts to access Fibre Channel storage using the iSCSI protocol.

Document Conventions

Command descriptions use these conventions:

boldface font	Commands and keywords are in boldface.
italic font	Arguments for which you supply values are in italics.
[ ]	Elements in square brackets are optional.
[x | y | z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.

Screen examples use these conventions:

`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
`< >`	Nonprinting characters, such as passwords, are in angle brackets.
`[ ]`	Default responses to system prompts are in square brackets.
`!, #`	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:

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

Send documentation comments to dcnm-san-docfeedback@cisco.com

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

Related Documentation

The documentation set for the Cisco MDS 9000 Family includes the following documents. To find a document online, use the Cisco MDS NX-OS Documentation Locator at:

http://www.cisco.com/en/US/docs/storage/san_switches/mds9000/roadmaps/doclocater.htm

Release Notes

- *Cisco MDS 9000 Family Release Notes for Cisco MDS NX-OS Releases*
- *Cisco MDS 9000 Family Release Notes for MDS SAN-OS Releases*
- *Cisco MDS 9000 Family Release Notes for Cisco MDS 9000 EPLD Images*
- *Cisco DCNM Release Notes*

Regulatory Compliance and Safety Information

- *Regulatory Compliance and Safety Information for the Cisco MDS 9000 Family*

Compatibility Information

- *Cisco Data Center Interoperability Support Matrix*
- *Cisco MDS 9000 NX-OS Hardware and Software Compatibility Information and Feature Lists*
- *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*

Hardware Installation

- *Cisco MDS 9500 Series Hardware Installation Guide*
- *Cisco MDS 9200 Series Hardware Installation Guide*
- *Cisco MDS 9100 Series Hardware Installation Guide*
- *Cisco MDS 9124 and Cisco MDS 9134 Multilayer Fabric Switch Quick Start Guide*

Software Installation and Upgrade

- *Cisco MDS 9000 NX-OS Software Upgrade and Downgrade Guide*

Send documentation comments to dcnm-san-docfeedback@cisco.com

Cisco NX-OS

- *Cisco MDS 9000 Family NX-OS Licensing Guide*
- *Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide*
- *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Quality of Service Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Security Configuration Guide*
- *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Intelligent Storage Services Configuration Guide*
- *Cisco MDS 9000 Family NX-OS High Availability and Redundancy Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Inter-VSAN Routing Configuration Guide*
- *Cisco MDS 9000 Family Cookbook for Cisco MDS SAN-OS*

Cisco DCNM-SAN

- *Cisco DCNM Fundamentals Guide, Release 5.x*
- *System Management Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *Interfaces Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *Fabric Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *Quality of Service Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *Security Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *IP Services Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *Intelligent Storage Services Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *High Availability and Redundancy Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *Inter-VSAN Routing Configuration Guide, Cisco DCNM for SAN, Release 5.x*
- *SMI-S and Web Services Programming Guide, Cisco DCNM for SAN, Release 5.x*

Command-Line Interface

- *Cisco MDS 9000 Family Command Reference*

Intelligent Storage Networking Services Configuration Guides

- *Cisco MDS 9000 Family I/O Acceleration Configuration Guide*
- *Cisco MDS 9000 Family SANTap Deployment Guide*
- *Cisco MDS 9000 Family Data Mobility Manager Configuration Guide*
- *Cisco MDS 9000 Family Storage Media Encryption Configuration Guide*

Send documentation comments to dcnm-san-docfeedback@cisco.com

Troubleshooting and Reference

- *Cisco MDS 9000 Family and Nexus 7000 Series System Messages Reference*
- *Cisco MDS 9000 Family SAN-OS Troubleshooting Guide*
- *Cisco MDS 9000 Family NX-OS MIB Quick Reference*
- *Cisco DCNM for SAN Database Schema Reference*

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

- Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

CHAPTER 1

Security Overview

The Cisco MDS 9000 NX-OS software supports advanced security features that provide security within a Storage Area Network (SAN). These features protect your network against deliberate or unintentional disruptions from internal or external threats.

This chapter includes the following sections:

FIPS

The Federal Information Processing Standards (FIPS) Publication 140-2, *Security Requirements for Cryptographic Modules*, details the U.S. government requirements for cryptographic modules. FIPS 140-2 specifies that a cryptographic module should be a set of hardware, software, firmware, or some combination that implements cryptographic functions or processes, including cryptographic algorithms and, optionally, key generation, and is contained within a defined cryptographic boundary. FIPS specifies certain crypto algorithms as secure, and it also identifies which algorithms should be used if a cryptographic module is to be called FIPS compliant.

For more information on configuring FIPS, see Chapter 2, “Configuring FIPS.”

Send documentation comments to dcnm-san-docfeedback@cisco.com

Users and Common Roles

Role-based authorization limits access to switch operations by assigning users to roles. All management access within the Cisco MDS 9000 Family is based upon roles. Users are restricted to performing the management operations that are explicitly permitted, by the roles to which they belong.

For information on configuring users and common roles, see Chapter 3, "Configuring Users and Common Role."

RADIUS and TACACS+

The authentication, authorization, and accounting (AAA) feature verifies the identity of, grants access to, and tracks the actions of users managing a switch. All Cisco MDS 9000 Family switches use RADIUS and TACACS+ protocols to provide solutions using remote AAA servers. This security feature provides a centralized user account management capability for AAA servers.

AAA uses security protocols to administer its security functions. If your router or access server is acting as a network access server, then the communication between your network access server and the RADIUS or TACACS+ security server is through AAA.

The chapters in this guide describe the following features:

- Switch management—A management security system that provides security to all management access methods, including the command-line interface (CLI) or Simple Network Management Protocol(SNMP).
- Switch AAA functionalities—A function by which you can configure AAA switch functionalities on any switch in the Cisco MDS 9000 Family, using the command-line interface (CLI) or Simple Network Management Protocol (SNMP).
- RADIUS—A distributed client and server system implemented through AAA that secures networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco routers and send authentication requests to a central RADIUS server that contains all user authentication and network service access information.
- TACACS+—A security application implemented through AAA that provides a centralized validation of users who are attempting to gain access to a router or network access server. TACACS+ services are maintained in a database on a TACACS+ daemon that typically runs on a UNIX or Windows NT workstation. TACACS+ provides for separate and modular authentication, authorization, and accounting facilities.

LDAP

The Lightweight Directory Access Protocol (LDAP) provides centralized validation of users attempting to gain access to a Cisco NX-OS device. LDAP services are maintained in a database on an LDAP daemon typically running on a UNIX or Windows NT workstation. You must have access to and must configure an LDAP server before the configured LDAP features on your Cisco NX-OS device are available.

LDAP provides for separate authentication and authorization facilities. LDAP allows for a single access control server (the LDAP daemon) to provide each service authentication and authorization independently. Each service can be connected to its own database to take advantage of other services available on that server or on the network, depending on the capabilities of the daemon.

Send documentation comments to dcnm-san-docfeedback@cisco.com

The LDAP client and server protocol uses TCP (TCP port 389) for transport requirements. Cisco NX-OS devices provide centralized authentication using the LDAP protocol.

For information on configuring RADIUS and TACACS+, see Chapter 4, "Configuring Security Features on an External AAA Server."

IP ACLs

IP access control lists (ACLs) provide basic network security on the out-of-band management Ethernet interface and the in-band IP management Interface. The Cisco MDS 9000 Family switches use IP ACLs to restrict traffic from unknown and untrusted sources and restrict network use based on user identity or device type.

For information on configuring IP ACLs, see Chapter 5, "Configuring IPv4 and IPv6 Access Control Lists".

PKI

The Public Key Infrastructure (PKI) allows an MDS 9000 switch to obtain and use digital certificates for secure communication in the network. PKI support provides manageability and scalability for applications, such as IPsec, IKE, and SSH, that support digital certificates.

For information on configuring PKI, see Chapter 6, "Configuring Certificate Authorities and Digital Certificates."

IPsec

IP Security (IPsec) protocol is a framework of open standards by the Internet Engineering Task Force (IETF) that provides data confidentiality, data integrity, and data origin authentication between participating peers. IPsec provides security services at the IP layer, including protecting one or more data flows between a pair of hosts, a pair of security gateways, or a security gateway and a host.

For information on configuring IPsec, see Chapter 7, "Configuring IPsec Network Security."

FC-SP and DHCHAP

Fibre Channel Security Protocol (FC-SP) capabilities provide switch to switch and hosts to switch authentication to overcome security challenges for enterprise-wide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) is an FC-SP protocol that provides authentication between Cisco MDS 9000 Family switches and other devices. DHCHAP consists of the CHAP protocol combined with the Diffie-Hellman exchange.

With FC-SP, switches, storage devices, and hosts are able to prove their identity through a reliable and manageable authentication mechanism. With FC-SP, Fibre Channel traffic can be secured on a frame-by-frame basis to prevent snooping and hijacking, even over untrusted links. A consistent set of policies and management actions are propagated through the fabric to provide a uniform level of security across the entire fabric.

For more information on configuring FS-SP and DHCHAP, see Chapter 8, "Configuring FC-SP and DHCHAP."

Send documentation comments to dcnm-san-docfeedback@cisco.com

Port Security

The port security feature prevents unauthorized access to a switch port by binding specific world-wide names (WWNs) that have access to one or more given switch ports.

When port security is enabled on a switch port, all devices connecting to that port must be in the port security database and must be listed in the database as bound to a given port. If both of these criteria are not met, the port will not achieve an operationally active state and the devices connected to the port will be denied access to the SAN.

For information on configuring port security, see Chapter 9, “Configuring Port Security.”

Fabric Binding

The fabric binding feature ensures Inter-Switch Links (ISLs) are enabled only between specified switches in the fabric binding configuration. This feature helps prevent unauthorized switches from joining the fabric or disrupting the current fabric operations. This feature uses the Exchange Fabric Membership Data (EEMD) protocol to ensure that the list of authorized switches is identical in all of the switches in a fabric.

For information on configuring fabric binding, see Chapter 10, “Configuring Fabric Binding.”

TrustSec Fibre Channel Link Encryption

Cisco TrustSec Fibre Channel Link Encryption is an extension of the Fibre Channel-Security Protocol (FC-SP) feature and uses the existing FC-SP architecture to provide integrity and confidentiality of transactions. Encryption is added to the peer authentication capability to provide security and prevent unwanted traffic interception. Peer authentication is implemented according to the FC-SP standard using the Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) protocol.

For information on configuring TrustSec Fibre Channel Link Encryption, see Chapter 11, “Configuring Cisco TrustSec Fibre Channel Link Encryption.”

Send documentation comments to dcnm-san-docfeedback@cisco.com

CHAPTER 2

Configuring FIPS

The Federal Information Processing Standards (FIPS) Publication 140-2, *Security Requirements for Cryptographic Modules*, details the U.S. government requirements for cryptographic modules. FIPS 140-2 specifies that a cryptographic module should be a set of hardware, software, firmware, or some combination that implements cryptographic functions or processes, including cryptographic algorithms and, optionally, key generation, and is contained within a defined cryptographic boundary.

FIPS specifies certain crypto algorithms as secure, and it also identifies which algorithms should be used if a cryptographic module is to be called FIPS compliant.

Note Cisco MDS SAN-OS Release 3.1(1) and NX-OS Release 4.1(1b) or later implements FIPS features and is currently in the certification process with the U.S. government, but it is not FIPS compliant at this time.

This chapter includes the following topics:

Information About FIPS Self-Tests

A cryptographic module must perform power-up self-tests and conditional self-tests to ensure that it is functional.

Note FIPS power-up self-tests automatically run when FIPS mode is enabled by entering the **fips mode enable** command. A switch is in FIPS mode only after all self-tests are successfully completed. If any of the self-tests fail, then the switch is rebooted.

Power-up self-tests run immediately after FIPS mode is enabled. A cryptographic algorithm test using a known answer must be run for all cryptographic functions for each FIPS 140-2-approved cryptographic algorithm implemented on the Cisco MDS 9000 Family.

Using a known-answer test (KAT), a cryptographic algorithm is run on data for which the correct output is already known, and then the calculated output is compared to the previously generated output. If the calculated output does not equal the known answer, the known-answer test fails.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Conditional self-tests must be run when an applicable security function or operation is invoked. Unlike the power-up self-tests, conditional self-tests are executed each time their associated function is accessed.

Conditional self-tests include the following:

- Pair-wise consistency test—This test is run when a public-private key-pair is generated.
- Continuous random number generator test—This test is run when a random number is generated.

Both of these tests automatically run when a switch is in FIPS mode.

Guidelines and Limitations

Follow these guidelines before enabling FIPS mode:

- Make your passwords a minimum of eight characters in length.
- Disable Telnet. Users should log in using SSH only.
- Disable remote authentication through RADIUS/TACACS+. Only users local to the switch can be authenticated.
- Disable SNMP v1 and v2. Any existing user accounts on the switch that have been configured for SNMPv3 should be configured only with SHA for authentication and AES/3DES for privacy.
- Disable VRRP.
- Delete all IKE policies that either have MD5 for authentication or DES for encryption. Modify the policies so they use SHA for authentication and 3DES/AES for encryption.
- Delete all SSH Server RSA1 key-pairs.

Enabling FIPS Mode

Detailed Steps

To enable FIPS mode using DCNM-SAN, follow these steps:

Step 1 Expand **Switches** from the Physical Attributes pane. Expand **Security** and then select **FIPS**.

You see the FIPS activation details in the Information pane as shown in Figure 2-1.

Figure 2-1 FIPS Activation in DCNM-SAN

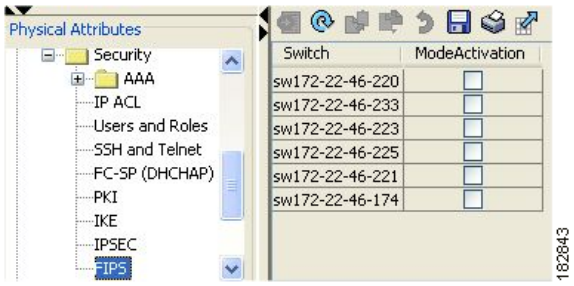

Step 2 Check the **ModeActivation** check box next to the switch for which you want to enable FIPS mode.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 3 Click **Apply Changes** to commit and distribute these changes.

Step 4 Click **Undo Changes** to discard any unsaved changes.

To enable FIPS mode using Device Manager, follow these steps:

Step 1 Choose **Physical > System** or right-click and select **Configure.**

You see the System dialog box as shown in Figure 2-2.

Figure 2-2 System Dialog Box

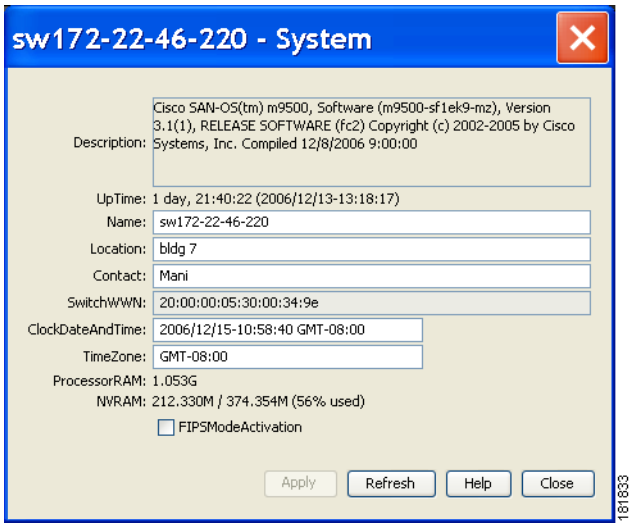

Step 2 Check the **FIPSModeActivation** check box to enable FIPS mode on the selected switch.

Step 3 Click **Apply** to save the changes.

Step 4 Click **Close** to close the dialog box.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field Descriptions for FIPS

FIPS

Field	Description
ModeActivation	To enable/disable FIPS mode on the device. FIPS 140-2 is a set of security requirements for cryptographic modules and it details the U.S. Government requirements for cryptographic modules. A module will comprise both hardware and software, eg a datacenter switching or routing module. The module is said to be in FIPS enabled mode when a request is recieved to enable the FIPS mode and a set of self-tests are successfully run in response to the request. If the self-tests fail, then an appropriate error is returned.

Send documentation comments to dcnm-san-docfeedback@cisco.com

CHAPTER 3

Configuring Users and Common Role

The CLI and SNMP use common roles in all switches in the Cisco MDS 9000 Family. You can use the CLI to modify a role that was created using SNMP 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, DCNM for SAN (DCNM-SAN or Device Manager) and vice versa.

This chapter includes the following topics:

Information About Role-Based Authorization

Switches in the Cisco MDS 9000 Family perform authentication based on roles. Role-based authorization limits access to switch operations by assigning users to roles. This kind of authentication restricts you to management operations based on the roles to which you have been added.

When you execute a command, perform command completion, or obtain context-sensitive help, the switch software allows the operation to progress if you have permission to access that command.

This section includes the following topics:

Send documentation comments to dcnm-san-docfeedback@cisco.com

About Roles

Each role can contain multiple users and each user can be part of multiple roles. For example, if role1 users are only allowed access to configuration commands, and role2 users are only allowed access to **debug** commands, then if Joe belongs to both role1 and role2, he can access configuration as well as **debug** commands.

Note If you belong to multiple roles, you can execute a union of all the commands permitted by these roles. Access to a command takes priority over being denied access to a command. For example, suppose you belong to a TechDocs group and you were denied access to configuration commands. However, you also belong to the engineering group and have access to configuration commands. In this case, you will have access to configuration commands.

Tip Any role, when created, does not allow access to the required commands immediately. The administrator must configure appropriate rules for each role to allow access to the required commands.

Rules and Features for Each Role

Up to 16 rules can be configured for each role. These rules reflect what CLI commands are allowed. The user-specified rule number determines the order in which the rules are applied. For example, rule 1 is applied before rule 2, which is applied before rule 3, and so on. A user not belonging to the network-admin role cannot perform commands related to roles.

For example, if user A is permitted to perform all **show** CLI commands, user A cannot view the output of the **show role** CLI command if user A does not belong to the network-admin role.

A rule specifies operations that can be performed by a specific role. Each rule consists of a rule number, a rule type (permit or deny), a CLI command type (for example, **config**, **clear**, **show**, **exec**, **debug**), and an optional feature name (for example, FSPF, zone, VSAN, fcping, or interface).

Note In this case, **exec** CLI commands refer to all commands in the EXEC mode that are not included in the **show**, **debug**, and **clear** CLI command categories.

Send documentation comments to dcnm-san-docfeedback@cisco.com

About the VSAN Policy

Configuring the VSAN policy requires the ENTERPRISE_PKG license (For more information, see *Cisco MDS 9000 Family NX-OS Licensing Guide*).

You can configure a role so that it only allows tasks to be performed for a selected set of VSANs. By default, the VSAN policy for any role is permit, which allows tasks to be performed for all VSANs. You can configure a role that only allows tasks to be performed for a selected set of VSANs. To selectively allow VSANs for a role, set the VSAN policy to deny, and then set the configuration to permit or the appropriate VSANs.

Note Users configured in roles where the VSAN policy is set to deny cannot modify the configuration for E ports. They can only modify the configuration for F or FL ports (depending on whether the configured rules allow such configuration to be made). This is to prevent such users from modifying configurations that may impact the core topology of the fabric.

Tip Roles can be used to create VSAN administrators. Depending on the configured rules, these VSAN administrators can configure MDS features (for example, zone, fcdomain, or VSAN properties) for their VSANs without affecting other VSANs. Also, if the role permits operations in multiple VSANs, then the VSAN administrators can change VSAN membership of F or FL ports among these VSANs.

Users belonging to roles in which the VSAN policy is set to deny are referred to as VSAN-restricted users.

Role Distributions

Role-based configurations use the Cisco Fabric Services (CFS) infrastructure to enable efficient database management, and to provide a single point of configuration for the entire fabric.

The following configurations are distributed:

- Role names and descriptions
- List of rules for the roles
- VSAN policy and the list of permitted VSANs

About Role Databases

Role-based configurations use two databases to accept and implement configurations.

- Configuration database—The running database currently enforced by the fabric.
- Pending database—Your subsequent configuration changes are stored in the pending database. If you modify the configuration, you need to commit or discard the pending database changes to the configuration database. The fabric remains locked during this period. Changes to the pending database are not reflected in the configuration database until you commit the changes.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Locking the Fabric

The first action that modifies the database creates the pending database and locks the feature in the entire fabric. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database along with the first change.

About Common Roles

The CLI and SNMP in all switches in the Cisco MDS 9000 Family use common roles. You can use SNMP to modify a role that was created using the CLI and vice versa (see Figure 3-1).

Figure 3-1 Common Roles

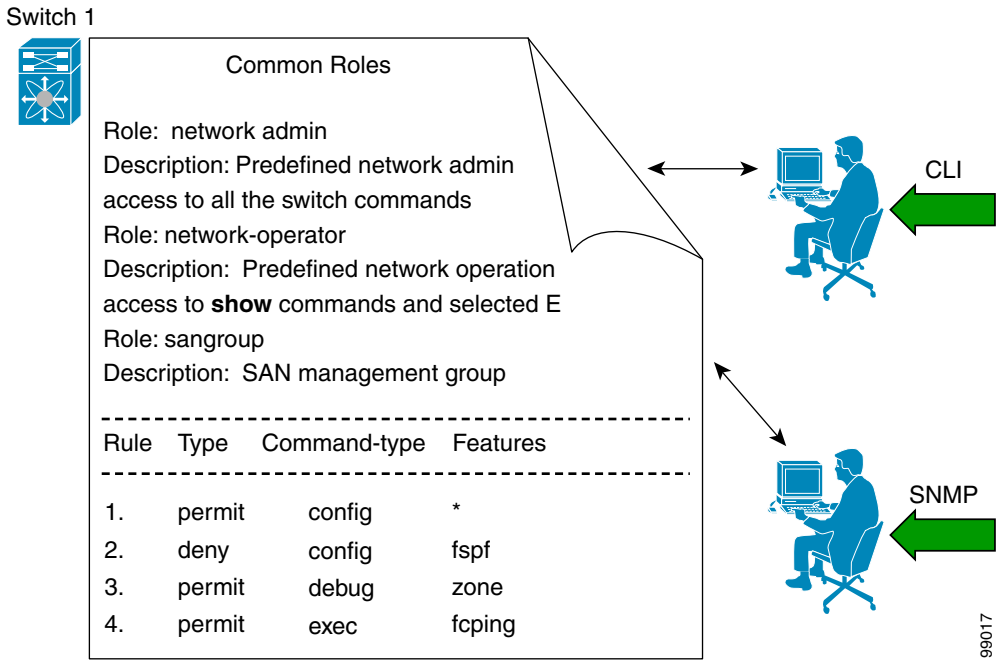

Each role in SNMP is the same as a role created or modified through the CLI (see the "Information About Role-Based Authorization" section on page 3-1).

Each role can be restricted to one or more VSANs as required.

You can create new roles or modify existing roles using SNMP or the CLI.

- SNMP—Use the CISCO-COMMON-ROLES-MIB to configure or modify roles. Refer to the *Cisco MDS 9000 Family MIB Quick Reference*.
- CLI—Use the **role name** command.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Mapping of CLI Operations to SNMP

SNMP has only three possible operations: GET, SET, and NOTIFY. The CLI has five possible operations: DEBUG, SHOW, CONFIG, CLEAR, and EXEC.

Note NOTIFY does not have any restrictions like the syslog messages in the CLI.

Table 3-1 explains how the CLI operations are mapped to the SNMP operations.

Table 3-1 CLI Operation to SNMP Operation Mapping

CLI Operation	SNMP Operation
DEBUG	Ignored
SHOW	GET
CONFIG	SET
CLEAR	SET
EXEC	SET

Creating Users Guidelines

The passphrase specified in the **snmp-server user** option and the password specified **username** option are synchronized.

By default, the user account does not expire unless you explicitly configure it to expire. The **expire** option determines the date on which the user account is disabled. The date is specified in the YYYY-MM-DD format.

When creating users, note the following guidelines:

- You can configure up to a maximum of 256 users on a switch.
- The following words are reserved and cannot be used to configure users: bin, daemon, adm, lp, sync, shutdown, halt, mail, news, uucp, operator, games, gopher, ftp, nobody, nscd, mailnull, rpc, rpcuser, xfs, gdm, mtsuser, ftpuser, man, and sys.
- User passwords are not displayed in the switch configuration file.
- If a password is trivial (short, easy-to-decipher), your password configuration is rejected. Be sure to configure a strong password as shown in the sample configuration. Passwords are case-sensitive. “admin” is no longer the default password for any Cisco MDS 9000 Family switch. You must explicitly configure a strong password.

Caution Cisco MDS NX-OS supports user names that are created with alphanumeric characters or specific special characters (+ [plus], = [equal], _ [underscore], - [hyphen], \ [backslash], and . [period]) whether created remotely (using TACACS+ or RADIUS) or locally. Local user names cannot be created with any special characters (apart from those specified). If a non-supported special character user name exists on an AAA server, and is entered during login, then the user is denied access.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Characteristics of Strong Passwords

A strong password has the following characteristics:

- At least eight characters long
- Does not contain many consecutive characters (such as “abcd”)
- Does not contain many repeating characters (such as “aaabbb”)
- Does not contain dictionary words
- Does not contain proper names
- Contains both upper- and lower-case characters
- Contains numbers

The following are examples of strong passwords:

- If2CoM18
- 2004AsdfLkj30
- Cb1955S21

About SSH

SSH provides secure communications to the Cisco NX-OS CLI. You can use SSH keys for the following SSH options:

- SSH2 using RSA
- SSH2 using DSA

Boot Mode SSH

Due to the increasing emphasis on security and security-related issues, the **ssh** command in this release runs in the Boot mode. SSH is a preferred and more secure method of data exchange over the network because it communicates over the secure channel, and the data is encrypted before sending on the channel.

Example 3-1 shows how to use the **ssh** command to conncet to a remote server from any switch.

Example 3-1 Connecting a Remote Server from Any Switch

```
switch# ssh admin @ hostname
```

SSH Authentication Using Digital Certificates

SSH authentication on the Cisco MDS 9000 Family switches provide X.509 digital certificate support for host authentication. An X.509 digital certificate is a data item that vouches for the origin and integrity of a message. It contains encryption keys for secured communications and is “signed” by a trusted certification authority (CA) to verify the identity of the presenter. The X.509 digital certificate support provides either DSA or RSA algorithms for authentication.

Send documentation comments to dcnm-san-docfeedback@cisco.com

The certificate infrastructure uses the first certificate that supports the Secure Socket Layer (SSL) and is returned by the security infrastructure, either through query or notification. Verification of certificates is successful if the certificates are from any of the trusted CAs.

You can configure your switch for either SSH authentication using an X.509 certificate or SSH authentication using a Public Key Certificate, but not both. If either of them is configured and the authentication fails, you will be prompted for a password.

For more information on CAs and digital certificates, see Chapter 6, "Configuring Certificate Authorities and Digital Certificates."

Passwordless File copy and SSH

Secure Shell (SSH) public key authentication can be used to achieve password-free logins. SCP and SFTP uses SSH in the background, which enables these copy protocols to be used for a password-free copy with public key authentication. The NX-OS version only supports the SCP and STFP client functionality.

You can create an RSA and DSA identity that can be used for authentication with SSH. The identity consists of two parts: public and private keys. The public and the private keys are generated by the switch or can be generated externally and imported to the switch. For import purposes, the keys should be in OPENSSH format.

To use the key on a host machine hosting an SSH server, you must transfer the public key file to the machine and add the contents of it to the authorized_keys file in your SSH directory (for example, $HOME/.ssh) on the server. For the import and export of private keys, the key is protected by encryption. You are asked to enter the passphrase for the keys. If you enter a passphrase, the private key is protected by encryption. If you leave the password field blank, the key will not be encrypted.

If you need to copy the keys to another switch, you will have to export the keys out of the switch to a host machine, and then import the keys to other switches from that machine.

The key files are persistent across reload.

Guidelines and Limitations

Fabric merge does not modify the role database on a switch. If two fabrics merge, and the fabrics have different role databases, the software generates an alert message.

See the "Merge Guidelines for RADIUS and TACACS+ Configurations" section on page 4-15 for detailed concepts.

- Verify that the role database is identical on all switches in the entire fabric.
- Be sure to edit the role database on any switch to the desired database and then commit it. This synchronizes the role databases on all the switches in the fabric.

Default Settings

Table 3-2 lists the default settings for all switch security features in any switch.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 3-2 Default Switch Security Settings

Parameters	Default
Roles in Cisco MDS Switches	Network operator (network-operator)
AAA configuration services	Local
Authentication port	1812
Accounting port	1813
Preshared key communication	Clear text
RADIUS server time out	1 (one) second
RADIUS server retries	Once
TACACS+	Disabled
TACACS+ servers	None configured
TACACS+ server timeout	5 seconds
AAA server distribution	Disabled
VSAN policy for roles	Permit
User account	No expiry (unless configured)
Password	None
Password-strength	Enabled
Accounting log size	250 KB
SSH service	Enabled
Telnet service	Disabled

Configuring Users and Common Role

This section includes the following topics:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring Roles and Profiles

Detailed Steps

Note Only users belonging to the network-admin role can create roles.

To create an additional role or to modify the profile for an existing role using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** from the Physical Attributes pane.

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

Step 3 Click **Create Row** to create a role in DCNM-SAN.

Step 4 Select the switches on which to configure a role.

Step 5 Enter the name of the role in the Name field.

Step 6 Enter the description of the role in the Description field.

Step 7 (Optional) Check the **Enable** check box to enable the VSAN scope and enter the list of VSANs in the Scope field to which you want to restrict this role.

Step 8 Click **Create** to create the role.

Note Device Manager automatically creates six roles that are required for Device Manager to display a view of a switch. These roles are **system**, **snmp**, **module**, **interface**, **hardware**, and **environment**.

Deleting Common Roles

Detailed Steps

To delete a common role using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** from the Physical Attributes pane.

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

Step 3 Click the role you want to delete.

Step 4 Click **Delete Row** to delete the common role.

Step 5 Click **Yes** to confirm the deletion or **No** to cancel it.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Modifying Rules

Detailed Steps

To modify the rules for an existing role using Device Manager, follow these steps:

Step 1 Choose **Security** > **Roles**.

Step 2 Click the role for which you want to edit the rules.

Step 3 Click **Rules** to view the rules for the role.

You see the Edit Role Rules dialog box.

Step 4 Edit the rules you want to enable or disable for the common role.

Step 5 Click **Apply** to apply the new rules.

Rule 1 is applied first, which permits, for example, sangroup users access to all **config** CLI commands. Rule 2 is applied next, denying FSPF configuration to sangroup users. As a result, sangroup users can perform all other **config** CLI commands, except the **fspf** CLI configuration commands.

Note The order of rule placement is important. If you had swapped these two rules and issued the **deny config feature fspf** rule first and issued the **permit config** rule next, you would be allowing all sangroup users to perform all configuration commands because the second rule globally overrode the first rule.

Modifying the VSAN Policy

Detailed Steps

To modify the VSAN policy for an existing role using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** from the Physical Attributes pane.

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

Step 3 Check the **Scope Enable** check box if you want to enable the VSAN scope and restrict this role to a subset of VSANs.

Step 4 Enter the list of VSANs in the Scope VSAN Id List field that you want to restrict this role to.

Step 5 Click **Apply Changes** to save these changes.

Committing Role-Based Configuration Changes

If you commit the changes made to the pending database, the configuration is committed to all the switches in the fabric. On a successful commit, the configuration change is applied throughout the fabric and the lock is released. The configuration database now contains the committed changes and the pending database is now cleared.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To commit role-based configuration changes using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** in the Physical Attributes pane.

Step 2 Click the **Roles CFS** tab in the Information pane.

Step 3 Set the Global drop-down menu to **enable** to enable CFS.

Step 4 Click the **Apply Changes** icon to save this change.

Step 5 Set the Config Action drop-down menu to **commit** to commit the roles using CFS.

Step 6 Click the **Apply Changes** icon to save this change.

Discarding Role-Based Configuration Changes

If you discard (abort) the changes made to the pending database, the configuration database remains unaffected and the lock is released.

Detailed Steps

To discard role-based configuration changes using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** in the Physical Attributes pane.

Step 2 Click the **Roles CFS** tab in the Information pane.

Step 3 Set the Config Action drop-down menu to **abort** to discard any uncommitted changes.

Step 4 Click the **Apply Changes** icon to save this change.

Enabling Role-Based Configuration Distribution

Detailed Steps

To enable role-based configuration distribution using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** in the Physical Attributes pane.

Step 2 Click the **Roles CFS** tab in the Information pane.

Step 3 Set the Global drop-down menu to **enable** to enable CFS distribution.

Step 4 Click the **Apply Changes** icon to save this change.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Clearing Sessions

Detailed Steps

To forcibly clear the existing role session in the fabric using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** in the Physical Attributes pane.

Step 2 Click the **Roles CFS** tab in the Information pane.

Step 3 Set the Config Action drop-down menu to **clear** to clear the pending database.

Step 4 Click the **Apply Changes** icon to save this change.

Caution Any changes in the pending database are lost when you clear a session.

Configuring Users

Before configuring users, make sure that you have configured roles to associate with the users that you are creating.

Note As of Cisco SAN-OS Release 3.1(2b), DCNM-SAN automatically checks whether encryption is enabled, which allows you to create users.

Detailed Steps

To configure a new user or to modify the profile of an existing user using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** from the Physical Attributes pane.

Step 2 Click the **Users** tab in the Information pane to see a list of users.

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

You see the Users - Create dialog box as shown in Figure 3-2.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 3-2 Users - Create Dialog Box

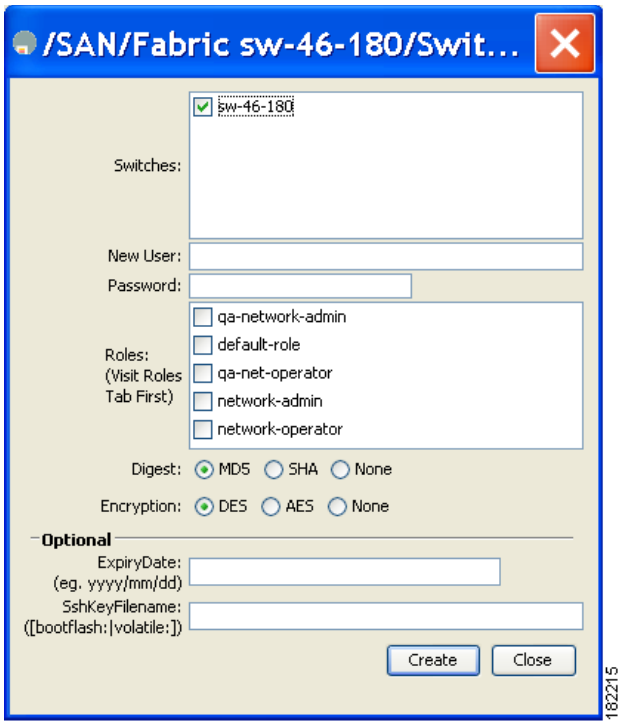

Step 4 (Optional) Alter the Switches check boxes to specify one or more switches.

Step 5 Enter the user name in the New User field.

Step 6 Enter the password for the user.

Step 7 Check the roles that you want to associate with this user.

See the "Rules and Features for Each Role" section on page 3-2.

Step 8 Select the appropriate option for the type of authentication protocol used. The default value is MD5.

Step 9 Select the appropriate option for the type of privacy protocol used. The default value is DES.

Step 10 (Optional) Enter the expiry date for this user.

Step 11 (Optional) Enter the SSH Key filename.

Step 12 Click **Create** to create the entry.

Deleting a User

Detailed Steps

To delete a user using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** from the Physical Attributes pane.

Step 2 Click the **Users** tab in the Information pane to see a list of users.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 3 Click the name of the user you want to delete.

Step 4 Click **Delete Row** to delete the selected user.

Step 5 Click **Apply Changes** to save this change.

Configuring SSH Services

A secure SSH connection with an RSA key is available as a default on all Cisco MDS 9000 Family switches. If you require a secure SSH connection with a DSA key, you need to disable the default SSH connection, Generate a DSA key and then enable the SSH connection (see the "Generating the SSH Server Key Pair" section on page 3-14).

Caution If you are logging in to a switch through SSH and you have issued the **aaa authentication login default none** command, you must enter one or more key strokes to log in. If you press the **Enter** key without entering at least one keystroke, your log in will be rejected.

This section includes the following topics:

Generating the SSH Server Key Pair

Ensure that you have an SSH server key pair with the appropriate version before enabling the SSH service. Generate the SSH server key pair according to the SSH client version used. The number of bits specified for each key pair ranges from 768 to 2048.

The SSH service accepts two types of key pairs for use by SSH version 2.

- The **dsa** option generates the DSA key pair for the SSH version 2 protocol.
- The **rsa** option generates the RSA keypair for the SSH version 2 protocol.

Caution If you delete all of the SSH keys, you cannot start a new SSH session.

Detailed Steps

To generate the SSH key pair using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security** and then select **SSH and Telnet**.

Step 2 Click the **Create Row** icon.

You see the SSH and Telnet Key - Create dialog box (see Figure 3-3).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 3-3 SSH and Telnet - Create Dialog Box

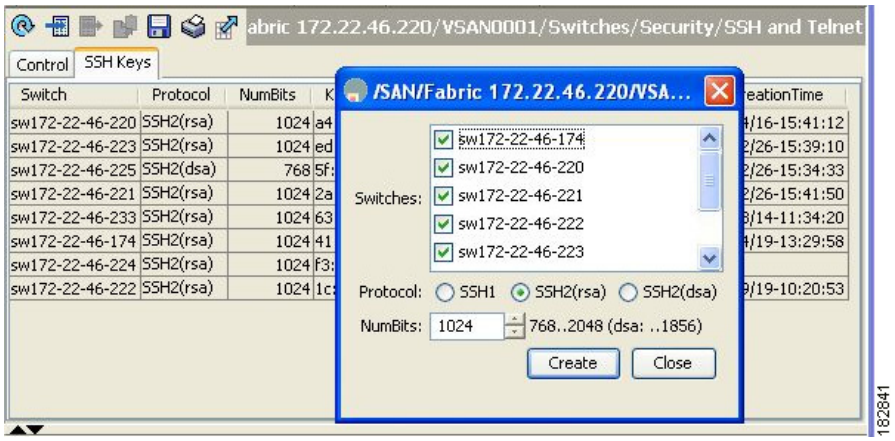

Step 3 Check the switches you want to assign to this SSH key pair.

Step 4 Choose the key pair option type from the listed Protocols. The listed protocols are SSH1, SSH2(rsa), and SSH2(dsa).

Step 5 Set the number of bits that will be used to generate the key pairs in the NumBits drop-down menu.

Step 6 Click **Create** to generate these keys.

Note 1856 DSA NumberKeys are not supported by switches that running Cisco MDS NX-OS software version 4.1(1) and later.

Overwriting a Generated Key Pair

If the SSH key pair option is already generated for the required version, you can force the switch to overwrite the previously generated key pair.

Detailed Steps

To overwrite the previously generated key pair using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security** and then select **SSH and Telnet**.

You see the configuration in the Information pane.

Step 2 Highlight the key that you want to overwrite and click **Delete Row**.

Step 3 Click the **Apply Changes** icon to save these changes.

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

You see the SSH and Telnet Key - Create dialog box.

Step 5 Check the switches you want to assign this SSH key pair.

Step 6 Choose the key pair option type from the Protocols radio buttons.

Step 7 Set the number of bits that will be used to generate the key pairs in the NumBits drop-down menu.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 8 Click **Create** to generate these keys.

Enabling SSH or Telnet Service

By default, the SSH service is enabled with the RSA key.

Detailed Steps

Note If you are logging in to a switch through SSH and you have issued the **aaa authentication login default none** CLI command, you must enter one or more key strokes to log in. If you press the **Enter** key without entering at least one keystroke, your log in will be rejected.

DCNM-SAN enables SSH automatically when you configure it.

To enable or disable SSH using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security** and then select **SSH and Telnet**.

Step 2 Select the **Control** tab and check an **SSH** check box or **Telnet** check box for each switch (see Figure 3-4).

Figure 3-4 Control Tab under SSH and Telnet

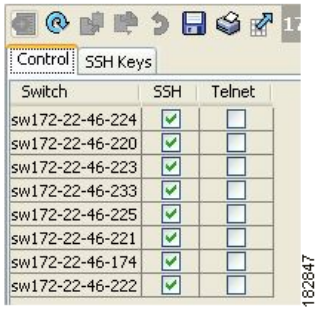

Step 3 Click the **Apply Changes** icon to save this change.

Changing Administrator Password Using DCNM-SAN

Detailed Steps

To change the administrator password in DCNM-SAN, follow these steps:

Step 1 Click the **Open** tab in the control panel.

Step 2 Choose the password field to change the password for an already existing user for the fabric.

Step 3 Click **Open** to open the fabric.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Note New password will be saved after the fabric is open. The user name and password fields are editable in the Fabric tab only after you unmanage the fabric.

Verifying Users and Common Role Configuration

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

This section includes the following topics:

Displaying Role-Based Information

The rules are displayed by rule number and are based on each role. All roles are displayed if the role name is not specified.

To view rules for a role using Device Manager, follow these steps:

Step 1 Click **Security > Roles.**

You see the Roles dialog box.

Step 2 Select a role name and click **Rules**.

You see the Rules dialog box.

Step 3 Click **Summary** to get a summarized view of the rules configured for this role.

Displaying Roles When Distribution is Enabled

To view the roles using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security** and then select **Users and Roles** in the Physical Attributes pane.

Step 2 Click the **Users** tab in the Information pane (see Figure 3-5).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 3-5 Roles CFS Tab

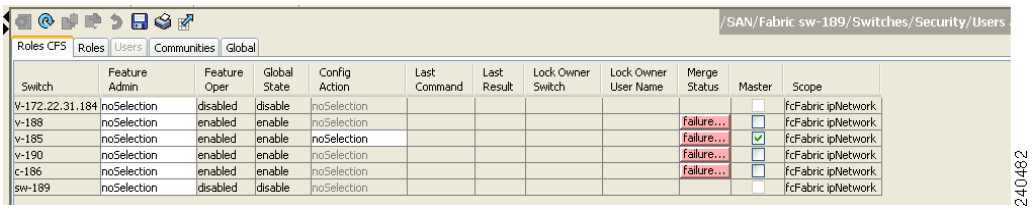

Step 3 Set the Config View As drop-down value to **pending** to view the pending database or set the Config View as drop-down menu to **running** to view the running database.

Step 4 Click **Apply Changes** to save this change.

Displaying User Account Information

To display information about configured user accounts using DCNM-SAN, follow these steps:

Step 1 Expand **Security** and then select **Users and Roles** in the Physical Attributes pane.

Step 2 Click the **Users** tab.

You see the list of SNMP users shown in Figure 3-6 in the Information pane.

Figure 3-6 Users Listed Under the Users Tab

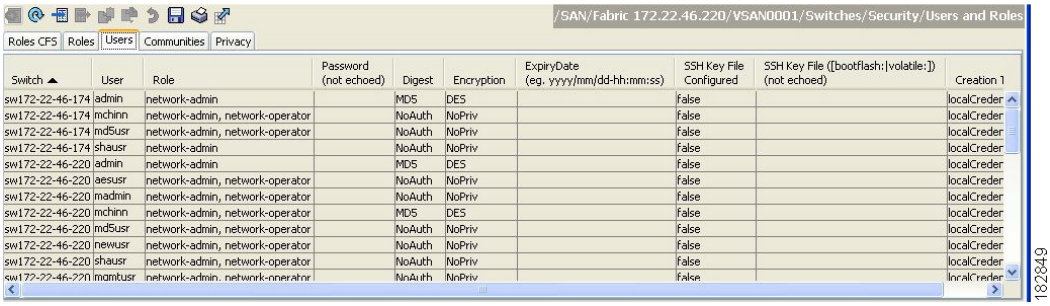

Field Descriptions for Users and Common Role

Common Roles

Note Common roles is not available in displayFCoE mode (use security roles).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Description	Description of the common role.
Enable	This specifies whether the common role has a VSAN restriction or not.
List	List of VSANs user is restricted to.

Feature History for Users and Common Role

Table 3-3 lists the release history for this feature. Only features that were introduced or modified in 5.x or a later release appear in the table.

Table 3-3 Feature History for FIPS

Feature Name	Releases	Feature Information
Changes to SSH	5.0(1a)	Boot Mode SSH, Passwordfree File copy, and SSH.
Role Distributions	5.0(1a)	Enabling role-based configuration distribution.
Creating Users Guidelines	5.0(1a)	Caution has been changed.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

CHAPTER 4

Configuring Security Features on an External AAA Server

The authentication, authorization, and accounting (AAA) feature verifies the identity of, grants access to, and tracks the actions of users managing a switch. All Cisco MDS 9000 Family switches use Remote Access Dial-In User Service (RADIUS) or Terminal Access Controller Access Control device Plus (TACACS+) protocols to provide solutions using remote AAA servers.

Based on the user ID and password combination provided, switches perform local authentication or authorization using the local database or remote authentication or authorization using a AAA server. A preshared secret key provides security for communication between the switch and AAA servers. This secret key can be configured for all AAA servers or for only a specific AAA server. This security feature provides a central management capability for AAA servers.

This chapter includes the following topics:

Information About Switch Management Security

Management security in any switch in the Cisco MDS 9000 Family provides security to all management access methods, including the command-line interface (CLI) or Simple Network Management Protocol (SNMP).

This section includes the following topics:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Security Options

You can access DCNM-SAN using TCP/UDP SNMP or HTTP traffic. For each management path (console, Telnet, and SSH), you can configure one or more of the following security control options: local, remote (RADIUS or TACACS+), or none.

- Remote security control
  - Using RADIUS

    See the "Configuring the RADIUS, TACACS+, and LDAP Server" section on page 4-16
  - Using TACACS+

    See the "Configuring the RADIUS, TACACS+, and LDAP Server" section on page 4-16
- Local security control.

  See the "Local AAA Services" section on page 4-14.

These security features can also be configured for the following scenarios:

- iSCSI authentication

  See the *IP Services Configuration Guide, Cisco DCNM for SAN*.
- Fibre Channel Security Protocol (FC-SP) authentication

  See Chapter 8, "Configuring FC-SP and DHCHAP."

Send documentation comments to dcnm-san-docfeedback@cisco.com

SNMP Security Options

The SNMP agent supports security features for SNMPv1, SNMPv2c, and SNMPv3. Normal SNMP security features apply to all applications that use SNMP (for example, Cisco MDS 9000 DCNM for SAN).

SNMP security options also apply to DCNM for SAN and Device Manager.

See the *Cisco MDS 9000 NX-OS Family System Management Configuration Guide* for more information on the SNMP security options.

Refer to the *Cisco DCNM Fundamentals Guide* for information on DCNM for SAN and Device Manager.

Switch AAA Functionalities

Using the CLI or DCNM for SAN (DCNM-SAN), or an SNMP application, you can configure AAA switch functionalities on any switch in the Cisco MDS 9000 Family.

This section includes the following topics:

Authentication

Authentication is the process of verifying the identity of the person or device accessing the switch. This identity verification is based on the user ID and password combination provided by the entity trying to access the switch. Cisco MDS 9000 Family switches allow you to perform local authentication (using the local lookup database) or remote authentication (using one or more RADIUS or TACACS+ servers).

Note When you log in to a Cisco MDS switch successfully using DCNM-SAN or Device Manager through Telnet or SSH and if that switch is configured for AAA server-based authentication, a temporary SNMP user entry is automatically created with an expiry time of one day. The switch authenticates the SNMPv3 protocol data units (PDUs) with your Telnet or SSH login name as the SNMPv3 user. The management station can temporarily use the Telnet or SSH login name as the SNMPv3 **auth** and **priv** passphrase. This temporary SNMP login is only allowed if you have one or more active MDS shell sessions. If you do not have an active session at any given time, your login is deleted and you will not be allowed to perform SNMPv3 operations.

Note DCNM-SAN does not support AAA passwords with trailing white space, for example “passwordA.”

Send documentation comments to dcnm-san-docfeedback@cisco.com

Authorization

The following authorization roles exist in all Cisco MDS switches:

- Network operator (network-operator)—Has permission to view the configuration only. The operator cannot make any configuration changes.
- Network administrator (network-admin)— Has permission to execute all commands and make configuration changes. The administrator can also create and customize up to 64 additional roles.
- Default-role—Has permission to use the GUI (DCNM-SAN and Device Manager). This access is automatically granted to all users for accessing the GUI.

These roles cannot be changed or deleted. You can create additional roles and configure the following options:

- Configure role-based authorization by assigning user roles locally or using remote AAA servers.
- Configure user profiles on a remote AAA server to contain role information. This role information is automatically downloaded and used when the user is authenticated through the remote AAA server.

Note If a user belongs only to one of the newly created roles and that role is subsequently deleted, then the user immediately defaults to the network-operator role.

Accounting

The accounting feature tracks and maintains a log of every management configuration used to access the switch. This information can be used to generate reports for troubleshooting and auditing purposes. Accounting logs can be stored locally or sent to remote AAA servers.

Remote AAA Services

Remote AAA services provided through RADIUS and TACACS+ protocols have the following advantages over local AAA services:

- User password lists for each switch in the fabric can be managed more easily.
- AAA servers are already deployed widely across enterprises and can be easily adopted.
- The accounting log for all switches in the fabric can be centrally managed.
- User role mapping for each switch in the fabric can be managed more easily.

Server Groups

You can specify remote AAA servers for authentication, authorization, and accounting using server groups. A server group is a set of remote AAA servers implementing the same AAA protocol. The purpose of a server group is to provide for failover servers in case a remote AAA server fails to respond. If the first remote server in the group fails to respond, the next remote server in the group is tried until one of the servers sends a response. If all the AAA servers in the server group fail to respond, then that server group option is considered a failure. If required, you can specify multiple server groups. If the Cisco MDS switch encounters errors from the servers in the first group, it tries the servers in the next server group.

Send documentation comments to dcnm-san-docfeedback@cisco.com

AAA Service Configuration Options

AAA configuration in Cisco MDS 9000 Family switches is service based. You can have separate AAA configurations for the following services:

- Telnet or SSH login (DCNM-SAN and Device Manager login)
- Console login
- iSCSI authentication (see the *IP Services Configuration Guide*, *Cisco DCNM for SAN*).
- FC-SP authentication (see Chapter 8, "Configuring FC-SP and DHCHAP").
- Accounting

In general, server group, local, and none are the three options that can be specified for any service in an AAA configuration. Each option is tried in the order specified. If all the options fail, local is tried.

Caution Cisco MDS NX-OS supports user names that are created with alphanumeric characters or specific special characters (+ [plus], = [equal], _ [underscore], - [hyphen] , \ [backslash], and . [period]) whether created remotely (using TACACS+ or RADIUS) or locally, provided the user name starts with an alphabetical character. Local user names cannot be created with all numbers or with any special characters (apart from those specified). If a numeric-only user name or a non-supported special character user name exists on an AAA server, and is entered during login, then the user is denied access.

Note Even if local is not specified as one of the options, it is tried by default if all AAA servers configured for authentication are unreachable. User has the flexibility to disable this fallback (See section "Configuring Fallback Mechanism for Authentication" section on page 4-18).

When RADIUS times out, local login is attempted depending on the fallback configuration. For this local login to be successful, a local account for the user with the same password should exist, and the RADIUS timeout and retries should take less than 40 seconds. The user is authenticated if the username and password exist in the local authentication configuration.

AAA Server Monitoring

An unresponsive AAA server introduces a delay in the processing of AAA requests. An MDS switch can periodically monitor an AAA server to check whether it is responding (or alive) to save time in processing AAA requests. The MDS switch marks unresponsive AAA servers as dead and does not send AAA requests to any dead AAA servers. An MDS switch periodically monitors dead AAA servers and brings them to the alive state once they are responding. This monitoring process verifies that an AAA server is in a working state before real AAA requests are sent its way. Whenever an AAA server changes to the dead or alive state, an SNMP trap is generated and the MDS switch warns the administrator that a failure is taking place before it can impact performance. See Figure 4-1 for AAA server states.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-1 AAA Server States

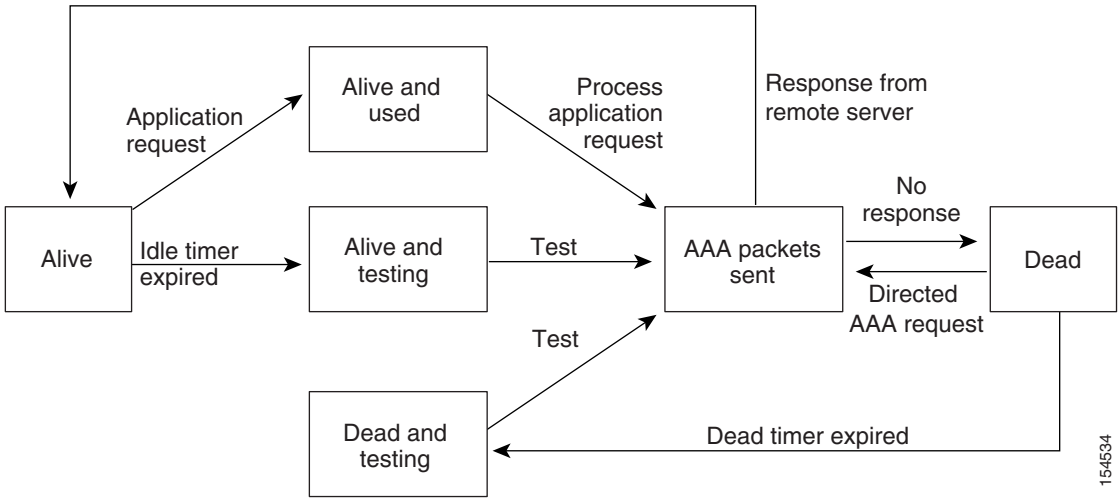

Note The monitoring interval for alive servers and dead servers is different and can be configured by the user. The AAA server monitoring is performed by sending a test authentication request to the AAA server.

The user name and password to be used in the test packet can be configured.

See the "Configuring the RADIUS, TACACS+, and LDAP Server" section on page 4-16.

Authentication and Authorization Process

Authentication is the process of verifying the identity of the person managing the switch. This identity verification is based on the user ID and password combination provided by the person managing the switch. The Cisco MDS 9000 Family switches allow you to perform local authentication (using the lookup database) or remote authentication (using one or more RADIUS servers or TACACS+ servers).

Authorization provides access control. It is the process of assembling a set of attributes that describe what the user is authorized to perform. Based on the user ID and password combination, the user is authenticated and authorized to access the network as per the assigned role. You can configure parameters that can prevent unauthorized access by an user, provided the switches use the TACACS+ protocol.

AAA authorization is the process of assembling a set of attributes that describe what the user is authorized to perform. Authorization in the Cisco NX-OS software is provided by attributes that are downloaded from AAA servers. Remote security servers, such as RADIUS and TACACS+, authorize users for specific rights by associating attribute-value (AV) pairs, which define those rights with the appropriate user.

Figure 4-2 shows a flow chart of the authorization and authentication process.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-2 Switch Authorization and Authentication Flow

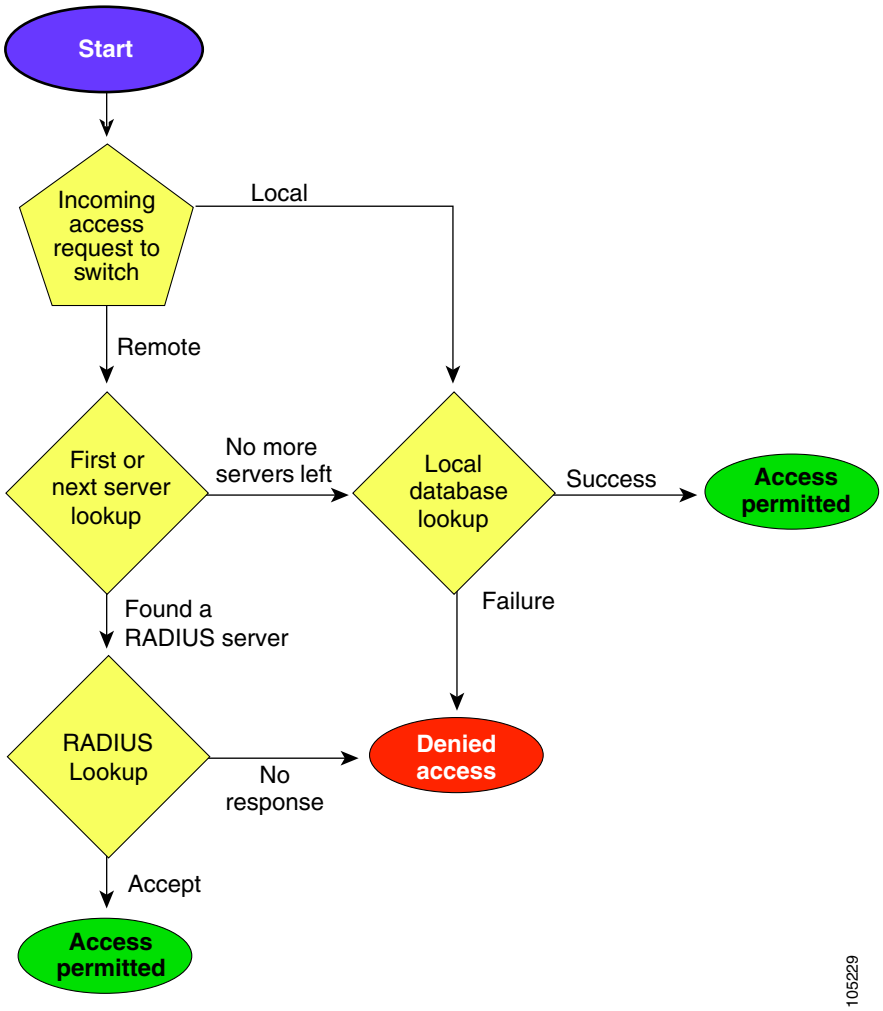

Note No more server groups left = no response from any server in all server groups.
No more servers left = no response from any server within this server group.

Global AAA Server Monitoring Parameters

The global AAA server monitoring parameters function as follows:

- When a new AAA server is configured it is monitored using the global test parameters, if defined.
- When global test parameters are added or modified, all the AAA servers, which do not have any test parameters configured, start getting monitored using the new global test parameters.
- When the server test parameters are removed for a server or when the idle-time is set to zero (default value) the server starts getting monitored using the global test parameters, if defined.
- If global test parameters are removed or global idle-time is set to zero, servers for which the server test parameters are present are not affected. However, monitoring stops for all other servers that were previously being monitored using global parameters.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- If the server monitoring fails with the user-specified server test parameters, the server monitoring does not fall back to global test parameters.

About RADIUS Server Default Configuration

DCNM-SAN allows you to set up a default configuration that can be used for any RADIUS server that you configure the switch to communicate with. The default configuration includes:

- Encryption type
- Timeout value
- Number of retransmission attempts
- Allowing the user to specify a RADIUS server at login

About the Default RADIUS Server Encryption Type and Preshared Key

You need to configure the RADIUS preshared key to authenticate the switch to the RADIUS server. The length of the key is restricted to 64 characters and can include any printable ASCII characters (white spaces are not allowed). You can configure a global key to be used for all RADIUS server configurations on the switch.

You can override this global key assignment by explicitly using the **key** option when configuring an individual RADIUS server.

About RADIUS Servers

You can add up to 64 RADIUS servers. RADIUS keys are always stored in encrypted form in persistent storage. The running configuration also displays encrypted keys. When you configure a new RADIUS server, you can use the default configuration or modify any of the parameters to override the default RADIUS configuration.

Configuring the Test Idle Timer

The test idle timer specifies the interval during which a RADIUS server receives no requests before the MDS switch sends out a test packet.

Note The default idle timer value is 0 minutes. When the idle time interval is 0 minutes, periodic RADIUS server monitoring is not performed.

To configure the test idle timer, see "Configuring an LDAP Server" section on page 4-19.

Configuring Test User Name

You can configure a username and password for periodic RADIUS server status testing. You do not need to configure the test username and password to issue test messages to monitor RADIUS servers. You can use the default test username (test) and default password (test).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Note We recommend that the test username not be the same as an existing username in the RADIUS database for security reasons.

To configure the optional username and password for periodic RADIUS server status testing, see "Configuring an LDAP Server" section on page 4-19.

About Validating a RADIUS Server

As of Cisco SAN-OS Release 3.0(1), you can periodically validate a RADIUS server. The switch sends a test authentication to the server using the username and password that you configure. If the server does not respond to the test authentication, then the server is considered non responding.

Note For security reasons we recommend that you do not use a username that is configured on your RADIUS server as a test username.

You can configure this option to test the server periodically, or you can run a one-time only test.

About Vendor-Specific Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating vendor-specific attributes (VSAs) between the network access server and the RADIUS server. The IETF uses attribute 26. VSAs allow vendors to support their own extended attributes that are not suitable for general use. The Cisco RADIUS implementation supports one vendor-specific option using the format recommended in the specification. The Cisco vendor ID is 9, and the supported option is vendor type 1, which is named **cisco-avpair.** The value is a string with the following format:

```
protocol : attribute separator value *
```

Where **protocol** is a Cisco attribute for a particular type of authorization, **separator** is `=` (equal sign) for mandatory attributes, and `*` (asterisk) is for optional attributes.

When you use RADIUS servers to authenticate yourself to a Cisco MDS 9000 Family switch, the RADIUS protocol directs the RADIUS server to return user attributes, such as authorization information, along with authentication results. This authorization information is specified through VSAs.

VSA Format

The following VSA protocol options are supported by the Cisco NX-OS software:

- **Shell** protocol—Used in Access-Accept packets to provide user profile information.
- **Accounting** protocol—Used in Accounting-Request packets. If a value contains any white spaces, it should be put within double quotation marks.

Send documentation comments to dcnm-san-docfeedback@cisco.com

The following attributes are supported by the Cisco NX-OS software:

- **roles**—This attribute lists all the roles to which the user belongs. The value field is a string storing the list of group names delimited by white space. For example, if you belong to roles **vsan-admin** and **storage-admin**, the value field would be "**vsan-admin storage-admin"**. This subattribute is sent in the VSA portion of the Access-Accept frames from the RADIUS server, and it can only be used with the shell protocol value. These are two examples using the roles attribute:

  ```
  shell:roles="network-admin vsan-admin"
  shell:roles*"network-admin vsan-admin"
  ```

  When an VSA is specified as **shell:roles*"network-admin vsan-admin"**, this VSA is flagged as an optional attribute, and other Cisco devices ignore this attribute.

- **accountinginfo**—This attribute stores additional accounting information besides the attributes covered by a standard RADIUS accounting protocol. This attribute is only sent in the VSA portion of the Account-Request frames from the RADIUS client on the switch, and it can only be used with the accounting protocol-related PDUs.

Specifying SNMPv3 on AAA Servers

The vendor/custom attribute **cisco-av-pair** can be used to specify user's role mapping using the format:

```
shell:roles="roleA roleB …"
```

If the role option in the **cisco-av-pair** attribute is not set, the default user role is network-operator.

The VSA format optionally specifies your SNMPv3 authentication and privacy protocol attributes also as follows:

```
shell:roles="roleA roleB..." snmpv3:auth=SHA priv=AES-128
```

The SNMPv3 authentication protocol options are SHA and MD5. The privacy protocol options are AES-128 and DES. If these options are not specified in the **cisco-av-pair** attribute on the ACS server, MD5 and DES are used by default.

One-Time Password Support

A one-time password (OTP) is a password that is valid for a single login session or transaction. OTPs avoid a number of shortcomings that are associated with usual (static) passwords. The most vital shortcoming that is addressed by OTPs is that, in contrast to static passwords, they are not at risk to replay attacks. If an intruder manages to record an OTP that was already used to log into a service or to conduct an operation, it will not be misused because it is no longer valid.

One-time password applies only to RADIUS and TACACS protocol daemons. In the case of the RADIUS protocol daemon, there is no configuration required from the switch side. In the case of the TACACS protocol, ASCII authentication mode needs to be enabled.

Send documentation comments to dcnm-san-docfeedback@cisco.com

About TACACS+

TACACS+ is a client/server protocol that uses TCP (TCP port 49) for transport requirements. All switches in the Cisco MDS 9000 Family provide centralized authentication using the TACACS+ protocol. The TACACS+ has the following advantages over RADIUS authentication:

- Provides independent, modular AAA facilities. Authorization can be done without authentication.
- Uses the TCP transport protocol to send data between the AAA client and server, making reliable transfers with a connection-oriented protocol.
- Encrypts the entire protocol payload between the switch and the AAA server to ensure higher data confidentiality. The RADIUS protocol only encrypts passwords.

About TACACS+ Server Default Configuration

DCNM-SAN allows you to set up a default configuration that can be used for any TACACS+ server that you configure the switch to communicate with. The default configuration includes:

- Encryption type
- Preshared key
- Timeout value
- Number of retransmission attempts
- Allowing the user to specify a TACACS+ server at login

About the Default TACACS+ Server Encryption Type and Preshared Key

You need to configure the TACACS+ preshared key to authenticate the switch to the TACACS+ server. The length of the key is restricted to 64 characters and can include any printable ASCII characters (white spaces are not allowed). You can configure a global key to be used for all TACACS+ server configurations on the switch.

You can override this global key assignment by explicitly using the **key** option when configuring an individual TACACS+ server.

About TACACS+ Servers

By default, the TACACS+ feature is disabled in all switches in the Cisco MDS 9000 Family. DCNM-SAN or Device Manager enables the TACACS+ feature automatically when you configure a TACACS+ server.

If a secret key is not configured for a configured server, a warning message is issued if a global key is not configured. If a server key is not configured, the global key (if configured) is used for that server.

Note Prior to Cisco MDS SAN-OS Release 2.1(2), you can use the dollar sign ($) in the key but the key must be enclosed in double quotes, for example “k$”. The percent sign (%) is not allowed. In Cisco MDS SAN-OS Release 2.1(2) and later, you can use the dollar sign ($) without double quotes and the percent sign (%) in global secret keys.

You can configure global values for the secret key for all TACACS+ servers.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Note If secret keys are configured for individual servers, those keys override the globally configured key.

Password Aging Notification through TACACS+ Server

Password aging notification is initiated when the user authenticates to a Cisco MDS 9000 switch via a TACACS+ account. The user is notified when a password is about to expire or has expired. If the password has expired, user is prompted to change the password.

Note As of Cisco MDS SAN-OS Release 3.2(1), only TACACS+ supports password aging notification. If you try to use RADIUS servers by enabling this feature, RADIUS generates a SYSLOG message and authentication falls back to the local database.

Password aging notification facilitates the following:

- Password change—You can change your password by entering a blank password.
- Password aging notification—Notifies password aging. Notification happens only if the AAA server is configured and MSCHAP and MSCHAPv2 is disabled.
- Password change after expiration—Initiates password change after the old password expires. Initiation happens from the AAA server.

Note Password aging notification fails if you do not disable MSCHAP and MSCHAPv2 authentication.

About Validating a TACACS+ Server

As of Cisco SAN-OS Release 3.0(1), you can periodically validate a TACACS+ server. The switch sends a test authentication to the server using the test username and test password that you configure. If the server does not respond to the test authentication, then the server is considered nonresponding.

Note We recommend that you do not configure the test user on your TACACS+ server for security reasons.

You can configure this option to test the server periodically, or you can run a one-time only test.

Periodically Validating a TACACS+ Server

To configure the switch to periodically test a TACACS+ server using DCNM-SAN, see the "Configuring the RADIUS, TACACS+, and LDAP Server" section on page 4-16.

Send documentation comments to dcnm-san-docfeedback@cisco.com

About Users Specifying a TACACS+ Server at Login

By default, an MDS switch forwards an authentication request to the first server in the TACACS+ server group. You can configure the switch to allow the user to specify which TACACS+ server to send the authenticate request. If you enable this feature, the user can log in as *username@hostname*, where the *hostname* is the name of a configured TACACS+ server.

About Bypassing a Nonresponsive Server

As of Cisco SAN-OS Release 3.0(1), you can bypass a nonresponsive AAA server within a server group. If the switch detects a nonresponsive server, it will bypass that server when authenticating users. Use this feature to minimize login delays caused by a faulty server. Instead of sending a request to a nonresponsive server and waiting for the authentication request to timeout, the switch sends the authentication request to the next server in the server group. If there are no other responding servers in the server group, the switch continues to attempt authentications against the nonresponsive server.

AAA Server Distribution

Configuration for RADIUS and TACACS+ AAA on an MDS switch can be distributed using the Cisco Fabric Services (CFS). The distribution is disabled by default (see the *System Management Configuration Guide, Cisco DCNM for SAN*).

After enabling the distribution, the first server or global configuration starts an implicit session. All server configuration commands entered thereafter are stored in a temporary database and applied to all switches in the fabric (including the originating one) when you explicitly commit the database. The various server and global parameters are distributed, except the server and global keys. These keys are unique secrets to a switch and should not be shared with other switches.

Note Server group configurations are not distributed.

Note For an MDS switch to participate in AAA server configuration distribution, it must be running Cisco MDS SAN-OS Release 2.0(1b) or later, or Cisco NX-OS Release 4.1(1).

Starting a Distribution Session on a Switch

A distribution session starts the moment you begin a RADIUS or TACACS+ server or global configuration. For example, the following tasks start an implicit session:

- Specifying the global timeout for RADIUS servers.
- Specifying the global timeout for TACACS+ servers.

Note After you issue the first configuration command related to AAA servers, all server and global configurations that are created (including the configuration that caused the distribution session start) are stored in a temporary buffer, not in the running configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

CHAP Authentication

Challenge Handshake Authentication Protocol (CHAP) is a challenge-response authentication protocol that uses the industry-standard Message Digest 5 (MD5) hashing scheme to encrypt the response. CHAP is used by various vendors of network access servers and clients. A server running routing and remote access supports CHAP so that remote access clients that require CHAP are authenticated. CHAP is supported as an authentication method in this release.

MSCHAP Authentication

Microsoft Challenge Handshake Authentication Protocol (MSCHAP) is the Microsoft version of CHAP.

Cisco MDS 9000 Family switches allow user logins to perform remote authentication using different versions of MSCHAP. MSCHAP is used for authentication on a RADIUS or TACACS+ server, while MSCHAPv2 is used for authentication on a RADIUS server.

About Enabling MSCHAP

By default, the switch uses Password Authentication Protocol (PAP) authentication between the switch and the remote server. If you enable MSCHAP, you need to configure your RADIUS server to recognize the MSCHAP vendor-specific attributes. See the "About Vendor-Specific Attributes" section on page 4-9. Table 4-1 shows the RADIUS vendor-specific attributes required for MSCHAP.

Table 4-1 MSCHAP RADIUS Vendor-Specific Attributes

Vendor-ID Number	Vendor-Type Number	Vendor-Specific Attribute	Description
311	11	MSCHAP-Challenge	Contains the challenge sent by an AAA server to an MSCHAP user. It can be used in both Access-Request and Access-Challenge packets.
211	11	MSCHAP-Response	Contains the response value provided by an MS-CHAP user in response to the challenge. It is only used in Access-Request packets.

Local AAA Services

The system maintains the username and password locally and stores the password information in encrypted form. You are authenticated based on the locally stored user information.

Accounting Services

Accounting refers to the log information that is kept for each management session in a switch. This information may be used to generate reports for troubleshooting and auditing purposes. Accounting can be implemented locally or remotely (using RADIUS). The default maximum size of the accounting log is 250,000 bytes and cannot be changed.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Tip The Cisco MDS 9000 Family switch uses interim-update RADIUS accounting-request packets to communicate accounting log information to the RADIUS server. The RADIUS server must be appropriately configured to log the information communicated in these packets. Several servers typically have log update/watchdog packets flags in the AAA client configuration. Turn on this flag to ensure proper RADIUS accounting.

Note Configuration operations are automatically recorded in the accounting log if they are performed in configuration mode. Additionally, important system events (for example, configuration save and system switchover) are also recorded in the accounting log.

Guidelines and Limitations

This section has the following topics:

Remote Authentication Guidelines

If you prefer using remote AAA servers, follow these guidelines:

- A minimum of one AAA server should be IP reachable.
- Be sure to configure a desired local AAA policy as this policy is used if all AAA servers are not reachable.
- AAA servers are easily reachable if an overlay Ethernet LAN is attached to the switch (see the *IP Services Configuration Guide, Cisco DCNM for SAN*). We recommend this method.

SAN networks connected to the switch should have at least one gateway switch connected to the Ethernet LAN reaching the AAA servers.

Merge Guidelines for RADIUS and TACACS+ Configurations

The RADIUS and TACACS+ server and global configuration are merged when two fabrics merge. The merged configuration is applied to CFS distribution-enabled switches.

When merging the fabric, be aware of the following conditions:

- The server groups are not merged.
- The server and global keys are not changed during the merge.
- The merged configuration contains all servers found on all CFS enabled switches.
- The timeout and retransmit parameters of the merged configuration are the largest values found per server and global configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Note Test parameter will be distributed via CFS for TACACS+ Daemon only. If the fabric contains only Cisco NX-OS Release 5.0 devices, then the test parameters will be distributed. If the fabric contains devices running Release 5.0 and some running Release 4.x, the test parameters are not distributed.

Caution If there is a conflict between two switches in the server ports configured, the merge fails.

Default Settings

Table 4-2 lists the default settings for all switch security features in any switch.

Table 4-2 Default Switch Security Settings

Parameters	Default
Roles in Cisco MDS switches	Network operator (network-operator)
AAA configuration services	Local
Authentication port	1812
Accounting port	1813
Preshared key communication	Clear text
RADIUS server timeout	1 (one) second
RADIUS server retries	Once
Authorization	Disabled
aaa user default role	enabled
RADIUS server directed requests	Disabled
TACACS+	Disabled
TACACS+ servers	None configured
TACACS+ server timeout	5 seconds
TACACS+ server directed requests	Disabled
AAA server distribution	Disabled
Accounting log size	250 KB

Configuring the RADIUS, TACACS+, and LDAP Server

Cisco MDS 9000 Family switches can use the RADIUS protocol to communicate with remote AAA servers. You can configure multiple RADIUS servers and server groups and set timeout and retry counts.

RADIUS is a distributed client/server protocol that secures networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco MDS 9000 Family switches and send authentication requests to a central RADIUS server that contains all user authentication and network service access information.

Send documentation comments to dcnm-san-docfeedback@cisco.com

This section defines the RADIUS operation, identifies its network environments, and describes its configuration possibilities.

A Cisco MDS switch uses the Terminal Access Controller Access Control System Plus (TACACS+) protocol to communicate with remote AAA servers. You can configure multiple TACACS+ servers and set timeout values.

This section includes the following topics:

Authorizing and Authenticating the Switch

Detailed Steps

To authorize and authenticate the switch, follow these steps:

Step 1 Log in to the required switch in the Cisco MDS 9000 Family, using the Telnet, SSH, DCNM-SAN or Device Manager, or console login options.

Step 2 When you have configured server groups using the server group authentication method, an authentication request is sent to the first AAA server in the group.

- If the AAA server fails to respond, then the next AAA server is contacted and so on until the remote server responds to the authentication request.
- If all AAA servers in the server group fail to respond, then the servers in the next server group are contacted.
- If all configured methods fail, then by default local database is used for authentication. The next section will describe the way to disable this fallback.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 3 When you are successfully authenticated through a remote AAA server, then the following possible actions are taken:

- If the AAA server protocol is RADIUS, then user roles specified in the **cisco-av-pair** attribute are downloaded with an authentication response.
- If the AAA server protocol is TACACS+, then another request is sent to the same server to get the user roles specified as custom attributes for the shell.
- If user roles are not successfully retrieved from the remote AAA server, then the user is assigned the network-operator role if the **show aaa user default-role** command is enabled. You are denied access if this command is disabled.

Step 4 When your user name and password are successfully authenticated locally, you are allowed to log in, and you are assigned the roles configured in the local database.

Configuring Fallback Mechanism for Authentication

You can enable or disable fallback to the local database in case the remote authentication is set and all of the AAA servers are unreachable (authentication error). The fallback is set to local by default in case of an authentication error. You can disable this fallback for both console and SSH or Telnet login. Disabling this fallback tightens the authentication security.

Detailed Steps

To configure the fallback mechanism, follow this step:

Step 1 Enter the **show run aaa all** command to verify that the default fallback is enabled for both the default and console login.

Disabling fallback will print a warning message.

Caution If fallback is disabled for both the default and console, remote authentication is enabled and servers are unreachable and then the switch will be locked.

Configuring the Default RADIUS Server Encryption Type and Preshared Key

Detailed Steps

To configure the default RADIUS server encryption type and preshared key, follow these steps:

Step 1 Expand **Switches > Security > AAA**, and then select **RADIUS.**

You see the RADIUS configuration in the Information pane.

Step 2 Click the **Defaults** tab.

Step 3 Select **plain** or **encrypted** from the AuthType drop-down menu.

Step 4 Set the key in the Auth Key field.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 5 Click the **Apply Changes** icon to save the changes.

Setting the Default RADIUS Server Timeout Interval and Retransmits

By default, a switch retries transmission to a RADIUS server only once before reverting to local authentication. You can increase this number up to a maximum of five retries per server. You can also configure the timeout value for the RADIUS server.

Detailed Steps

To configure the number of retransmissions and the time between retransmissions to the RADIUS servers, follow these steps:

Step 1 Expand **Switches > Security > AAA,** and then select **RADIUS.**

You see the RADIUS configuration in the Information pane.

Step 2 Choose the **Defaults** tab.

You see the RADIUS default settings.

Step 3 Fill in the Timeout and Retransmits fields for authentication attempts.

Step 4 Click the **Apply Changes** icon to save the changes.

Configuring an LDAP Server

Detailed Steps

To configure an LDAP server and all of its options, follow these steps:

Step 1 Expand **Switches > Security > AAA,** and then select **LDAP**.

You see the LDAP configuration in the Information pane.

Step 2 Click the **Servers** tab.

You see any existing RADIUS servers.

Step 3 Click **Create Row** to add a new LDAP server.

You see the Create LDAP Server dialog box.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-3 LDAP Server Creation

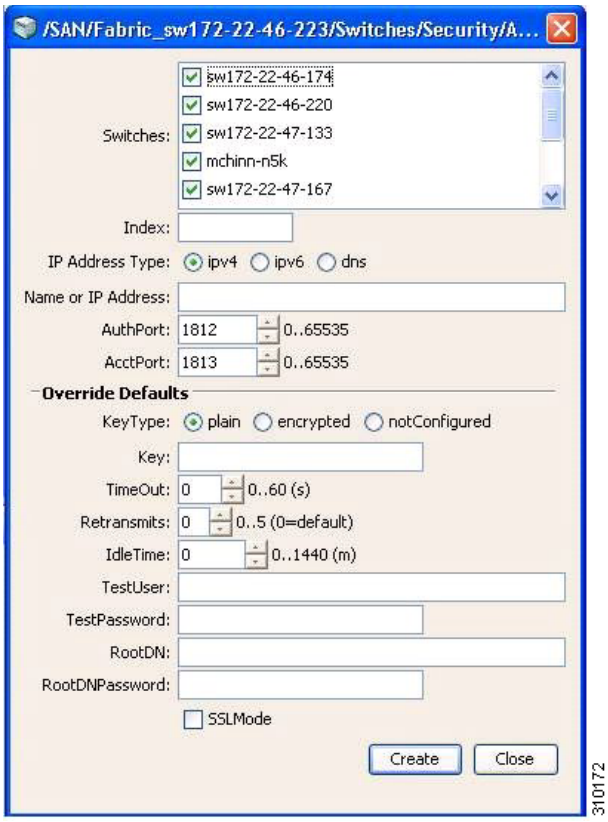

Step 4 Select the switches that you want to assign as LDAP servers.

Step 5 Assign an index number to identify the LDAP server.

Step 6 Select the IP address type for the LDAP server.

Step 7 Fill in the IP address or name for the LDAP server.

Step 8 (Optional) Modify the authentication and accounting ports used by this LDAP server.

Step 9 Select the appropriate key type for the LDAP server.

Step 10 Select the TimeOut value in seconds. The valid range is 0 to 60 seconds.

Step 11 Select the number of times the switch tries to connect to an LDAP server(s) before reverting to local authentication.

Step 12 Enter the test idle time interval value in minutes. The valid range is 1 to 1440 minutes.

Step 13 Enter the test user with the default password. The default username is test.

Step 14 Click **Create** to save these changes.

Creating LDAP Search Map

To create an LDAP search map, follow these steps:

Step 1 Expand **Switches > Security > AAA,** and then select **LDAP**.

Send documentation comments to dcnm-san-docfeedback@cisco.com

You see the LDAP configuration in the Information pane.

Step 2 Click the **Search Map** tab.

Step 3 Click **Create Row** to add a new LDAP search map.

Step 4 Enter the LDAP search map name for the **Name** field.

Step 5 Select the appropriate search type for the **Type** field.

Step 6 Enter the base domain name for the **BaseDN** field.

Step 7 Enter the filter value for the **Filter** field.

Step 8 Enter the attribute value for the **Attribute** field.

Step 9 Click **Create** to save the changes.

Configuring a RADIUS Server

Detailed Steps

To configure a RADIUS server and all its options, follow these steps:

Step 1 Expand **Switches > Security > AAA,** and then select **RADIUS**.

You see the RADIUS configuration in the Information pane.

Step 2 Click the **Servers** tab.

You see any existing RADIUS servers.

Step 3 Click **Create Row** to add a new RADIUS server.

You see the Create RADIUS Server dialog box shown in Figure 4-4.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-4 Create RADIUS Server

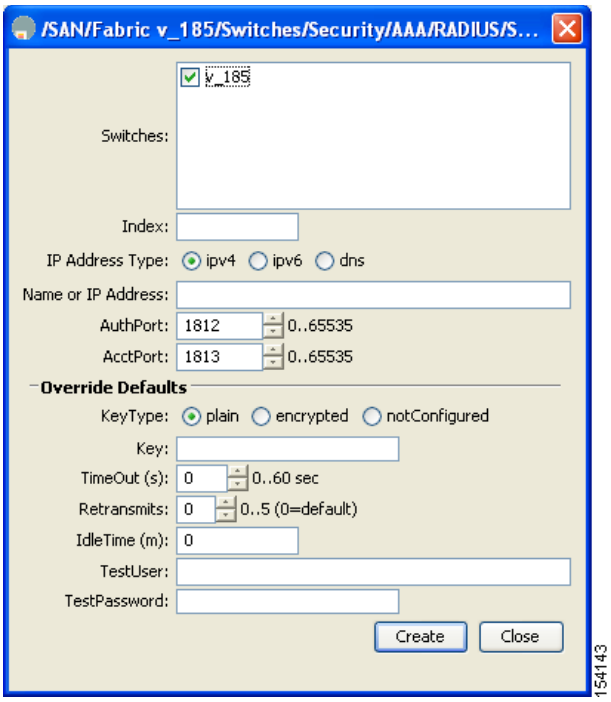

Step 4 Select the switches that you want to assign as RADIUS servers.

Step 5 Assign an index number to identify the RADIUS server.

Step 6 Select the IP address type for the RADIUS server.

Step 7 Fill in the IP address or name for the RADIUS server.

Step 8 (Optional) Modify the authentication and accounting ports used by this RADIUS server.

Step 9 Select the appropriate key type for the RADIUS server.

Step 10 Select the TimeOut value in seconds. The valid range is 0 to 60 seconds.

Step 11 Select the number of times the switch tries to connect to a RADIUS server(s) before reverting to local authentication.

Step 12 Enter the test idle time interval value in minutes. The valid range is 1 to 1440 minutes.

Step 13 Enter the test user with the default password. The default username is test.

Step 14 Click **Create** to save these changes.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Validating a RADIUS Server

Detailed Steps

To configure the switch to periodically test a RADIUS server, follow these steps:

Step 1 Expand **Switches > Security > AAA**, and then select **RADIUS**.

You see the RADIUS configuration in the Information pane.

Step 2 Click the **Servers** tab.

You see any existing RADIUS servers.

Step 3 Click **Create Row** to add a new RADIUS server.

You see the Create RADIUS Server dialog box (see Figure 4-4).

Step 4 Fill in the IP address.

Step 5 Modify the authentication and accounting ports used by this RADIUS server.

Step 6 Fill in the TestUser field and, optionally, the TestPassword field. The default password for the test is **Cisco**.

Step 7 Set the IdleTime field for the time that the server is idle before you send a test authentication.

Step 8 Click **Create** to save these changes.

Allowing Users to Specify a RADIUS Server at Login

By default, an MDS switch forwards an authentication request to the first server in the RADIUS server group. You can configure the switch to allow the user to specify which RADIUS server to send the authenticate request by enabling the directed request option. If you enable this option, the user can log in as *username@hostname*, where the *hostname* is the name of a configured RADIUS server.

Detailed Steps

To allow users logging into an MDS switch to select a RADIUS server for authentication, follow these steps:

Step 1 Expand **Switches > Security > AAA**, and then select **RADIUS**.

You see the RADIUS configuration in the Information pane.

Step 2 Click the **Defaults** tab.

You see the RADIUS default settings.

Step 3 Check the **DirectedReq** check box for the RADIUS server.

Step 4 Click the **Apply Changes** icon to save the changes.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Setting the Default TACACS+ Server Encryption Type and Preshared Key

Detailed Steps

To configure the default TACACS+ server encryption type and preshared key, follow these steps:

Step 1 Expand **Switches > Security > AAA,** and then select **TACACS+.**

You see the TACACS+ configuration in the Information pane.

Step 2 If the Defaults tab is dimmed, click the **CFS** tab.

Step 3 Click the **Defaults** tab.

You see the TACACS+ default settings.

Step 4 Select **plain** or **encrypted** from the AuthType drop-down menu and set the key in the Auth Key field.

Step 5 Click the **Apply Changes** icon to save the changes.

Setting the Default TACACS+ Server Timeout Interval and Retransmits

By default, a switch retries a TACACS+ server only once. This number can be configured. The maximum is five retries per server. You can also configure the timeout value for the TACACS+ server.

Detailed Steps

To configure the number of retransmissions and the time between retransmissions to the TACACS+ servers, follow these steps:

Step 1 Expand **Switches > Security > AAA,** and then select **TACACS+.**

You see the TACACS+ configuration in the Information pane.

Step 2 Click the **Defaults** tab. (If the **Defaults** tab is disabled, click the **CFS** tab first).

You see the TACACS+ default settings.

Step 3 Supply values for the Timeout and Retransmits fields for authentication attempts.

Step 4 Click the **Apply Changes** icon to save the changes.

Configuring a TACACS+ Server

Detailed Steps

To configure a TACACS+ server and all its options using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security > AAA,** and then select **TACACS+**.

You see the TACACS+ configuration in the Information pane.

Step 2 Click the **Servers** tab.

Send documentation comments to dcnm-san-docfeedback@cisco.com

You see any existing TACACS+ servers.

Step 3 Click **Create Row** to add a new TACACS+ server.

You see the Create TACACS+ Server dialog box as shown in Figure 4-5.

Figure 4-5 Create TACACS+ Server Dialog Box

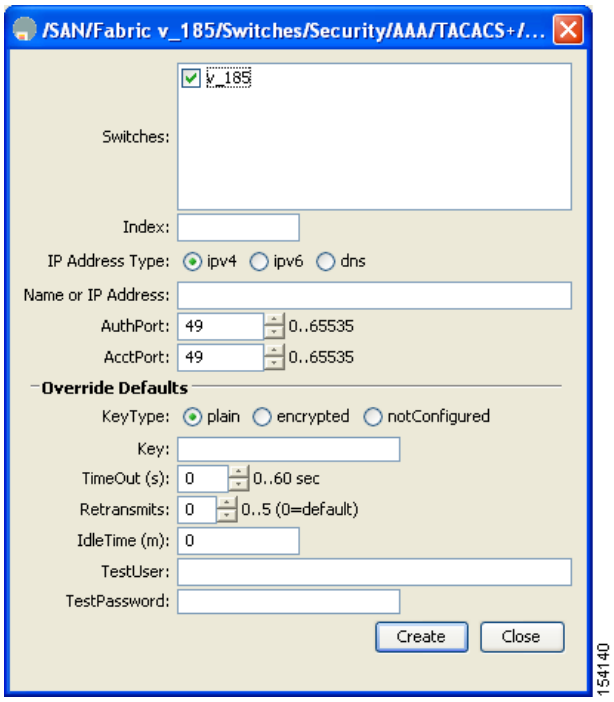

Step 4 Select the switches that you want to assign as TACACS servers.

Step 5 Assign an index number to identify the TACACS server.

Step 6 Select the IP address type for the TACACS server.

Step 7 Fill in the IP address or name for the TACACS server.

Step 8 Modify the authentication and accounting ports used by this TACACS server.

Step 9 Select the appropriate key type for the TACACS server.

Step 10 Select the TimeOut value in seconds. The valid range is 0 to 60 seconds.

Step 11 Select the number of times the switch tries to connect to a TACACS server(s) before reverting to local authentication.

Step 12 Enter the test idle time interval value in minutes. The valid range is 1 to 1440 minutes.

Step 13 Enter the test user with the default password. The default username is test.

Step 14 Click **Create** to save these changes.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Allowing Users to Specify a TACACS+ Server at Login

Detailed Steps

To configure the switch to allow users to specify a TACACS+ server at login using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security > AAA**, and then select **TACACS+**.

You see the TACACS+ configuration in the Information pane.

Step 2 Click the **Defaults** tab.

You see the TACACS+ default settings.

Step 3 Check the **DirectedReq** check box.

Step 4 Click the **Apply Changes** icon to save the changes.

Configuring Server Groups

You can specify one or more remote AAA servers to authenticate users using server groups. All members of a group must belong to the same protocol, either RADIUS or TACACS+. The servers are tried in the same order in which you configure them.

The AAA server monitoring feature can mark an AAA server as dead. You can configure a period of time in minutes to elapse before the switch sends requests to a dead AAA server. (See the "AAA Server Monitoring" section on page 4-5).

Restrictions

You can configure these server groups at any time but they only take effect when you apply them to an AAA service. You configure AAA policies for CLI users or DCNM-SAN or Device Manager users.

Note Configuration of a TACACS+ group fails if MSCHPv2 authentication is not disabled.

Detailed Steps

To configure a RADIUS, TACACS+, or LDAP server group using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security,** and then select **AAA**.

You see the AAA configuration in the Information pane. If you do not see the screen, click the **Server Groups** tab.

You see the RADIUS, TACACS+, or LDAP server groups configured.

Step 2 Click **Create Row** to create a server group.

You see the Create Server dialog box.

Step 3 Click the **radius** radio button to add a RADIUS server group, the **tacacs+** radio button to add a TACACS+ server group, and the **ldap** radio button to add a LDAP server group.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 4 Supply server names for the ServerIdList field.

Step 5 When you chose LDAP, enter the LDAP search map name for the LDAPSearchMapName.

- LDAPSSLMODE—Specifies if the TLS tunnel should be setup before binding with the LDAP server.
- LDAPBindFirst—Specifies if the user bind should be completed before the search.

Step 6 Click the **plain** radio button to select the plain authentication method, click the **kerberos** button to select the kerberos authentication method, and click **md5digest** to select the md5digest authentication method.

Step 7 Enter the password for the **LDAPComparePasswd** field:

- LDAPCertDNBind—Specifies if the User Certification Bind needs to be checked while doing PKI SSH certificate authorization.
- LDAPUserServerBind—Specifies if the User Server Bind should be checked as part of SSH PKI authorization.

Step 8 Set the DeadTime field for the number of minutes that a server can be nonresponsive before it is marked as bypassed. See the "About Bypassing a Nonresponsive Server" section on page 4-13.

Step 9 Click **Create** to create this server group.

The LDAP Server Group displays LDAP-specific parameters.

Step 10 Click the **Applications** tab to assign this server group to an application.

You can associate a server group with all applications or you can specify specific applications.

Step 11 Click the **General** tab to assign the type of authentication to this server group.

Check either the MSCHAP or MSCHAPv2 check box based on the type of server group.

Step 12 Click the **Apply Changes** icon to save the changes.

Once the LDAP Server group is created, the configuration information is displayed in two tabs:

- Server Groups—Displays common data shared by all AAA protocols (RADIUS, TACACS+, and LDAP).
- LDAP Server Group—Displays only LDAP-specific protocols.

Enabling AAA Server Distribution

Restrictions

Only switches where distribution is enabled can participate in the distribution activity.

Detailed Steps

To enable RADIUS server distribution, follow these steps:

Step 1 Expand **Switches > Security > AAA,** and then select **RADIUS**.

You see the RADIUS configuration in the Information pane.

Step 2 Click the **CFS** tab. You see the RADIUS CFS configuration.

Step 3 Choose **enable** from the Admin drop-down list for all switches that you want to enable CFS for RADIUS.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 4 Click **Apply Changes** to distribute these changes through the fabric.

To enable TACACS+ server distribution, follow these steps:

Step 1 Expand **Switches > Security > AAA,** and then select **TACACS+**.

You see the TACACS+ configuration in the Information pane.

Step 2 Click the **CFS** tab.

You see the TACACS+ CFS configuration.

Step 3 Choose **enable** from the Admin drop-down list for all switches that you want to enable CFS on for TACACS+.

Step 4 Click **Apply Changes** to distribute these changes through the fabric.

Committing the Distribution

The RADIUS or TACACS+ global and/or server configuration stored in the temporary buffer can be applied to the running configuration across all switches in the fabric (including the originating switch).

Detailed Steps

To distribute a RADIUS or TACACS+ configuration, follow these steps:

Step 1 Expand **Switches > Security > AAA,** and then select either **RADIUS** or **TACACS+**. You see the RADIUS or TACACS+ configuration in the Information pane.

Step 2 Click the **CFS** tab. You see the RADIUS or TACACS+ CFS configuration.

Step 3 Choose **commitChanges** in the Config Action drop-down list for all switches that you want to enable CFS for RADIUS or TACACS+.

Step 4 Click **Apply Changes** to distribute the changes through the fabric.

Discarding the Distribution Session

Discarding the distribution of a session in progress causes the configuration in the temporary buffer to be dropped. The distribution is not applied.

Detailed Steps

To discard RADIUS or TACACS+ distribution, follow these steps:

Step 1 Expand **Switches > Security > AAA,** and then select either **RADIUS** or **TACACS+**. You see either the RADIUS or TACACS+ configuration in the Information pane.

Step 2 Click the **CFS** tab. You see either the RADIUS or TACACS+ CFS configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 3 Choose **abort** from the Config Action drop-down list for each switch that should discard the pending RADIUS or TACACS+ distribution.

Step 4 Click **Apply Changes**.

Clearing Sessions

Detailed Steps

To clear a RADIUS or TACACS+ distribution, follow these steps:

Step 1 Expand **Switches > Security > AAA** and then select either **RADIUS** or **TACACS+**.

You see either the RADIUS or TACACS+ configuration in the Information pane.

Step 2 Choose the **CFS** tab. You see either the RADIUS or TACACS+ CFS configuration.

Step 3 Choose **clear** from the Config Action drop-down list for each switch that should clear the pending RADIUS or TACACS+ distribution.

Step 4 Click **Apply Changes**.

Enabling MSCHAP Authentication

Detailed Steps

Note Password aging, MSCHAPv2, and MSCHAP authentication can fail if one of these authentication is not disabled.

To enable MSCHAP authentication using Device Manager, follow these steps:

Step 1 Click **Security > AAA.**

You see the AAA configuration in the Information pane (see Figure 4-6).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-6 AAA Configuration in Device Manager

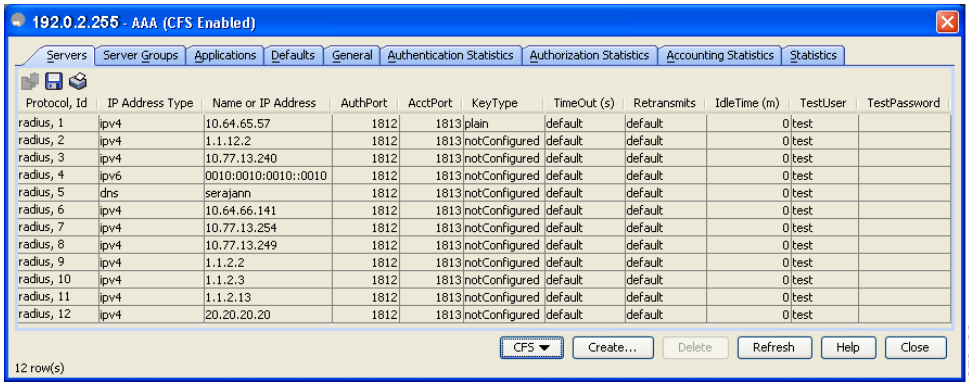

Step 2 Click the **General** tab.

You see the MSCHAP configuration (see Figure 4-7).

Figure 4-7 MSCHAP Configuration

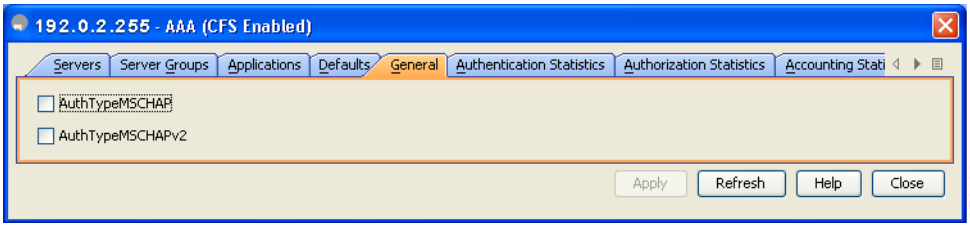

Step 3 Check the **AuthTypeMSCHAP** or **AuthTypeMSCHAPv2** check box to use MSCHAP or MSCHAPv2 to authenticate users on the switch.

Step 4 Click **Apply Changes** to save the changes.

Configuring Cisco Access Control Servers

The Cisco Access Control Server (ACS) uses TACACS+ and RADIUS protocols to provide AAA services that ensure a secure environment. When using the AAA server, user management is normally done using Cisco ACS. Figure 4-8, Figure 4-9, Figure 4-10, and Figure 4-11 display ACS server user setup configurations for network-admin roles and multiple roles using either RADIUS or TACACS+.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-8 Configuring the network-admin Role When Using RADIUS

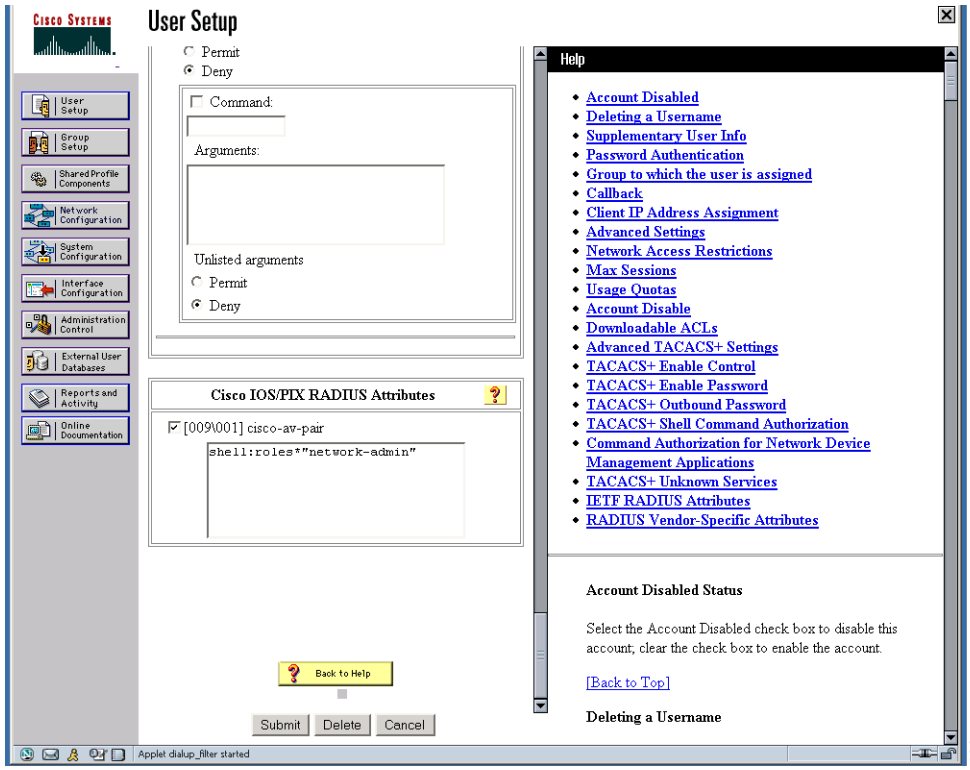

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-9 Configuring Multiple Roles with SNMPv3 Attributes When Using RADIUS

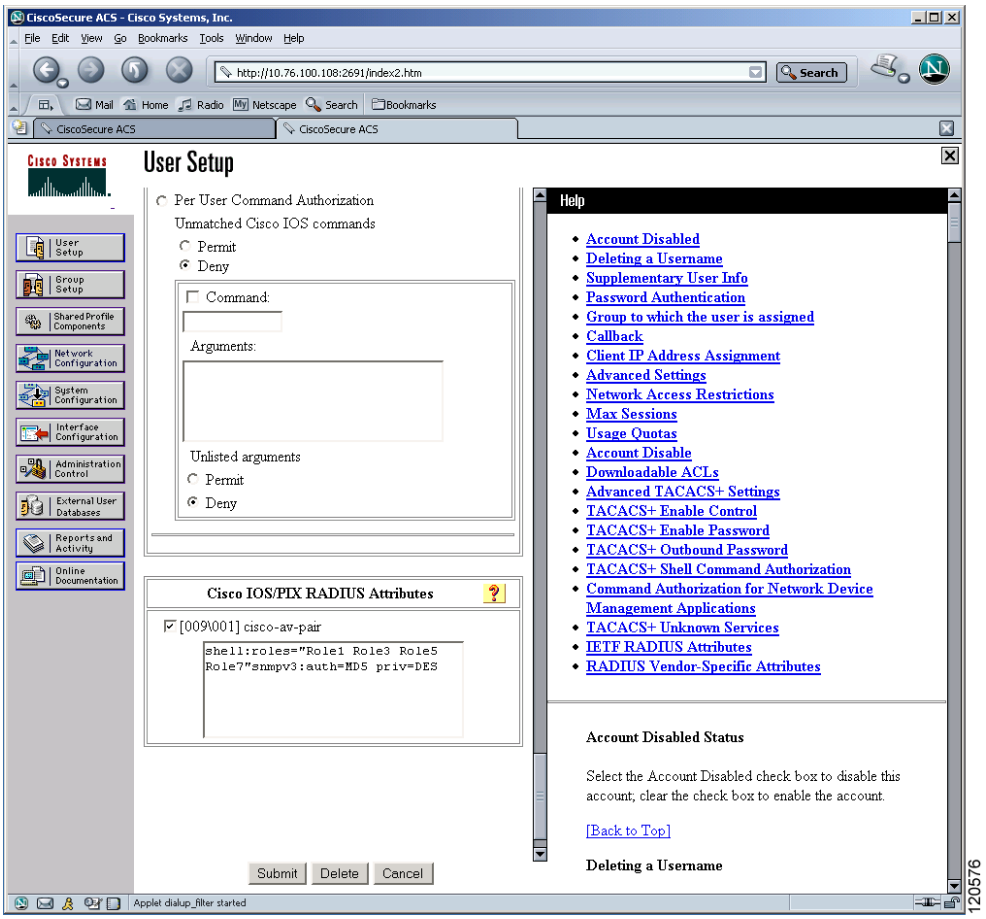

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-10 Configuring the network-admin Role with SNMPv3 Attributes When Using TACACS+

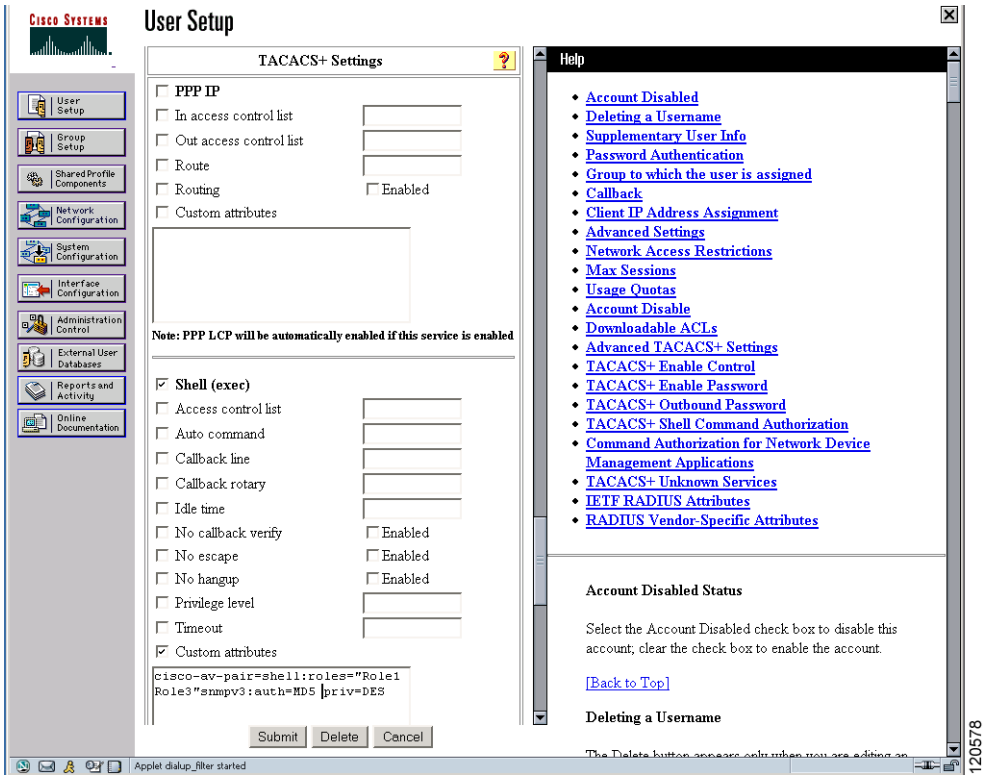

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 4-11 Configuring Multiple Roles with SNMPv3 Attributes When Using TACACS+

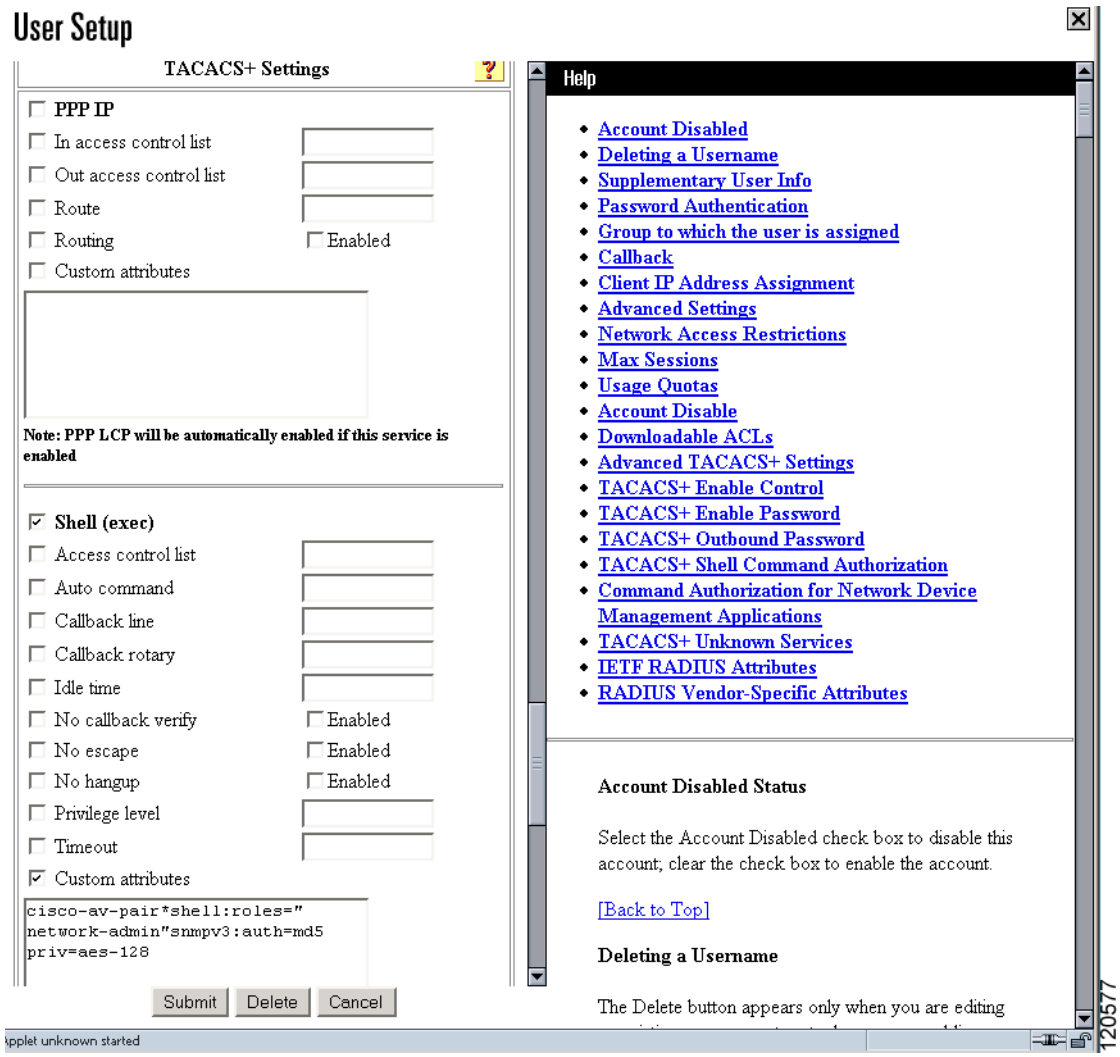

Verifying RADIUS and TACACS+ Configuration

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

- Displaying RADIUS Server Statistics, page 4-35
- Displaying TACACS+ Server Statistics, page 4-35
- Displaying the Pending Configuration to be Distributed, page 4-35

Send documentation comments to dcnm-san-docfeedback@cisco.com

Displaying RADIUS Server Statistics

To display RADIUS server statistics, follow these steps:

Step 1 Expand **Switches > Security > AAA,** and then select **RADIUS**.

You see the RADIUS configuration in the Information pane.

Step 2 Click the **Statistics** tab.

You see the RADIUS server statistics.

Displaying TACACS+ Server Statistics

To display TACACS+ server statistics, follow these steps:

Step 1 Expand **Switches > Security > AAA,** and then select **TACACS+**.

You see the TACACS+ configuration in the Information pane.

Step 2 Choose the **Statistics** tab.

You see the TACACS+ server statistics.

Displaying the Pending Configuration to be Distributed

To display the RADIUS or TACACS+ global and/or server configuration stored in the temporary buffer, follow these steps:

Step 1 Expand **Switches > Security > AAA,** and then select **RADIUS** or select **TACACS+.**

Step 2 Click the **CFS** tab.

You see the distribution status on the CFS tab.

Step 3 Click the **pending** or **running** radio button.

Step 4 Click **Apply Changes** to save the changes.

Step 5 Click the **Servers** tab to view the pending or running configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Feature History for RADIUS, TACACS+, and LDAP

Table 4-3 lists the release history for this feature. Only features that were introduced or modified in Cisco NX-OS Release 5.x or a later release appear in the table.

Table 4-3 Feature History for RADIUS, TACACS+, and LDAP

Feature Name	Releases	Feature Information
Configuring LDAP	5.2	Added LDAP server and server groups configuration.
Switch AAA Functionalities	5.0(1a)	Added configuring fallback mechanism for authentication, configuring AAA server monitoring parameters globally.
Configuring TACACS+ Server Monitoring Parameters	5.0(1a)	Added CHAP authentication.
OTP Authentication	5.0(1a)	Added one-time password support
Merge Guidelines for RADIUS and TACACS+ Configurations	5.0(1a)	TACACS test parameters have to be distributed via CFS; a note has been changed.
AAA Service Configuration Options	5.0(1a)	Note has been changed.

Send documentation comments to dcnm-san-docfeedback@cisco.com

CHAPTER 5

Configuring IPv4 and IPv6 Access Control Lists

Cisco MDS 9000 Family switches can route IP version 4 (IPv4) traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature routes traffic between VSANs. To do so, each VSAN must be in a different IPv4 subnetwork. Each Cisco MDS 9000 Family switch provides the following services for network management systems (NMS):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.
- IP forwarding on the in-band Fibre Channel interface using the IP over Fibre Channel (IPFC) function—IPFC specifies how IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.
- IP routing (default routing and static routing)—If your configuration does not need an external router, you can configure a default route using static routing.

Switches are compliant with RFC 2338 standards for Virtual Router Redundancy Protocol (VRRP) features. VRRP is a restartable application that provides a redundant, alternate path to the gateway switch.

IPv4 Access Control Lists (IPv4-ACLs and IPv6-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IPv4-ACLs and IPv6-ACLs restrict IP-related traffic based on the configured IP filters. A filter contains the rules to match an IP packet, and if the packet matches, the rule also stipulates if the packet should be permitted or denied.

Each switch in the Cisco MDS 9000 Family can have a maximum total of 128 IPv4-ACLs or 128 IPv6-ACLs and each IPv4-ACL or IPv6-ACL can have a maximum of 256 filters.

This chapter includes the following topics:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Information About IPv4 and IPv6 Access Control Lists

Cisco MDS 9000 Family switches can route IP version 4 (IPv4) traffic between Ethernet and Fibre Channel interfaces. The IP static routing feature routes traffic between VSANs. To do so, each VSAN must be in a different IPv4 subnetwork. Each Cisco MDS 9000 Family switch provides the following services for network management systems (NMS):

- IP forwarding on the out-of-band Ethernet interface (mgmt0) on the front panel of the supervisor modules.
- IP forwarding on the in-band Fibre Channel interface using the IP over Fibre Channel (IPFC) function—IPFC specifies how IP frames can be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so NMS information can cross the Fibre Channel network without using an overlay Ethernet network.
- IP routing (default routing and static routing)—If your configuration does not need an external router, you can configure a default route using static routing.

IPv4 Access Control Lists (IPv4-ACLs and IPv6-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IPv4-ACLs and IPv6-ACLs restrict IP-related traffic based on the configured IP filters. A filter contains the rules to match an IP packet, and if the packet matches, the rule also stipulates if the packet should be permitted or denied.

Each switch in the Cisco MDS 9000 Family can have a maximum total of 128 IPv4-ACLs or 128 IPv6-ACLs and each IPv4-ACL or IPv6-ACL can have a maximum of 256 filters.

This section contains the following topics:

About Filter Contents

An IP filter contains rules for matching an IP packet based on the protocol, address, port, ICMP type, and type of service (ToS).

Protocol Information

The protocol information is required in each filter. It identifies the name or number of an IP protocol. You can specify the IP protocol in one of two ways:

- Specify an integer ranging from 0 to 255. This number represents the IP protocol.
- Specify the name of a protocol including, but not restricted to, Internet Protocol (IP), Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Internet Control Message Protocol (ICMP).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Note When configuring IPv4-ACLs or IPv6-ACLs on Gigabit Ethernet interfaces, only use the TCP or ICMP options.

Address Information

The address information is required in each filter. It identifies the following details:

- Source—The address of the network or host from which the packet is being sent.
- Source-wildcard—The wildcard bits applied to the source.
- Destination—The number of the network or host to which the packet is being sent.
- Destination-wildcard—The wildcard bits applied to the destination.

Specify the source and source-wildcard or the destination and destination-wildcard in one of two ways:

- Using the 32-bit quantity in four-part, dotted decimal format (10.1.1.2/0.0.0.0 is the same as host 10.1.1.2).
  - Each wildcard bit set to zero indicates that the corresponding bit position in the packet's IPv4 address must exactly match the bit value in the corresponding bit position in the source.
  - Each wildcard bit set to one indicates that both a zero bit and a one bit in the corresponding position of the packet's IPv4 or IPv6 address will be considered a match to this access list entry. Place ones in the bit positions you want to ignore. For example, 0.0.255.255 requires an exact match of only the first 16 bits of the source. Wildcard bits set to one do not need to be contiguous in the source-wildcard. For example, a source-wildcard of 0.255.0.64 would be valid.
- Using the **any** option as an abbreviation for a source and source-wildcard or destination and destination-wildcard (0.0.0.0/255.255.255.255)

Port Information

The port information is optional. To compare the source and destination ports, use the **eq** (equal) option, the **gt** (greater than) option, the **lt** (less than) option, or the **range** (range of ports) option. You can specify the port information in one of two ways:

- Specify the number of the port. Port numbers range from 0 to 65535. Table 5-1 displays the port numbers recognized by the Cisco NX-OS software for associated TCP and UDP ports.
- Specify the name of a TCP or UDP port as follows:
  - TCP port names can only be used when filtering TCP.
  - UDP port names can only be used when filtering UDP.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 5-1 TCP and UDP Port Numbers

Protocol	Port	Number
UDP	dns	53
	tftp	69
	ntp	123
	radius accounting	1646 or 1813
	radius authentication	1645 or 1812
	snmp	161
	snmp-trap	162
	syslog	514
TCP[1]	ftp	20
	ftp-data	21
	ssh	22
	telnet	23
	smtp	25
	tasacs-ds	65
	www	80
	sftp	115
	http	143
	wbem-http	5988
	wbem-https	5989

1. If the TCP connection is already established, use the **established** option to find matches. A match occurs if the TCP datagram has the ACK, FIN, PSH, RST, or URG control bit set.

ICMP Information

IP packets can be filtered based on the following optional ICMP conditions:

- icmp-type—The ICMP message type is a number from 0 to 255.
- icmp-code—The ICMP message code is a number from 0 to 255.

Table 5-2 displays the value for each ICMP type.

Table 5-2 ICMP Type Value

ICMP Type[1]	Code
echo	8
echo-reply	0
destination unreachable	3
traceroute	30
time exceeded	11

Send documentation comments to dcnm-san-docfeedback@cisco.com

1. ICMP redirect packets are always rejected.

ToS Information

IP packets can be filtered based on the following optional ToS conditions:

- ToS level—The level is specified by a number from 0 to 15.
- ToS name—The name can be max-reliability, max-throughput, min-delay, min-monetary-cost, and normal.

Guidelines and Limitations

Follow these guidelines when configuring IPv4-ACLs or IPv6-ACLs in any switch or director in the Cisco MDS 9000 Family:

- You can apply IPv4-ACLs or IPv6-ACLs to VSAN interfaces, the management interface, Gigabit Ethernet interfaces on IPS modules and MPS-14/2 modules, and Ethernet PortChannel interfaces.

Tip If IPv4-ACLs or IPv6-ACLs are already configured in a Gigabit Ethernet interface, you cannot add this interface to an Ethernet PortChannel group. See the *IP Services Configuration Guide, Cisco DCNM for SAN* for guidelines on configuring IPv4-ACLs.

Caution Do not apply IPv4-ACLs or IPv6-ACLs to only one member of a PortChannel group. Apply IPv4-ACLs or IPv6-ACLs to the entire channel group.

- Configure the order of conditions accurately. As the IPv4-ACL or the IPv6-ACL filters are sequentially applied to the IP flows, only the first match determines the action taken. Subsequent matches are not considered. Be sure to configure the most important condition first. If no conditions match, the software drops the packet.
- Configure explicit deny on the IP Storage Gigabit Ethernet ports to apply IP ACLs because implicit deny does not take effect on these ports.

Configuring IPv4-ACLs or IPv6-ACLs

This section contains the following topics:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Creating IPv4-ACLs or IPv6-ACLs with the IP-ACL Wizard

Traffic coming into the switch is compared to IPv4-ACL or IPv6-ACL filters based on the order that the filters occur in the switch. New filters are added to the end of the IPv4-ACL or the IPv6-ACL. The switch keeps looking until it has a match. If no matches are found when the switch reaches the end of the filter, the traffic is denied. For this reason, you should have the frequently hit filters at the top of the filter. There is an *implied deny* for traffic that is not permitted. A single-entry IPv4-ACL or IPv6-ACL with only one deny entry has the effect of denying all traffic.

Detailed Steps

To configure an IPv4-ACL or an IPv6-ACL, follow these steps:

Step 1 Create an IPv4-ACL or an IPv6-ACL by specifying a filter name and one or more access condition(s). Filters require the source and destination address to match a condition. Use optional keywords to configure finer granularity.

Note The filter entries are executed in sequential order. You can only add the entries to the end of the list. Take care to add the entries in the correct order.

Step 2 Apply the access filter to specified interfaces.

To create an ordered list of IP filters in a named IPv4-ACL or IPv6-ACL profile using the IPv4-ACL Wizard, follow these steps:

Detailed Steps

Step 1 Click the **IP ACL Wizard** icon from the DCNM-SAN toolbar (see Figure 5-1).

Figure 5-1 IP ACL Wizard

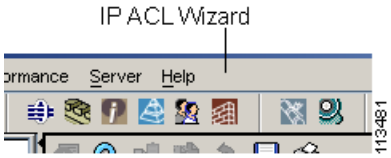

You see the IP ACL Wizard.

Step 2 Enter a name for the IP-ACL.

Note If you are creating an IPv6-ACL, check the IPv6 check box.

Step 3 Click **Add** to add a new rule to this IP-ACL. You see a new rule in the table with default values.

Step 4 Modify the Source IP and Source Mask as necessary for your filter.

Note The IP-ACL Wizard only creates inbound IP filters.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 5 Choose the appropriate filter type from the Application drop-down list.

Step 6 Choose **permit** or **deny** from the Action drop-down list.

Step 7 Repeat Step 3 through Step 6 for additional IP filters.

Step 8 Click **Up** or **Down** to order the filters in this IP-ACL.

Tip Order the IP filters carefully. Traffic is compared to the IP filters in order. The first match is applied and the rest are ignored.

Step 9 Click **Next**.

You see a list of switches that you can apply this IP-ACL.

Step 10 Uncheck any switches that you do not want to apply this IP-ACL.

Step 11 Select the **Interface** you want to apply this IP-ACL.

Step 12 Click **Finish** to create this IP-ACL and apply it to the selected switches.

Creating IPv4-ACLs or IPv6-ACLs

Detailed Steps

To add entries to an existing IPv4-ACL or an IPv6-ACL using Device Manager, follow these steps:

Step 1 Choose **Security > IP ACL.**

Step 2 Click **Create** to create an IP-ACL profile.

You see the Create IP ACL Profiles dialog box. Enter an IP-ACL profile name.

Step 3 Click **Create** and then click **Close**.

This creates a new IP-ACL profile.

Step 4 Click the IP-ACL you created and click **Rules**.

After you create an IPv4-ACL or an IPv6-ACL, you can add subsequent IP filters at the end of the IPv4-ACL or the IPv6-ACL if you are using Device Manager. DCNM-SAN allows you to reorder existing rules for a profile. You cannot insert filters in the middle of an IPv4-ACL or an IPv6-ACL. Each configured entry is automatically added to the end of a IPv4-ACL or an IPv6-ACL.

Step 5 Click **Create** to create an IP filter.

Step 6 Choose either **permit** or **deny** for the Action and set the IP Number in the Protocol field. The drop-down menu provides common filtered protocols.

Step 7 Set the source IP address you want this filter to match against and the wildcard mask, or check the **any** check box to match this filter against any IP address.

This creates an IP filter that will check the source IP address of frames.

Note The wildcard mask denotes a subset of the IP address you want to match against. This allows a range of addresses to match against this filter.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 8 Set the transport layer source port range if the protocol chosen is TCP or UDP.

Step 9 Repeat Step 7 and Step 8 for the destination IP address and port range.

This creates an IP filter that will check the destination IP address of frames.

Step 10 Set the ToS, ICMPType, and ICMPCode fields as appropriate.

Step 11 Check the **TCPEstablished** check box if you want to match TCP connections with ACK,FIN,PSH,RST,SYN or URG control bits set.

Step 12 Check the **LogEnabled** check box if you want to log all frames that match this IP filter.

Step 13 Click **Create** to create this IP filter and add it to your IP-ACL.

Removing IP Filters from an Existing IPv4-ACL or IPv6-ACL

Detailed Steps

To remove configured entries from an IPv4-ACL or an IPv6-ACL using Device Manager, follow these steps:

Step 1 Choose **Security > IP ACLs**.

You see the IP-ACL dialog box.

Step 2 Click the IP-ACL you want to modify and click **Rules**.

You see the list of IP filters associated with this IP-ACL.

Step 3 Select the filter that you want to delete and click **Delete** to delete that IP filter.

Deleting IP-ACLs

Prerequisites

You must delete the association between the IP-ACL and interfaces before deleting the IP-ACL.

Detailed Steps

To delete an IP-ACL, follow these steps:

Step 1 Expand **Switches > Security,** and then select **IP ACL** from the Physical Attributes pane.

You see the IP-ACL configuration in the Information pane.

Step 2 Click the **Profiles** tab.

You see a list of switches, ACLs, and profile names.

Step 3 Select the row you want to delete. To delete multiple rows, hold down the Shift key while selecting rows.

Step 4 Click **Delete Row**. The IP-ACLs are deleted.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Reading the IP-ACL Log Dump

Use the LogEnabled check box option during IP filter creation to log information about packets that match this filter. The log output displays the ACL number, permit or deny status, and port information.

For the input ACL, the log displays the raw MAC information. The keyword “MAC=” does not refer to showing an Ethernet MAC frame with MAC address information. It refers to the Layer 2 MAC-layer information dumped to the log. For the output ACL, the raw Layer 2 information is not logged.

The following example is an input ACL log dump:

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN=vsan1 OUT=
MAC=10:00:00:05:30:00:47:df:10:00:00:05:30:00:8a:1f:aa:aa:03:00:00:00:08:00:45:00:00:54:00
:00:40:00:40:01:0e:86:0b:0b:0b:0c:0b:0b:0b:02:08:00:ff:9c:01:15:05:00:6f:09:17:3f:80:02:01
:00:08:09:0a:0b:0c:0d:0e:0f:10:11:12:13:14:15:16:17:18:19:1a:1b:1c:1d:1e:1f:20:21:22:23:24
:25:26:27:28:29:2a:2b SRC=11.11.11.12 DST=11.11.11.2 LEN=84 TOS=0x00 PREC=0x00 TTL=64 ID=0
DF PROTO=ICMP TYPE=8 CODE=0 ID=277 SEQ=1280
```

The following example is an output ACL log dump:

```
Jul 17 20:38:44 excal-2
%KERN-7-SYSTEM_MSG:
%IPACL-7-DENY:IN= OUT=vsan1 SRC=11.11.11.2 DST=11.11.11.12 LEN=84 TOS=0x00 PREC=0x00
TTL=255 ID=38095 PROTO=ICMP TYPE=0 CODE=0 ID=277 SEQ=1280
```

Applying an IP-ACL to an Interface

You can define IP-ACLs without applying them. However, the IP-ACLs will have no effect until they are applied to an interface on the switch. You can apply IP-ACLs to VSAN interfaces, the management interface, Gigabit Ethernet interfaces on IPS modules and MPS-14/2 modules, and Ethernet PortChannel interfaces.

Tip Apply the IP-ACL on the interface closest to the source of the traffic.

When you are trying to block traffic from source to destination, you can apply an inbound IPv4-ACL to M0 on Switch 1 instead of an outbound filter to M1 on Switch 3 (see Figure 5-2).

Figure 5-2 Denying Traffic on the Inbound Interface

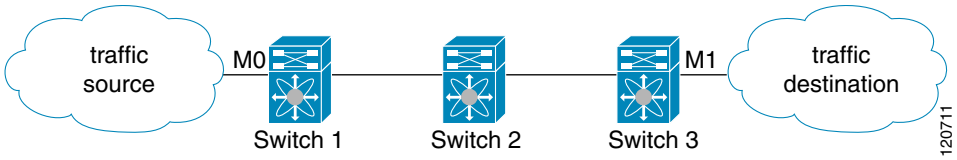

The **access-group** option controls access to an interface. Each interface can only be associated with one IP-ACL per direction. The ingress direction can have a different IP-ACL than the egress direction. The IP-ACL becomes active when applied to the interface.

Tip Create all conditions in an IP-ACL before applying it to the interface.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Caution If you apply an IP-ACL to an interface before creating it, all packets in that interface are dropped because the IP-ACL is empty.

The terms *in, out, source*, and *destination* are used as referenced by the switch:

- In—Traffic that arrives at the interface and goes through the switch; the source is where it transmitted from and the destination is where it is transmitted to (on the other side of the router).

Tip The IP-ACL applied to the interface for the ingress traffic affects both local and remote traffic.

- Out—Traffic that has already been through the switch and is leaving the interface; the source is where it transmitted from and the destination is where it is transmitted to.

Tip The IP-ACL applied to the interface for the egress traffic only affects local traffic.

Applying an IP-ACL to mgmt0

A system default ACL called mgmt0 exists on the mgmt0 interface. This ACL is not visible to the user, so mgmt0 is a reserved ACL name that cannot be used. The mgmt0 ACL blocks most ports and only allows access to required ports in compliance to accepted security policies.

Detailed Steps

To apply an IP-ACL to an interface, follow these steps:

Step 1 Expand **Switches > Security,** and then select **IP ACL** in the Physical Attributes pane.

You see the IP-ACL configuration in the Information pane.

Step 2 Click the **Interfaces** tab.

You see a list of interfaces and associated IP-ACLs.

Step 3 Click **Create Row**.

Step 4 (Optional) Remove the switches you do not want to include in the IP-ACL by unchecking the check boxes next to the switch addresses.

Set the **interface** you want associated with an IPv4-ACL or IPv6-ACL in the Interface field.

Step 5 Choose a ProfileDirection (either **inbound** or **outbound**).

Step 6 Enter the IP-ACL name in the Profile Name field.

Note This IP-ACL name must have already been created using the Create Profiles dialog box. If not, no filters will be enabled until you to go to the Create Profiles dialog box and create the profile.

Step 7 Click **Create** to associate the IP-ACL.

You see the newly associated access list in the list of IP-ACLs.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuration Examples for IP-ACL

To define an IP-ACL that restricts management access using Device Manager, follow these steps:

Step 1 Choose **Security > IP ACL.**

You see the IP-ACL dialog box.

Step 2 Click **Create** to create an IP-ACL.

You see the Create IP ACL Profiles dialog box.

Step 3 Enter **RestrictMgmt** as the profile name and click **Create**.

This creates an empty IP-ACL named RestrictMgmt (see Figure 5-3).

Figure 5-3 RestrictMgmt Profile Added to the List

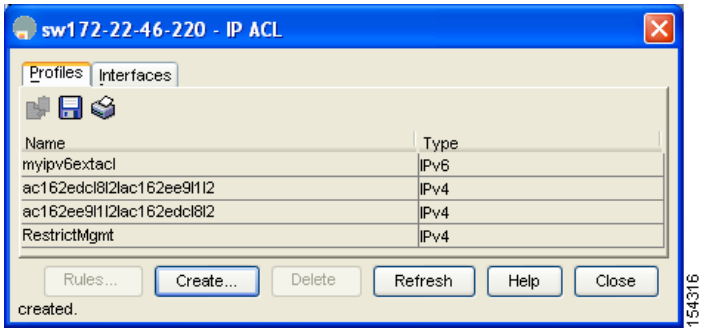

Step 4 Select **RestrictMgmt** and click **Rules**.

You see an empty list of IP filters associated with this IP-ACL.

Step 5 Click **Create** to create the first IP filter.

You see the Create IP Filter dialog box.

Step 6 Create an IP filter to allow management communications from a trusted subnet:

a. Choose the **permit** Action and select **0 IP** from the Protocol drop-down menu.

b. Set the source IP address to 10.67.16.0 and the wildcard mask to 0.0.0.255.

Note The wildcard mask denotes a subset of the IP address you want to match against. This allows a range of addresses to match against this filter.

c. Check the **any** check box for the destination address.

d. Click **Create** to create this IP filter and add it to the RestrictMgmt IP-ACL.

Repeat Step a through Step d to create an IP filter that allows communications for all addresses in the 10.67.16.0/24 subnet.

Step 7 Create an IP filter to allow ICMP ping commands:

a. Choose the **permit** Action and select **1-ICMP** from the Protocol drop-down menu.

b. Check the **any** check box for the source address.

c. Check the **any** check box for the destination address.

Send documentation comments to dcnm-san-docfeedback@cisco.com

d. Select **8 echo** from the ICMPType drop-down menu.

e. Click **Create** to create this IP filter and add it to the RestrictMgmt IP-ACL.

Repeat Step a through Step e to create an IP filter that allows ICMP ping.

Step 8 Create a final IP Filter to block all other traffic:

a. Choose the **deny** Action and select **0 IP** from the Protocol drop-down menu.

b. Check the **any** check box for the source address.

c. Check the **any** check box for the destination address.

d. Click **Create** to create this IP filter and add it to the RestrictMgmt IP-ACL.

e. Click **Close** to close the Create IP Filter dialog box.

Repeat Step a through Step d to create an IP filter that blocks all other traffic.

Step 9 Apply the RestrictMgmt IP ACL to the mgmt0 interface:

a. Click **Security,** select **IP ACL,** and then click the **Interfaces** tab in the IP ACL dialog box.

b. Click **Create**.

You see the Create IP-ACL Interfaces dialog box.

c. Select **mgmt0** from the Interfaces drop-down menu.

d. Select the **inbound** Profile Director.

e. Select **RestrictMgmt** from the ProfileName drop-down menu.

f. Click **Create** to apply the RestrictMgmt IP-ACL to the mgmt0 interface.

Repeat Step a through Step f to apply the new IP-ACL to the mgmt0 interface.

Field Descriptions for IPv4 and IPv6 Access Control Lists

The following are the field descriptions for IPv4 and IPv6 access control lists:

IP ACL Profiles

Field	Description
Name	This is the unique IP protocol filter profile identifier.
Type	This object determines the usage type for this filter profile. This usage type cannot be changed after the profile has been created.

IP ACL Interfaces

Field	Description
ProfileName	This is the unique IP protocol filter profile identifier.

Send documentation comments to dcnm-san-docfeedback@cisco.com

IP Filter Profiles

Field	Description
Action	If it is set to deny, all frames matching this filter will be discarded and scanning of the remainder of the filter list will be aborted. If it is set to permit, all frames matching this filter will be allowed for further bridging or routing processing.
Protocol	This filter protocol value matches the Internet Protocol Number in the frames. These IP numbers are defined in the Network Working Group Request for Comments (RFC) documents. Setting this to '-1' will make the filtering match any IP number.
Address	The source IP address to be matched for this filter. A value of 0 causes all source address to match.
Mask	This is the wildcard mask for the SrcAddress bits that must match. 0 bits in the mask indicate the corresponding bits in the SrcAddress must match in order for the matching to be successful, and 1 bits are don't care bits in the matching. A value of 0 causes only IP frames of source address the same as SrcAddress to match.
PortLow	If Protocol is UDP or TCP, this is the inclusive lower bound of the transport-layer source port range that is to be matched, otherwise it is ignored during matching. This value must be equal to or less than the value specified for this entry in SrcPortHigh.
PortHigh	If Protocol is UDP or TCP, this is the inclusive upper bound of the transport-layer source port range that is to be matched, otherwise it is ignored during matching. This value must be equal to or greater than the value specified for this entry in SrcPortLow. If this value is '0', the UDP or TCP port number is ignored during matching.
Address	The destination IP address to be matched for this filter. A value of 0 causes all source address to match.
Mask	This is the wildcard mask for the DestAddress bits that must match. 0 bits in the mask indicate the corresponding bits in the DestAddress must match in order for the matching to be successful, and 1 bits are don't care bits in the matching. A value of 0 causes only IP frames of source address the same as SrcAddress to match.
PortLow	If Protocol is UDP or TCP, this is the inclusive lower bound of the transport-layer destination port range that is to be matched, otherwise it is ignored during matching. This value must be equal to or less than the value specified for this entry in PortHigh.
PortHigh	If Protocol is UDP or TCP, this is the inclusive upper bound of the transport-layer destination port range that is to be matched, otherwise it is ignored during matching. This value must be equal to or greater than the value specified for this entry in DestPortLow. If this value is '0', the UDP or TCP port number is ignored during matching.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Precedence	The IP traffic precedence parameters in each frame are used to guide the selection of the actual service parameters when transmitting a datagram through a particular network. Most network treats high precedence traffic as more important than other traffic. The IP Precedence value ranges from '0' to '7', with '7' the highest precedence and '0' the lowest precedence. The value '-1' means to match frames of any IP precedence. In other words, the IP precedence parameter will not to checked if this value is '-1'. The precedence level are: • routine(0) - Routine traffic precedence • priority(1) - Priority traffic precedence • immediate(2) - Immediate traffic precedence • flash(3) - Flash traffic precedence • flashOverride(4) - Flash-override traffic precedence • critical(5) - Critical precedence • internet(6) - Internetwork control traffic precedence • network(7) - Network control traffic precedence.
TOS	The Type of Service (TOS) of the frame. The TOS values ranges from '0' to '15'. The value '-1' matches any TOS value.
ICMPType	This filter specifies the ICMP message type to be matched. Setting this value to '-1' will make the filtering match any ICMP message type.
ICMPCode	This filter specifies the ICMP message code to be matched. Setting this value to '-1' will make the filtering match any ICMP code.
TCPEstablished	This filter if true specifies that for TCP protocol, in an established connection, a match occurs if the TCP datagram has the ACK,FIN,PSH,RST,SYN or URG control bits set. If false, a match will occur for any TCP datagram.
LogEnabled	Specifies whether filtered frames will be logged by the filtering subsystem or not. If true, then all frames will be logged. If false, then no frame will be logged.

Send documentation comments to dcnm-san-docfeedback@cisco.com

CHAPTER 6

Configuring Certificate Authorities and Digital Certificates

This chapter includes the following topics:

Information About Certificate Authorities and Digital Certificates

Public Key Infrastructure (PKI) support provides the means for the Cisco MDS 9000 Family switches to obtain and use digital certificates for secure communication in the network. PKI support provides manageability and scalability for IPsec/IKE and SSH.

Certificate Authorities (CAs) manage certificate requests and issue certificates to participating entities such as hosts, network devices, or users. The CAs provide centralized key management for the participating entities.

Digital signatures, based on public key cryptography, digitally authenticate devices and individual users. In public key cryptography, such as the RSA encryption system, each device or user has a key-pair containing both a private key and a public key. The private key is kept secret and is known only to the owning device or user only. However, the public key is known to everybody. The keys act as complements. Anything encrypted with one of the keys can be decrypted with the other. A signature is formed when data is encrypted with a sender's private key. The receiver verifies the signature by decrypting the message with the sender's public key. This process relies on the receiver having a copy of the sender's public key and knowing with a high degree of certainty that it really does belong to the sender and not to someone pretending to be the sender.

This section includes the following topics:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Purpose of CAs and Digital Certificates

CAs manage certificate requests and issue certificates to participating entities such as hosts, network devices, or users. The CAs provide centralized key management for the participating entities.

Digital signatures, based on public key cryptography, digitally authenticate devices and individual users. In public key cryptography, such as the RSA encryption system, each device or user has a key-pair containing both a private key and a public key. The private key is kept secret and is known only to the owning device or user only. However, the public key is known to everybody. The keys act as complements. Anything encrypted with one of the keys can be decrypted with the other. A signature is formed when data is encrypted with a sender's private key. The receiver verifies the signature by decrypting the message with the sender's public key. This process relies on the receiver having a copy of the sender's public key and knowing with a high degree of certainty that it really does belong to the sender and not to someone pretending to be the sender.

Digital certificates link the digital signature to the sender. A digital certificate contains information to identify a user or device, such as the name, serial number, company, department, or IP address. It also contains a copy of the entity's public key. The certificate is itself signed by a CA, a third party that is explicitly trusted by the receiver to validate identities and to create digital certificates.

To validate the signature of the CA, the receiver must first know the CA's public key. Normally this process is handled out-of-band or through an operation done at installation. For instance, most web browsers are configured with the public keys of several CAs by default. The Internet Key Exchange (IKE), an essential component of IPsec, can use digital signatures to scalably authenticate peer devices before setting up security associations.

Trust Model, Trust Points, and Identity CAs

The trust model used in PKI support is hierarchical with multiple configurable trusted CAs. Each participating entity is configured with a list of CAs to be trusted so that the peer's certificate obtained during the security protocol exchanges can be verified, provided it has been issued by one of the locally trusted CAs. To accomplish this, the CA's self-signed root certificate (or certificate chain for a subordinate CA) is locally stored. The process of securely obtaining a trusted CA's root certificate (or the entire chain in the case of a subordinate CA) and storing it locally is called *CA authentication* and is a mandatory step in trusting a CA.

The information about a trusted CA that is locally configured is called the *trust point* and the CA itself is called a *trust point CA*. This information consists of CA certificate (or certificate chain in case of a subordinate CA) and the certificate revocation checking information.

The MDS switch can also enroll with a trust point to obtain an identity certificate (for example, for IPsec/IKE). This trust point is called an *identity CA*.

Send documentation comments to dcnm-san-docfeedback@cisco.com

RSA Key-Pairs and Identity Certificates

You can generate one or more RSA key-pairs and associate each RSA key-pair with a trust point CA where the MDS switch intends to enroll to obtain an identity certificate. The MDS switch needs only one identity per CA, which consists of one key-pair and one identity certificate per CA.

Cisco MDS NX-OS allows you to generate RSA key-pairs with a configurable key size (or modulus). The default key size is 512. You can also configure an RSA key-pair label. The default key label is the switch fully qualified domain name (FQDN).

The following list summarizes the relationship between trust points, RSA key-pairs, and identity certificates:

- A trust point corresponds to a specific CA that the MDS switch trusts for peer certificate verification for any application (such as IKE or SSH).
- An MDS switch can have many trust points and all applications on the switch can trust a peer certificate issued by any of the trust point CAs.
- A trust point is not restricted to a specific application.
- An MDS switch enrolls with the CA corresponding to the trust point to obtain an identity certificate. You can enroll your switch with multiple trust points thereby obtaining a separate identity certificate from each trust point. The identity certificates are used by applications depending upon the purposes specified in the certificate by the issuing CA. The purpose of a certificate is stored in the certificate as certificate extensions.
- When enrolling with a trust point, you must specify an RSA key-pair to be certified. This key-pair must be generated and associated to the trust point before generating the enrollment request. The association between the trust point, key-pair, and identity certificate is valid until it is explicitly removed by deleting the certificate, key-pair, or trust point.
- The subject name in the identity certificate is the fully qualified domain name for the MDS switch.
- You can generate one or more RSA key-pairs on a switch and each can be associated to one or more trust points. But no more than one key-pair can be associated to a trust point, which means only one identity certificate is allowed from a CA.
- If multiple identity certificates (each from a distinct CA) have been obtained, the certificate that an application selects to use in a security protocol exchange with a peer is application specific.
- You do not need to designate one or more trust points for an application. Any application can use any certificate issued by any trust point as long as the certificate purpose satisfies the application requirements.
- You do not need more than one identity certificate from a trust point or more than one key-pair to be associated to a trust point. A CA certifies a given identity (name) only once and does not issue multiple certificates with the same subject name. If you need more than one identity certificate for a CA, then define another trust point for the same CA, associate another key-pair to it, and have it certified, provided CA allows multiple certificates with the same subject name.

Multiple Trusted CA Support

An MDS switch can be configured to trust multiple CAs by configuring multiple trust points and associating each with a distinct CA. With multiple trusted CAs, you do not have to enroll a switch with the specific CA that issued a certificate to a peer. Instead, you configure the switch with multiple trusted CAs that the peer trusts. A switch can then use a configured trusted CA to verify certificates offered by a peer that were not issued by the same CA defined in the identity of the switch.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring multiple trusted CAs allows two or more switches enrolled under different domains (different CAs) to verify the identity of each other when using IKE to set up IPsec tunnels.

PKI Enrollment Support

Enrollment is the process of obtaining an identity certificate for the switch that is used for applications such as IPsec/IKE or SSH. It occurs between the switch requesting the certificate and the certificate authority.

The PKI enrollment process for a switch involves the following steps:

1. Generate an RSA private and public key-pair on the switch.
2. Generate a certificate request in standard format and forward it to the CA.
3. Manual intervention at the CA server by the CA administrator may be required to approve the enrollment request, when it is received by the CA.
4. Receive the issued certificate back from the CA, signed with the CA's private key.
5. Write the certificate into a nonvolatile storage area on the switch (bootflash).

Manual Enrollment Using Cut-and-Paste Method

Cisco MDS NX-OS supports certificate retrieval and enrollment using a manual cut-and-paste method. Cut-and-paste enrollment literally means you must cut and paste the certificate requests and resulting certificates between the switch and the CA, as follows:

1. Create an enrollment certificate request, which is displayed in base64-encoded text form.
2. Cut and paste the encoded certificate request text in an e-mail message or in a web form and send it to the CA.
3. Receive the issued certificate (in base64-encoded text form) from the CA in an e-mail message or in a web browser download.
4. Cut and paste the issued certificate to the switch using the certificate import facility.

Note DCNM-SAN does not support cut and paste. Instead, it allows the enrollment request (certificate signing request) to be saved in a file to be sent manually to the CA.

Multiple RSA Key-Pair and Identity CA Support

Multiple identity CA support enables the switch to enroll with more than one trust point. This results in multiple identity certificates; each from a distinct CA. This allows the switch to participate in IPsec and other applications with many peers using certificates issued by appropriate CAs that are acceptable to those peers.

The multiple RSA key-pair support feature allows the switch to maintain a distinct key pair for each CA with which it is enrolled. Thus, it can match policy requirements for each CA without conflicting with the requirements specified by the other CAs, such as key length. The switch can generate multiple RSA key-pairs and associate each key-pair with a distinct trust point. Thereafter, when enrolling with a trust point, the associated key-pair is used to construct the certificate request.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Peer Certificate Verification

The PKI support on an MDS switch provides the means to verify peer certificates. The switch verifies certificates presented by peers during security exchanges pertaining to applications, such as IPsec/IKE and SSH. The applications verify the validity of the peer certificates presented to them. The peer certificate verification process involves the following steps:

- Verifies that the peer certificate is issued by one of the locally trusted CAs.
- Verifies that the peer certificate is valid (not expired) with respect to current time.
- Verifies that the peer certificate is not yet revoked by the issuing CA.

For revocation checking, two methods are supported: certificate revocation list (CRL) and Online Certificate Status Protocol (OCSP). A trust point uses one or both of these methods to verify that the peer certificate has not been revoked.

CRL Downloading, Caching, and Checking Support

Certificate revocation lists (CRLs) are maintained by CAs to give information of prematurely revoked certificates, and the CRLs are published in a repository. The download URL is made public and also specified in all issued certificates. A client verifying a peer's certificate should obtain the latest CRL from the issuing CA and use it to determine if the certificate has been revoked. A client can cache the CRLs of some or all of its trusted CAs locally and use them later if necessary until the CRLs expire.

Cisco MDS NX-OS allows the manual configuration of pre-downloaded of CRLs for the trust points, and then caches them in the switch bootflash (cert-store). During the verification of a peer certificate by IPsec or SSH, the issuing CA's CRL is consulted only if the CRL has already been cached locally and the revocation checking is configured to use CRL. Otherwise, CRL checking is not performed and the certificate is considered to be not revoked if no other revocation checking methods are configured. This mode of CRL checking is called CRL optional.

OCSP Support

Online Certificate Status Protocol (OCSP) facilitates online certificate revocation checking. You can specify an OCSP URL for each trust point. Applications choose the revocation checking mechanisms in a specified order. The choices are CRL, OCSP, none, or a combination of these methods.

Import and Export Support for Certificates and Associated Key-Pairs

As part of the CA authentication and enrollment process, the subordinate CA certificate (or certificate chain) and identity certificates can be imported in standard PEM (base64) format.

The complete identity information in a trust point can be exported to a file in the password-protected PKCS#12 standard format. It can be later imported to the same switch (for example, after a system crash) or to a replacement switch. The information in a PKCS#12 file consists of the RSA key-pair, the identity certificate, and the CA certificate (or chain).

Maximum Limits

Table 6-1 lists the maximum limits for CAs and digital certificate parameters.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 6-1 Maximum Limits for CA and Digital Certificate

Feature	Maximum Limit
Trust points declared on a switch	16
RSA key-pairs generated on a switch	16
Identity certificates configured on a switch	16
Certificates in a CA certificate chain	10
Trust points authenticated to a specific CA	10

Default Settings

Table 6-2 lists the default settings for CAs and digital certificate parameters.

Table 6-2 Default CA and Digital Certificate Parameters

Parameters	Default
Trust point	None
RSA key-pair	None
RSA key-pair label	Switch FQDN
RSA key-pair modulus	512
RSA key-pair exportable	Yes
Revocation check method of trust point	CRL

Configuring CAs and Digital Certificates

This section describes the tasks you must perform to allow CAs and digital certificates your Cisco MDS switch device to interoperate. This section includes the following sections:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring the Host Name and IP Domain Name

You must configure the host name and IP domain name of the switch if they are not already configured. This is required because switch FQDN is used as the subject in the identity certificate. Also, the switch FQDN is used as a default key label when none is specified during key-pair generation. For example, a certificate named SwitchA.example.com is based on a switch host name of SwitchA and a switch IP domain name of example.com.

Caution Changing the host name or IP domain name after generating the certificate can invalidate the certificate.

Detailed Steps

To configure the host name and IP domain name, refer to the *Cisco MDS 9000 NX-OS Fundamental Configuration Guide*.

Generating an RSA Key Pair

RSA key-pairs are used to sign and/or encrypt and decrypt the security payload during security protocol exchanges for applications such as IKE/IPsec and SSH, and they are required before you can obtain a certificate for your switch.

Detailed Steps

To generate an RSA key-pair, follow these steps:

Step 1 Expand **Switches > Security** and then select **PKI** in the Information pane.

Step 2 Click the **RSA Key-Pair** tab.

Step 3 Click **Create Row**.

Step 4 Select the switches for which you want to create the RSA key-pair.

Step 5 Assign a name to the RSA key-pair.

Step 6 Select the Size or modulus values. Valid modulus values are 512, 768, 1024, 1536, and 2048.

Note The security policy (or requirement) at the local site (MDS switch) and at the CA (where enrollment is planned) are considered in deciding the appropriate key modulus.

Note The maximum number of key-pairs you can configure on a switch is 16.

Step 7 Check the **Exportable** check box if you want the key to be exportable.

Caution The exportability of a key-pair cannot be changed after key-pair generation.

Note Only exportable key-pairs can be exported in PKCS#12 format.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 8 Click **Create** to create the RSA key pair.

Creating a Trust Point CA Association

Detailed Steps

To create a trust point CA association, follow these steps:

Step 1 Expand **Switches** > **Security,** and then select **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point** tab in the Information Pane.

Step 3 Click **Create Row**.

Step 4 Select the switch for which you are creating the trust point CA from the **Switch** drop-down menu.

Step 5 Assign a name to the trust point CA.

Step 6 Select a key-pair name to be associated with this trust point for enrollment. It was generated earlier in the "Generating an RSA Key Pair" section on page 6-7. Only one RSA key-pair can be specified per CA.

Step 7 From the RevokeCheckMethod drop-down menu, select the certificate revocation method that you would like to use. You can use CRL, OCSP, CRL OCSP, or OCSP CRL to check for certificate revocation.

The CRL OCSP option checks for revoked certificates first in the locally stored CRL. If not found, the switch uses OCSP to check the revoked certificates on the URL specified in Step 7.

Step 8 Enter the OCSP URL if you selected an OCSP certificate revocation method.

Note The OSCP URL must be configured before configuring the revocation checking method.

Step 9 Click **Create** to successfully create the trust point CA.

Copying Files to Bootflash

Detailed Steps

To copy files to bootflash using Device Manager, follow these steps:

Step 1 Choose **Admin** > **Flash Files**.

Step 2 Select bootflash in the Device field.

Step 3 Click **Copy**.

Step 4 Select **tftp** as the Protocol field.

Step 5 Click the **Browse** button to locate the appropriate file to copy to bootflash.

Step 6 Click **Apply** to apply these changes.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Authenticating the CA

The configuration process of trusting a CA is complete only when the CA is authenticated to the MDS switch. The switch must authenticate the CA. It does this by obtaining the self-signed certificate of the CA in PEM format, which contains the public key of the CA. Because the certificate of the CA is self-signed (the CA signs its own certificate) the public key of the CA should be manually authenticated by contacting the CA administrator to compare the fingerprint of the CA certificate.

Note If the CA being authenticated is not a self-signed CA (that is, it is a subordinate CA to another CA, which itself may be a subordinate to yet another CA, and so on, finally ending in a self-signed CA), then the full list of the CA certificates of all the CAs in the certification chain needs to be input during the CA authentication step. This is called the *CA certificate chain* of the CA being authenticated. The maximum number of certificates in a CA certificate chain is 10.

Detailed Steps

To authenticate a CA, follow these steps:

Step 1 Expand **Switches > Security,** and then select **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information pane.

Step 3 From the Command field drop-down menu, select the appropriate option.

Available options are **caauth**, **cadelete**, **certreq**, **certimport**, **certdelete**, **pkcs12import**, and **pkcs12export**. The **caauth** option is provided to authenticate a CA and install its CA certificate or certificate chain in a trust point.

Step 4 Click the **Browse** button in the URL field and select the appropriate import certificate file from the **Bootflash Files** dialog box. It is the file name containing the CA certificate or chain in the bootflash:filename format.

Note You can authenticate a maximum of 10 trust points to a specific CA.

Note If you do not see the required file in the Import Certificate dialog box, make sure that you copy the file to bootflash. See "Copying Files to Bootflash" section on page 8.

Step 5 Click **Apply Changes** to save the changes.

Authentication is then confirmed or not confirmed depending on whether or not the certificate can be accepted after manual verification of its fingerprint.

Confirming CA Authentication

As mentioned in step 5 of "Authenticating the CA" section on page 6-9, CA authentication is required to be followed by CA confirmation in order to accept the CA certificate based on its fingerprint verification.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To confirm CA authentication, follow these steps:

Step 1 Expand **Switches > Security,** and then select **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information Pane.

Step 3 Make a note of the CA certificate fingerprint displayed in the IssuerCert FingerPrint column for the trust point row in question. Compare the CA certificate fingerprint with the fingerprint already communicated by the CA (obtained from the CA web site).

If the fingerprints match exactly, accept the CA with the **certconfirm** command in the Command drop-down menu. Otherwise, reject the CA with the **certnoconfirm** command.

Step 4 If you selected **certconfirm** in step 3, click **Command** and select the **certconfirm** action from the drop-down menu. Click **Apply Changes.**

If you selected **certnoconfirm** in step 3, click **Command** and select the **certnoconfirm** action drop-down menu. Click **Apply Changes**.

Configuring Certificate Revocation Checking Methods

During security exchanges with a client (for example, an IKE peer or SSH user), the MDS switch performs the certificate verification of the peer certificate sent by the client and the verification process may involve certificate revocation status checking.

You can use different methods for checking for revoked sender certificates. You can configure the switch to check the CRL downloaded from the CA (see the "Configuring a CRL" section on page 6-14), you can use OSCP if it is supported in your network, or both. Downloading the CRL and checking locally does not generate traffic in your network. However, certificates can be revoked between downloads and your switch would not be aware of the revocation. OCSP provides the means to check the current CRL on the CA. However, OCSP can generate network traffic that can impact network efficiency. Using both local CRL checking and OCSP provides the most secure method for checking for revoked certificates.

Note You must authenticate the CA before configuring certificate revocation checking.

DCNM-SAN allows you to configure certificate revocation checking methods when you are creating a trust point CA. See "Creating a Trust Point CA Association" section on page 6-8.

Generating Certificate Requests

You must generate a request to obtain identity certificates from the associated trust point CA for each of your switch's RSA key pairs. You must then cut and paste the displayed request into an e-mail message or in a website form for the CA.

Detailed Steps

To generate a request for signed certificates from the CA, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 1 Expand **Switches > Security,** and then select **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information pane.

Step 3 Select the **certreq** option from the Command drop-down menu.

This action generates a PKCS#10 certificate signing request (CSR) needed for an identity certificate from the CA corresponding to this trust point entry. This entry requires an associated key-pair. The CA certificate or certificate chain should already be configured through the **caauth** action. See "Authenticating the CA" section on page 6-9.

Step 4 Enter the output file name for storing the generated certificate request.

It will be used to store the CSR generated in PEM format. Use the format bootflash:filename. This CSR should be submitted to the CA to get the identity certificate. Once the identity certificate is obtained, it should be installed in this trust point. See "Installing Identity Certificates" section on page 6-11.

Step 5 Enter the *challenge* password to be included in the CSR.

Note The challenge password is not saved with the configuration. This password is required in the event that your certificate needs to be revoked, so you must remember this password.

Step 6 Click **Apply Changes** to save the changes.

Installing Identity Certificates

You receive the identity certificate from the CA by e-mail or through a web browser in base64 encoded text form. You must install the identity certificate from the CA by cutting and pasting the encoded text using the CLI import facility.

Detailed Steps

To install an identity certificate received from the CA, follow these steps:

Step 1 Expand **Switches > Security,** and then select **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab, in the Information pane.

Step 3 Select the **certimport** option from the Command drop-down menu to import an identity certificate in this trust point. The identity certificate is obtained from the corresponding CA for a CSR that was generated previously (see "Generating Certificate Requests" section on page 6-10).

Note The identity certificate should be available in PEM format in a file in bootflash.

Step 4 Enter the name of the certificate file that should have been copied to bootflash in the URL field in the bootflash:filename format.

Step 5 Click **Apply Changes** to save your changes.

Send documentation comments to dcnm-san-docfeedback@cisco.com

If successful, the values of the identity certificate and its related objects, like the certificate file name, are automatically updated with the appropriate values as per the corresponding attributes in the identity certificate.

Saving Your Configuration

Save your work when you make configuration changes or the information is lost when you exit.

Detailed Steps

To save your configuration, follow these steps:

Step 1 Expand **Switches,** and then select **Copy Configuration** in the Physical Attributes pane.

Step 2 Select the switch configuration including the RSA key pairs and certificates.

Step 3 Click **Apply Changes** to save the changes.

Ensuring Trust Point Configurations Persist Across Reboots

The trust point configuration is a normal Cisco NX-OS configuration that persists across system reboots only if you copy it explicitly to the startup configuration. The certificates, key pairs, and CRL associated with a trust point are automatically persistent if you have already copied the trust point configuration in the startup configuration. Conversely, if the trust point configuration is not copied to the startup configuration, the certificates, key pairs, and CRL associated with it are not persistent since they require the corresponding trust point configuration after a reboot. Always copy the running configuration to the startup configuration to ensure that the configured certificates, key pairs, and CRLs are persistent. Also, save the running configuration after deleting a certificate or key pair to ensure that the deletions are permanent.

The certificates and CRL associated with a trust point automatically become persistent when imported (that is, without an explicitly copying to the startup configuration) if the specific trust point is already saved in startup configuration.

We also recommend that you create a password-protected backup of the identity certificates and save it to an external server (see the "Exporting and Importing Identity Information in PKCS#12 Format" section on page 6-13).

Note Copying the configuration to an external server does include the certificates and key pairs.

Monitoring and Maintaining CA and Certificates Configuration

The tasks in the section are optional. This section includes the following topics:

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Deleting Certificates from the CA Configuration, page 6-14
- Deleting RSA Key Pairs from Your Switch, page 6-15

Exporting and Importing Identity Information in PKCS#12 Format

You can export the identity certificate along with the RSA key pair and CA certificate (or the entire chain in the case of a subordinate CA) of a trust point to a PKCS#12 file for backup purposes. You can later import the certificate and RSA key pair to recover from a system crash on your switch or when you replace the supervisor modules.

Note Only the **bootflash:***filename* format local syntax is supported when specifying the export and import URL.

Detailed Steps

To export a certificate and key pair to a PKCS#12-formatted file, follow these steps:

Step 1 Expand **Switches > Security,** and then select **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information Pane (see Figure 6-1).

Step 3 Select the **pkcs12export** option in the Command drop-down menu to export the key pair, identity certificate, and the CA certificate or certificate chain in PKCS#12 format from the selected trust point.

Figure 6-1 Pkcs12export Option Exports a Key Pair

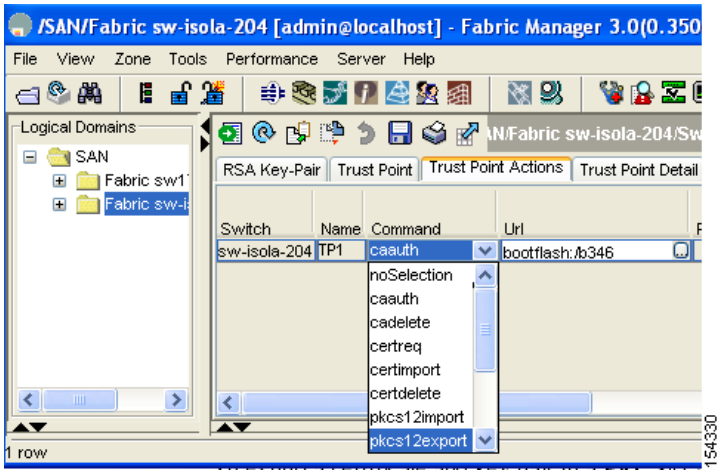

Step 4 Enter the output file name as bootflash:filename to store the exported PKCS#12 identity.

Step 5 Enter the required password. The password is set for encoding the PKCS#12 data. On successful completion, the exported data is available in bootflash in the specified file.

Step 6 Click **Apply Changes** to save the changes.

To import a certificate and key pair formatted as a PKCS#12 formatted file, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 1 Expand **Switches** > **Security,** and then select **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information pane (see Figure 6-1).

Step 3 Select the **pkcs12import** option from the Command drop-down menu to import the key-pair, identity certificate, and the CA certificate or certificate chain in the PKCS#12 format to the selected trust point.

Step 4 Enter the input in the bootflash:filename format containing the PKCS#12 identity.

Step 5 Enter the required password. The password is set for decoding the PKCS#12 data. On completion, the imported data is available in bootflash in the specified file.

Step 6 Click **Apply Changes** to save the changes.

On completion the trust point is created in the RSA key-pair table corresponding to the imported key pair. The certificate information is updated in the trust point.

Note The trust point must be empty (with no RSA key pair associated with it and no CA is associated with it using CA authentication) for the PKCS#12 file import to succeed.

Configuring a CRL

Detailed Steps

To configure the CRL from a file to a trust point, follow these steps:

Step 1 Click **Switches** > **Security** > **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information pane.

Step 3 Select the **crlimport** option from the Command drop-down menu to import the CRL to the selected trust point.

Step 4 Enter the input file name with the CRL in the bootflash:filename format, in the URL field.

Step 5 Click **Apply Changes** to save the changes.

Deleting Certificates from the CA Configuration

You can delete the identity certificates and CA certificates that are configured in a trust point. You must first delete the identity certificate, followed by the CA certificates. After deleting the identity certificate, you can disassociate the RSA key-pair from a trust point. The certificate deletion is necessary to remove expired or revoked certificates, certificates whose key-pairs are compromised (or suspected to be compromised) or CAs that are no longer trusted.

Detailed Steps

To delete the CA certificate (or the entire chain in the case of a subordinate CA) from a trust point using DCNM-SAN, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 1 Click **Switches > Security > PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information pane.

Step 3 Select the **cadelete** option from the Command drop-down menu to delete the identity certificate from a trust point.

Note If the identity certificate being deleted is the last-most or only identity certificate in the device, you must use the **forcecertdelete** action to delete it. This ensures that the administrator does not mistakenly delete the last-most or only identity certificate and leave the applications (such as IKE and SSH) without a certificate to use.

Step 4 Click **Apply Changes** to save the changes.

To delete the identity certificate, click the **Trust Point Actions** tab and select the **certdelete** or **forcecertdelete** in the Command drop-down menu.

Deleting RSA Key Pairs from Your Switch

Under certain circumstances you may want to delete your switch's RSA key pairs. For example, if you believe the RSA key pairs were compromised in some way and should no longer be used, you should delete the key pairs.

Detailed Steps

To delete RSA key pairs from your switch, follow these steps:

Step 1 Expand **Switches > Security,** and then select **PKI** in the Physical Attributes pane.

Step 2 Click the **RSA Key-Pair** tab in the Information pane.

Step 3 Click **Delete Row**.

Step 4 Click **Yes** or **No** in the Confirmation dialog box.

Note After you delete RSA key pairs from a switch, ask the CA administrator to revoke your switch's certificates at the CA. You must supply the challenge password you created when you originally requested the certificates. See "Generating Certificate Requests" section on page 6-10.

Configuration Examples

This section shows an example of the tasks you can use to configure certificates and CRLs on the Cisco MDS 9000 Family switches using the Microsoft Windows Certificate server.

This section includes the following topics:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring Certificates on the MDS Switch

To configure certificates on an MDS switch, follow these steps:

Step 1 Choose **Switches** and set the LogicalName field to configure the switch host name.

Step 2 Choose **Switches > Interfaces > Management > DNS** and set the DefaultDomainName field to configure.

Step 3 To create an RSA key-pair for the switch, follow these steps:

a. Choose **Switches > Security > PKI** and select the RSA **Key-Pair** tab.

b. Click **Create Row** and set the name and size field.

c. Check the **Exportable** check box and click **Create**.

Step 4 To create a trust point and associate the RSA key-pairs with it, follow these steps:

a. Choose **Switches > Security > PKI** and select the **Trustpoints** tab.

b. Click **Create Row** and set the TrustPointName field.

c. Select the RSA key-pairs from the KeyPairName drop-down menu.

d. Select the certificates revocation method from the CARevoke drop-down menu.

e. Click **Create**.

Step 5 Choose **Switches > Copy Configuration** and click **Apply Changes** to copy the running to startup configuration and save the trustpoint and key pair.

Step 6 Download the CA certificate from the CA that you want to add as the trustpoint CA.

Step 7 To authenticate the CA that you want to enroll to the trust point, follow these steps:.

a. Using Device Manager, choose **Admin > Flash Files** and select **Copy** and tftp copy the CA certificate to bootflash.

b. Using DCNM-SAN, choose **Switches > Security > PKI** and select the **TrustPoint Actions** tab.

c. Select **cauth** from the Command drop-down menu.

d. Click **...** in the URL field and select the CA certificate from bootflash.

e. Click **Apply Changes** to authenticate the CA that you want to enroll to the trust point.

f. Click the **Trust Point Actions** tab in the Information Pane.

g. Make a note of the CA certificate fingerprint displayed in the IssuerCert FingerPrint column for the trust point row in question. Compare the CA certificate fingerprint with the fingerprint already communicated by the CA (obtained from the CA web site). If the fingerprints match exactly, accept the CA by performing the **certconfirm** trust point action. Otherwise, reject the CA by performing the **certnoconfirm** trust point action.

Send documentation comments to dcnm-san-docfeedback@cisco.com

h. If you select **certconfirm** in step g, click the **Trust Point Actions** tab, select **certconfirm** from the command drop-down menu, and then click **Apply Changes**.

i. If you select **certnoconfirm** in step g, click the **Trust Point Actions** tab, select the **certnoconfirm** from the command drop-down menu and then click **Apply Changes**.

Step 8 To generate a certificate request for enrolling with that trust point, follow these steps:

a. Click the **Trust Point Actions** tab in the Information pane.

b. Select **certreq** from the Command drop-down menu. This generates a PKCS#10 certificate signing request (CSR) needed for an identity certificate from the CA corresponding to this trust point entry.

c. Enter the output file name for storing the generated certificate request. It should be specified in the bootflash:filename format and will be used to store the CSR generated in PEM format.

d. Enter the *challenge* password to be included in the CSR. The challenge password is not saved with the configuration. This password is required in the event that your certificate needs to be revoked, so you must remember this password.

e. Click **Apply Changes** to save the changes.

Step 9 Request an identity certificate from the CA.

Note The CA may require manual verification before issuing the identity certificate.

Step 10 To import the identity certificate, follow these steps:

a. Using Device Manager, choose **Admin > Flash Files** and select **Copy** and use TFTP to copy the CA certificate to bootflash.

b. Using DCNM-SAN, choose **Switches > Security > PKI** and click the **TrustPoint Actions** tab.

c. Select the **certimport** option from the Command drop-down menu to import an identity certificate in this trust point.

Note The identity certificate should be available in PEM format in a file in bootflash.

d. Enter the name of the certificate file which was copied to bootflash, in the URL field in the bootflash:filename format.

e. Click **Apply Changes** to save your changes.

If successful, the values of the identity certificate and its related objects, like the certificate file name, are automatically updated with the appropriate values as per the corresponding attributes in the identity certificate.

Downloading a CA Certificate

To download a CA certificate from the Microsoft Certificate Services web interface, follow these steps:

Step 1 Click the **Retrieve the CA certificate or certificate revocation task** radio button in the Microsoft Certificate Services web interface and click the **Next button**.

Step 2 Select the CA certificate file to download from the displayed list. Click the **Base 64 encoded** radio button, and choose the **Download CA certificate** link.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 3 Click the **Open** button in the File Download dialog box.

Step 4 Click the **Copy to File** button in the Certificate dialog box and click **OK**.

Step 5 Select the **Base-64 encoded X.509 (CER)** on the Certificate Export Wizard dialog box and click **Next**.

Step 6 Enter the destination file name in the File name: text box on the Certificate Export Wizard dialog box and click **Next**.

Step 7 Click the **Finish** button on the Certificate Export Wizard dialog box.

Step 8 Display the CA certificate stored in Base-64 (PEM) format using the Microsoft Windows **type** command.

Requesting an Identity Certificate

To request an identify certificate from a Microsoft Certificate server using a PKCS#10 certificate signing request (CRS), follow these steps:

Step 1 Click the Request an identity certificate radio button on the Microsoft Certificate Services web interface and click **Next**.

Step 2 Click the **Advanced Request** radio button and click **Next**.

Step 3 Click the **Submit a certificate request using a base64 encoded PKCS#10 file or a renewal request using a base64 encoded PKCS#7 file** radio button and click **Next**.

Step 4 Paste the base64 PKCS#10 certificate request in the Saved Request text box and click **Next**.

The certificate request is copied from the MDS switch console (see the "Generating Certificate Requests" section on page 6-10 and "Configuring Certificates on the MDS Switch" section on page 6-16).

Step 5 Wait one or two days until the certificate is issued by the CA administrator.

Step 6 The CA administrator approves the certificate request.

Step 7 Click the **Check on a pending certificate** radio button on the Microsoft Certificate Services web interface and click **Next**.

Step 8 Select the certificate request you want to check and click **Next**.

Step 9 Select **Base 64 encoded** and click the **Download CA certificate** link.

Step 10 Click **Open** on the File Download dialog box.

Step 11 Click the **Details** tab on the Certificate dialog and click the **Copy to File** button. Click the **Base-64 encoded X.509 (.CER)** radio button on the Certificate Export Wizard dialog box and click **Next**.

Step 12 Enter the destination file name in the File name: text box on the Certificate Export Wizard dialog box, then click **Next**.

Step 13 Click **Finish**.

Step 14 Display the identity certificate in base64-encoded format using the Microsoft Windows **type** command.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Revoking a Certificate

To revoke a certificate using the Microsoft CA administrator program, follow these steps:

Step 1 Click the **Issued Certificates** folder on the Certification Authority tree. From the list, right-click the certificate you want to revoke.

Step 2 Select **All Tasks > Revoke Certificate**.

Step 3 Select a reason for the revocation from the Reason code drop-down list, and click **Yes**.

Step 4 Click the **Revoked Certificates** folder to list and verify the certificate revocation.

Generating and Publishing the CRL

To generate and publish the CRL using the Microsoft CA administrator program, follow these steps:

Step 1 Select **Action > All Tasks > Publish** on the Certification Authority screen.

Step 2 Click **Yes** on the Certificate Revocation List dialog box to publish the latest CRL.

Downloading the CRL

To download the CRL from the Microsoft CA website, follow these steps:

Step 1 Click **Request the CA certificate or certificate revocation list** radio button on the Microsoft Certificate Services web interface and click **Next**.

Step 2 Click the **Download latest certificate revocation list** link.

Step 3 Click **Save** in the File Download dialog box.

Step 4 Enter the destination file name in the Save As dialog box and click **Save**.

Step 5 Display the CRL using the Microsoft Windows **type** command.

Importing the CRL

To import the CRL to the trust point corresponding to the CA, follow these steps:

Step 1 Click **Switches** > **Security** > **PKI** in the Physical Attributes pane.

Step 2 Click the **Trust Point Actions** tab in the Information pane.

Step 3 Select the **crlimport** option from the Command drop-down menu to import the CRL to the selected trust point.

Step 4 Enter the input file name with the CRL in the bootflash:filename format, in the URL field.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 5 Click **Apply Changes** to save the changes.

Note The identity certificate for the switch that was revoked (serial number 0A338EA1000000000074) is listed at the end.

Send documentation comments to dcnm-san-docfeedback@cisco.com

CHAPTER 7

Configuring IPsec Network Security

IP security (IPsec) protocol is a framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. It is developed by the Internet Engineering Task Force (IETF). IPsec provides security services at the IP layer, including protecting one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host. The overall IPsec implementation is the latest version of RFC 2401. Cisco NX-OS IPsec implements RFC 2402 through RFC 2410.

IPsec uses the Internet Key Exchange (IKE) protocol to handle protocol and algorithm negotiation and to generate the encryption and authentication keys used by IPsec. While IKE can be used with other protocols, its initial implementation is with the IPsec protocol. IKE provides authentication of the IPsec peers, negotiates IPsec security associations, and establishes IPsec keys. IKE uses RFCs 2408, 2409, 2410, and 2412, and additionally implements the draft-ietf-ipsec-ikev2-16.txt draft.

The term IPsec is sometimes used to describe the entire protocol of IPsec data services and IKE security protocols and is other times used to describe only the data services.

This chapter includes the following topics:

Information About IPsec Network Security

IP security (IPsec) protocol is a framework of open standards that provides data confidentiality, data integrity, and data authentication between participating peers. It is developed by the Internet Engineering Task Force (IETF). IPsec provides security services at the IP layer, including protecting one or more data flows between a pair of hosts, between a pair of security gateways, or between a security gateway and a host. The overall IPsec implementation is the latest version of RFC 2401. Cisco NX-OS IPsec implements RFC 2402 through RFC 2410.

Send documentation comments to dcnm-san-docfeedback@cisco.com

IPsec uses the Internet Key Exchange (IKE) protocol to handle protocol and algorithm negotiation and to generate the encryption and authentication keys used by IPsec. While IKE can be used with other protocols, its initial implementation is with the IPsec protocol. IKE provides authentication of the IPsec peers, negotiates IPsec security associations, and establishes IPsec keys. IKE uses RFCs 2408, 2409, 2410, and 2412, and additionally implements the draft-ietf-ipsec-ikev2-16.txt draft.

IPsec provides security for transmission of sensitive information over unprotected networks such as the Internet. IPsec acts at the network layer, protecting and authenticating IP packets between participating IPsec devices (peers).

Note IPsec is not supported by the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter.

IPsec provides security for transmission of sensitive information over unprotected networks such as the Internet. IPsec acts at the network layer, protecting and authenticating IP packets between participating IPsec devices (peers).

IPsec provides the following network security services. In general, the local security policy dictates the use of one or more of these services between two participating IPsec devices:

- Data confidentiality—The IPsec sender can encrypt packets before transmitting them across a network.
- Data integrity—The IPsec receiver can authenticate packets sent by the IPsec sender to ensure that the data has not been altered during transmission.
- Data origin authentication—The IPsec receiver can authenticate the source of the IPsec packets sent. This service is dependent upon the data integrity service.
- Anti-replay protection—The IPsec receiver can detect and reject replayed packets.

Note The term *data authentication* is generally used to mean data integrity and data origin authentication. Within this chapter it also includes anti-replay services, unless otherwise specified.

With IPsec, data can be transmitted across a public network without fear of observation, modification, or spoofing. This enables applications such as Virtual Private Networks (VPNs), including intranets, extranets, and remote user access.

IPsec as implemented in Cisco NX-OS software supports the Encapsulating Security Payload (ESP) protocol. This protocol encapsulates the data to be protected and provides data privacy services, optional data authentication, and optional anti-replay services.

Note The Encapsulating Security Payload (ESP) protocol is a header inserted into an existing TCP/IP packet, the size of which depends on the actual encryption and authentication algorithms negotiated. To avoid fragmentation, the encrypted packet fits into the interface maximum transmission unit (MTU). The path MTU calculation for TCP takes into account the addition of ESP headers, plus the outer IP header in tunnel mode, for encryption. The MDS switches allow 100 bytes for packet growth for IPsec encryption.

Note When using IPsec and IKE, each Gigabit Ethernet interface on the IPS module (either on 14+2 LC or 18+4 LC) must be configured in its own IP subnet. If there are multiple Gigabit Ethernet interfaces configured with IP address or network-mask in the same IP subnet, IKE packets may not be sent to the right peer and thus IPsec tunnel will not come up.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 7-1 shows different IPsec scenarios.

Figure 7-1 FCIP and iSCSI Scenarios Using MPS-14/2 Modules

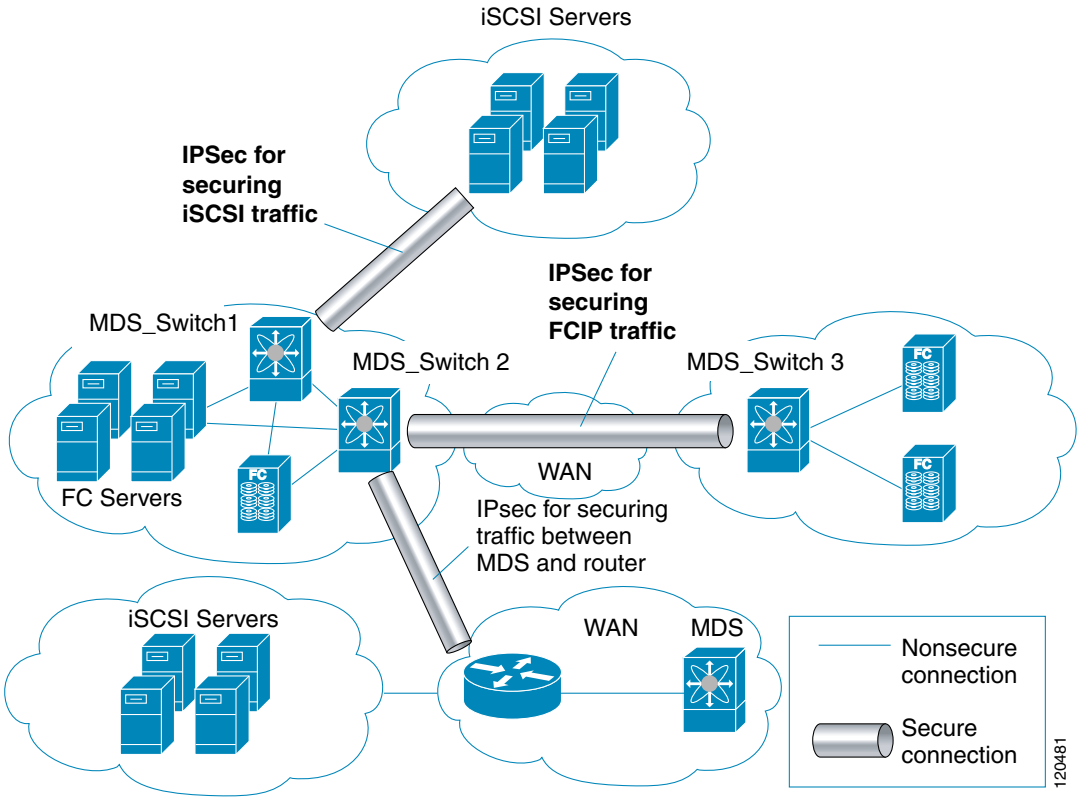

This section includes the following topics:

- About IKE, page 7-4
- IPsec Compatibility, page 7-4
- IPsec and IKE Terminology, page 7-5
- Supported IPsec Transforms and Algorithms, page 7-6
- Supported IKE Transforms and Algorithms, page 7-6
- About IPsec Digital Certificate Support, page 7-7
- About IKE Initialization, page 7-9
- About the IKE Domain, page 7-10
- About IKE Tunnels, page 7-10
- About IKE Policy Negotiation, page 7-10
- Optional IKE Parameter Configuration, page 7-11
- About Crypto IPv4-ACLs, page 7-12
- About Transform Sets in IPsec, page 7-14
- About Crypto Map Entries, page 7-15
- About SA Lifetime Negotiation, page 7-16
- About the AutoPeer Option, page 7-17

Send documentation comments to dcnm-san-docfeedback@cisco.com

About IKE

IKE automatically negotiates IPsec security associations and generates keys for all switches using the IPsec feature. Specifically, IKE provides these benefits:

- Allows you to refresh IPsec SAs.
- Allows IPsec to provide anti-replay services.
- Supports a manageable, scalable IPsec configuration.
- Allows dynamic authentication of peers.

IKE is not supported on the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeSystem.

IPsec Compatibility

IPsec features are compatible with the following Cisco MDS 9000 Family hardware:

- Cisco 18/4-port Multi-Service Module (MSM-18/4) modules and MDS 9222i Module-1 modules.
- Cisco 14/2-port Multiprotocol Services (MPS-14/2) modules in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors
- Cisco MDS 9216i Switch with the 14/2-port multiprotocol capability in the integrated supervisor module. Refer to the *Cisco MDS 9200 Series Hardware Installation Guide* for more information on the Cisco MDS 9216i Switch.
- The IPsec feature is not supported on the management interface.

IPsec features are compatible with the following fabric setup:

- Two connected Cisco MDS 9200 Switches or Cisco MDS 9500 Directors running Cisco MDS SAN-OS Release 2.0(1b) or later, or Cisco NX-OS Release 4.1(1).
- A Cisco MDS 9200 Switches or Cisco MDS 9500 Directors running Cisco MDS SAN-OS Release 2.0(1b) or later, or Cisco NX-OS Release 4.1(1) connected to any IPsec compliant device.
- The following features are not supported in the Cisco NX-OS implementation of the IPsec feature:
  - Authentication Header (AH)
  - Transport mode
  - Security association bundling
  - Manually configuring security associations
  - Per host security association option in a crypto map
  - Security association idle timeout
  - Dynamic crypto maps

Send documentation comments to dcnm-san-docfeedback@cisco.com

Note Any reference to crypto maps in this document, only refers to static crypto maps.

IPsec and IKE Terminology

The terms used in this chapter are explained in this section.

- Security association (SA)— An agreement between two participating peers on the entries required to encrypt and decrypt IP packets. Two SAs are required for each peer in each direction (inbound and outbound) to establish bidirectional communication between the peers. Sets of bidirectional SA records are stored in the SA database (SAD). IPsec uses IKE to negotiate and bring up SAs. Each SA record includes the following information:
  - Security parameter index (SPI)—A number which, together with a destination IP address and security protocol, uniquely identifies a particular SA. When using IKE to establish the SAs, the SPI for each SA is a pseudo-randomly derived number.
  - Peer—A switch or other device that participates in IPsec. For example, a Cisco MDS switch or other Cisco routers that support IPsec.
  - Transform—A list of operations done to provide data authentication and data confidentiality. For example, one transform is the ESP protocol with the HMAC-MD5 authentication algorithm.
  - Session key—The key used by the transform to provide security services.
  - Lifetime—A lifetime counter (in seconds and bytes) is maintained from the time the SA is created. When the time limit expires the SA is no longer operational and, if required, is automatically renegotiated (rekeyed).
  - Mode of operation—Two modes of operation are generally available for IPsec: tunnel mode and transport mode. The Cisco NX-OS implementation of IPsec only supports the tunnel mode. The IPsec tunnel mode encrypts and authenticates the IP packet, including its header. The gateways encrypt traffic on behalf of the hosts and subnets.
    The Cisco NX-OS implementation of IPsec does not support transport mode.

  Note The term *tunnel mode* is different from the term *tunnel*, which is used to indicate a secure communication path between two peers, such as two switches connected by an FCIP link.

- Anti-replay—A security service where the receiver can reject old or duplicate packets to protect itself against replay attacks. IPsec provides this optional service by use of a sequence number combined with the use of data authentication.
- Data authentication—Data authentication can refer either to integrity alone or to both integrity and authentication (data origin authentication is dependent on data integrity).
  - Data integrity—Verifies that data has not been altered.
  - Data origin authentication—Verifies that the data was actually sent by the claimed sender.
- Data confidentiality—A security service where the protected data cannot be observed.
- Data flow—A grouping of traffic, identified by a combination of source address and mask or prefix, destination address mask or prefix length, IP next protocol field, and source and destination ports, where the protocol and port fields can have any of these values. Traffic matching a specific combination of these values is logically grouped together into a data flow. A data flow can represent a single TCP connection between two hosts, or it can represent traffic between two subnets. IPsec protection is applied to data flows.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Perfect forward secrecy (PFS)—A cryptographic characteristic associated with a derived shared secret value. With PFS, if one key is compromised, previous and subsequent keys are not compromised, because subsequent keys are not derived from previous keys.
- Security Policy Database (SPD)—An ordered list of policies applied to traffic. A policy decides if a packet requires IPsec processing, if it should be allowed in clear text, or if it should be dropped.
  - The IPsec SPDs are derived from user configuration of crypto maps.
  - The IKE SPD is configured by the user.

Supported IPsec Transforms and Algorithms

The component technologies implemented for IPsec include the following transforms:

- Advanced Encrypted Standard (AES) is an encryption algorithm. It implements either 128 or 256 bits using Cipher Block Chaining (CBC) or counter mode.
- Data Encryption Standard (DES) is used to encrypt packet data and implements the mandatory 56-bit DES-CBC. CBC requires an initialization vector (IV) to start encryption. The IV is explicitly given in the IPsec packet.
- Triple DES (3DES) is a stronger form of DES with 168-bit encryption keys that allow sensitive information to be transmitted over untrusted networks.

Note Cisco NX-OS images with strong encryption are subject to United States government export controls, and have a limited distribution. Images to be installed outside the United States require an export license. Customer orders might be denied or subject to delay due to United States government regulations. Contact your sales representative or distributor for more information, or send e-mail to export@cisco.com.

- Message Digest 5 (MD5) is a hash algorithm with the HMAC variant. HMAC is a keyed hash variant used to authenticate data.
- Secure Hash Algorithm (SHA-1) is a hash algorithm with the Hash Message Authentication Code (HMAC) variant.
- AES-XCBC-MAC is a Message Authentication Code (MAC) using the AES algorithm.

Supported IKE Transforms and Algorithms

The component technologies implemented for IKE include the following transforms:

- Diffie-Hellman (DH) is a public-key cryptography protocol that allows two parties to establish a shared secret over an unsecure communications channel. Diffie-Hellman is used within IKE to establish session keys. Group 1 (768-bit), Group 2 (1024-bit), and Group 5 (1536-bit) are supported.
- Advanced Encrypted Standard (AES) is an encryption algorithm. It implements either 128 bits using Cipher Block Chaining (CBC) or counter mode.
- Data Encryption Standard (DES) is used to encrypt packet data and implements the mandatory 56-bit DES-CBC. CBC requires an initialization vector (IV) to start encryption. The IV is explicitly given in the IPsec packet.
- Triple DES (3DES) is a stronger form of DES with 168-bit encryption keys that allow sensitive information to be transmitted over untrusted networks.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Note Cisco NX-OS images with strong encryption are subject to United States government export controls, and have a limited distribution. Images to be installed outside the United States require an export license. Customer orders might be denied or subject to delay due to United States government regulations. Contact your sales representative or distributor for more information, or send e-mail to export@cisco.com.

- Message Digest 5 (MD5) is a hash algorithm with the HMAC variant. HMAC is a keyed hash variant used to authenticate data.
- Secure Hash Algorithm (SHA-1) is a hash algorithm with the Hash Message Authentication Code (HMAC) variant.
- The switch authentication algorithm uses the preshared keys based on the IP address

About IPsec Digital Certificate Support

This section describes the advantages of using certificate authorities (CAs) and digital certificates for authentication.

Implementing IPsec Without CAs and Digital Certificates

Without a CA and digital certificates, enabling IPsec services (such as encryption) between two Cisco MDS switches requires that each switch has the key of the other switch (such as an RSA public key or a shared key). You must manually specify either the RSA public keys or preshared keys on each switch in the fabric using IPsec services. Also, each new device added to the fabric will require manual configuration of the other switches in the fabric to support secure communication. Each (see Figure 7-2) switch uses the key of the other switch to authenticate the identity of the other switch; this authentication always occurs when IPsec traffic is exchanged between the two switches.

If you have multiple Cisco MDS switches in a mesh topology and want to exchange IPsec traffic passing among all of those switches, you must first configure shared keys or RSA public keys among all of those switches.

Figure 7-2 Two IPsec Switches Without CAs and Digital Certificates

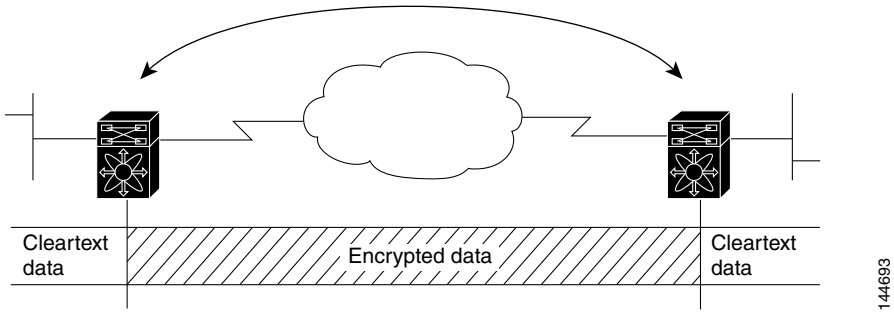

Every time a new switch is added to the IPsec network, you must configure keys between the new switch and each of the existing switches. (In Figure 7-3, four additional two-part key configurations are required to add a single encrypting switch to the network).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Consequently, the more devices that require IPsec services, the more involved the key administration becomes. This approach does not scale well for larger, more complex encrypting networks.

Figure 7-3 Four IPsec Switches Without a CA and Digital Certificates

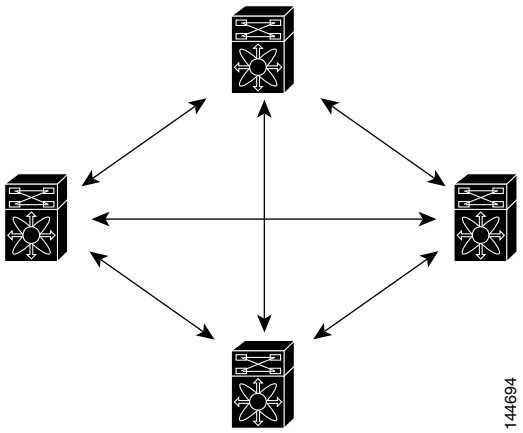

Implementing IPsec with CAs and Digital Certificates

With CA and digital certificates, you do not have to configure keys between all the encrypting switches. Instead, you individually enroll each participating switch with the CA, requesting a certificate for the switch. When this has been accomplished, each participating switch can dynamically authenticate all the other participating switches. When two devices want to communicate, they exchange certificates and digitally sign data to authenticate each other. When a new device is added to the network, you simply enroll that device with a CA, and none of the other devices needs modification. When the new device attempts an IPsec connection, certificates are automatically exchanged and the device can be authenticated.

Figure 7-4 shows the process of dynamically authenticating the devices.

Figure 7-4 Dynamically Authenticating Devices with a CA

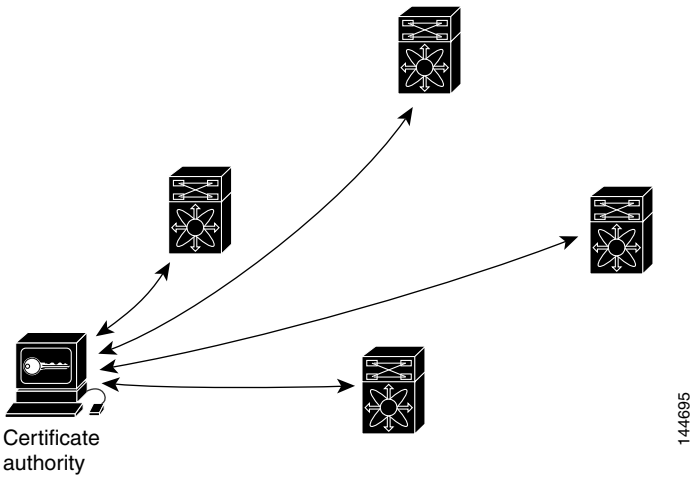

Send documentation comments to dcnm-san-docfeedback@cisco.com

To add a new IPsec switch to the network, you need only configure that new switch to request a certificate from the CA, instead of making multiple key configurations with all the other existing IPsec switches.

How CA Certificates Are Used by IPsec Devices

When two IPsec switches want to exchange IPsec-protected traffic passing between them, they must first authenticate each other—otherwise, IPsec protection cannot occur. The authentication is done with IKE.

IKE can use two methods to authenticate the switches, using preshared keys without a CA and using RSA key-pairs with a CA. Both methods require that keys must be preconfigured between the two switches.

Without a CA, a switch authenticates itself to the remote switch using either RSA-encrypted preshared keys.

With a CA, a switch authenticates itself to the remote switch by sending a certificate to the remote switch and performing some public key cryptography. Each switch must send its own unique certificate that was issued and validated by the CA. This process works because the certificate of each switch encapsulates the public key of the switch, each certificate is authenticated by the CA, and all participating switches recognize the CA as an authenticating authority. This scheme is called IKE with an RSA signature.

Your switch can continue sending its own certificate for multiple IPsec sessions, and to multiple IPsec peers until the certificate expires. When the certificate expires, the switch administrator must obtain a new one from the CA.

CAs can also revoke certificates for devices that will no longer participate in IPsec. Revoked certificates are not recognized as valid by other IPsec devices. Revoked certificates are listed in a certificate revocation list (CRL), which each peer may check before accepting a certificate from another peer.

Certificate support for IKE has the following considerations:

- The switch FQDN (host name and domain name) must be configured before installing certificates for IKE.
- Only those certificates that are configured for IKE or general usage are used by IKE.
- The first IKE or general usage certificate configured on the switch is used as the default certificate by IKE.
- The default certificate is for all IKE peers unless the peer specifies another certificate.
- If the peer asks for a certificate which is signed by a CA that it trusts, then IKE uses that certificate, if it exists on the switch, even if it is not the default certificate.
- If the default certificate is deleted, the next IKE or general usage certificate, if any exists, is used by IKE as the default certificate.
- Certificate chaining is not supported by IKE.
- IKE only sends the identity certificate, not the entire CA chain. For the certificate to be verified on the peer, the same CA chain must also exist there.

About IKE Initialization

The IKE feature must first be enabled and configured so the IPsec feature can establish data flow with the required peer. DCNM-SAN initializes IKE when you first configure it.

Send documentation comments to dcnm-san-docfeedback@cisco.com

You cannot disable IKE if IPsec is enabled. If you disable the IKE feature, the IKE configuration is cleared from the running configuration.

About the IKE Domain

You must apply the IKE configuration to an IPsec domain to allow traffic to reach the supervisor module in the local switch. DCNM-SAN sets the IPsec domain automatically when you configure IKE.

About IKE Tunnels

An IKE tunnel is a secure IKE session between two endpoints. IKE creates this tunnel to protect IKE messages used in IPsec SA negotiations.

Two versions of IKE are used in the Cisco NX-OS implementation.

- IKE version 1 (IKEv1) is implemented using RFC 2407, 2408, 2409, and 2412.
- IKE version 2 (IKEv2) is a simplified and more efficient version and does not interoperate with IKEv1. IKEv2 is implemented using the draft-ietf-ipsec-ikev2-16.txt draft.

About IKE Policy Negotiation

To protect IKE negotiations, each IKE negotiation begins with a common (shared) IKE policy. An IKE policy defines a combination of security parameters to be used during the IKE negotiation. By default, no IKE policy is configured. You must create IKE policies at each peer. This policy states which security parameters will be used to protect subsequent IKE negotiations and mandates how peers are authenticated. You can create multiple, prioritized policies at each peer to ensure that at least one policy will match a remote peer's policy.

You can configure the policy based on the encryption algorithm (DES, 3DES, or AES), the hash algorithm (SHA or MD5), and the DH group (1, 2, or 5). Each policy can contain a different combination of parameter values. A unique priority number identifies the configured policy. This number ranges from 1 (highest priority) to 255 (lowest priority). You can create multiple policies in a switch. If you need to connect to a remote peer, you must ascertain that at least one policy in the local switch contains the identical parameter values configured in the remote peer. If several policies have identical parameter configurations, the policy with the lowest number is selected.

Table 7-1 provides a list of allowed transform combinations.

Table 7-1 IKE Transform Configuration Parameters

Parameter	Accepted Values	Keyword	Default Value
encryption algorithm	56-bit DES-CBC	**des**	**3des**
	168-bit DES	**3des**	
	128-bit AES	**aes**	
hash algorithm	SHA-1 (HMAC variant)	**sha**	**sha**
	MD5 (HMAC variant)	**md5**	

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 7-1 IKE Transform Configuration Parameters (continued)

Parameter	Accepted Values	Keyword	Default Value
authentication method	Preshared keys	Not configurable	Preshared keys
DH group identifier	768-bit DH	**1**	**1**
	1024-bit DH	**2**	
	1536-bit DH	**5**	

The following table lists the supported and verified settings for IPsec and IKE encryption authentication algorithms on the Microsoft Windows and Linux platforms:

Platform	IKE	IPsec
Microsoft iSCSI initiator, Microsoft IPsec implementation on Microsoft Windows 2000 platform	3DES, SHA-1 or MD5, DH group 2	3DES, SHA-1
Cisco iSCSI initiator, Free Swan IPsec implementation on Linux platform	3DES, MD5, DH group 1	3DES, MD5

Note When you configure the hash algorithm, the corresponding HMAC version is used as the authentication algorithm.

When the IKE negotiation begins, IKE looks for an IKE policy that is the same on both peers. The peer that initiates the negotiation will send all its policies to the remote peer, and the remote peer will try to find a match. The remote peer looks for a match by comparing its own highest priority policy against the other peer's received policies. The remote peer checks each of its policies in order of its priority (highest priority first) until a match is found.

A match is found when the two peers have the same encryption, hash algorithm, authentication algorithm, and DH group values. If a match is found, IKE completes the security negotiation and the IPsec SAs are created.

If an acceptable match is not found, IKE refuses negotiation and the IPsec data flows will not be established.

Optional IKE Parameter Configuration

You can optionally configure the following parameters for the IKE feature:

- The lifetime association within each policy—The lifetime ranges from 600 to 86,400 seconds. The default is 86,400 seconds (equals one day). The lifetime association within each policy is configured when you are creating an IKE policy. See the "Configuring an IKE Policy" section on page 7-24.
- The keepalive time for each peer if you use IKEv2—The keepalive ranges from 120 to 86,400 seconds. The default is 3,600 seconds (equals one hour).

Send documentation comments to dcnm-san-docfeedback@cisco.com

- The initiator version for each peer—IKE v1 or IKE v2 (default). Your choice of initiator version does not affect interoperability when the remote device initiates the negotiation. Configure this option if the peer device supports IKEv1 and you can play the initiator role for IKE with the specified device. Use the following considerations when configuring the initiator version with FCIP tunnels:
  - If the switches on both sides of an FCIP tunnel are running MDS SAN-OS Release 3.0(1) or later, or Cisco NX-OS 4.1(1) you must configure initiator version IKEv1 on both sides of an FCIP tunnel to use only IKEv1. If one side of an FCIP tunnel is using IKEv1 and the other side is using IKEv2, the FCIP tunnel uses IKEv2.
  - If the switch on one side of an FCIP tunnel is running MDS SAN-OS Release 3.0(1) or later, or Cisco NX-OS 4.1(1b) and the switch on the other side of the FCIP tunnel is running MDS SAN-OS Release 2.x, configuring IKEv1 on either side (or both) results in the FCIP tunnel using IKEv1.

Note Only IKE v1 is supported to build IPsec between 2.x and 3.x MDS switches.

Caution You may need to configure the initiator version even when the switch does not behave as an IKE initiator under normal circumstances. Always using this option guarantees a faster recovery of traffic flows in case of failures.

Tip The keepalive time only applies to IKEv2 peers and not to all peers.

Note When IPsec implementations in the host prefer to initiate the IPsec rekey, be sure to configure the IPsec lifetime value in the Cisco MDS switch to be higher than the lifetime value in the host.

About Crypto IPv4-ACLs

IP access control lists (IPv4-ACLs) provide basic network security to all switches in the Cisco MDS 9000 Family. IPv4 IP-ACLs restrict IP-related traffic based on the configured IP filters. See Chapter 5, "Configuring IPv4 and IPv6 Access Control Lists" for details on creating and defining IPv4-ACLs.

In the context of crypto maps, IPv4-ACLs are different from regular IPv4-ACLs. Regular IPv4-ACLs determine what traffic to forward or block at an interface. For example, IPv4-ACLs can be created to protect all IP traffic between subnet A and subnet Y or Telnet traffic between host A and host B.

Crypto IPv4-ACLs are used to define which IP traffic requires crypto protection and which traffic does not.

Crypto IPv4-ACLs associated with IPsec crypto map entries have four primary functions:

- Select outbound traffic to be protected by IPsec (permit = protect).
- Indicate the data flow to be protected by the new SAs (specified by a single permit entry) when initiating negotiations for IPsec SAs.
- Process inbound traffic to filter out and discard traffic that should have been protected by IPsec.
- Determine whether or not to accept requests for IPsec SAs on behalf of the requested data flows when processing IKE negotiation from the IPsec peer.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Tip If you want some traffic to receive one type of IPsec protection (for example, encryption only) and other traffic to receive a different type of IPsec protection (for example, both authentication and encryption), create two IPv4-ACLs. Use both IPv4-ACLs in different crypto maps to specify different IPsec policies.

Note IPsec does not support IPv6-ACLs.

Mirror Image Crypto IPv4-ACLs

For every crypto IPv4-ACL specified for a crypto map entry defined at the local peer, define a mirror image crypto IPv4-ACL at the remote peer. This configuration ensures that IPsec traffic applied locally can be processed correctly at the remote peer.

Tip The crypto map entries themselves must also support common transforms and must refer to the other system as a peer.

Figure 7-5 shows some sample scenarios with and without mirror image IPv4-ACLs.

Figure 7-5 IPsec Processing of Mirror Image Configuration

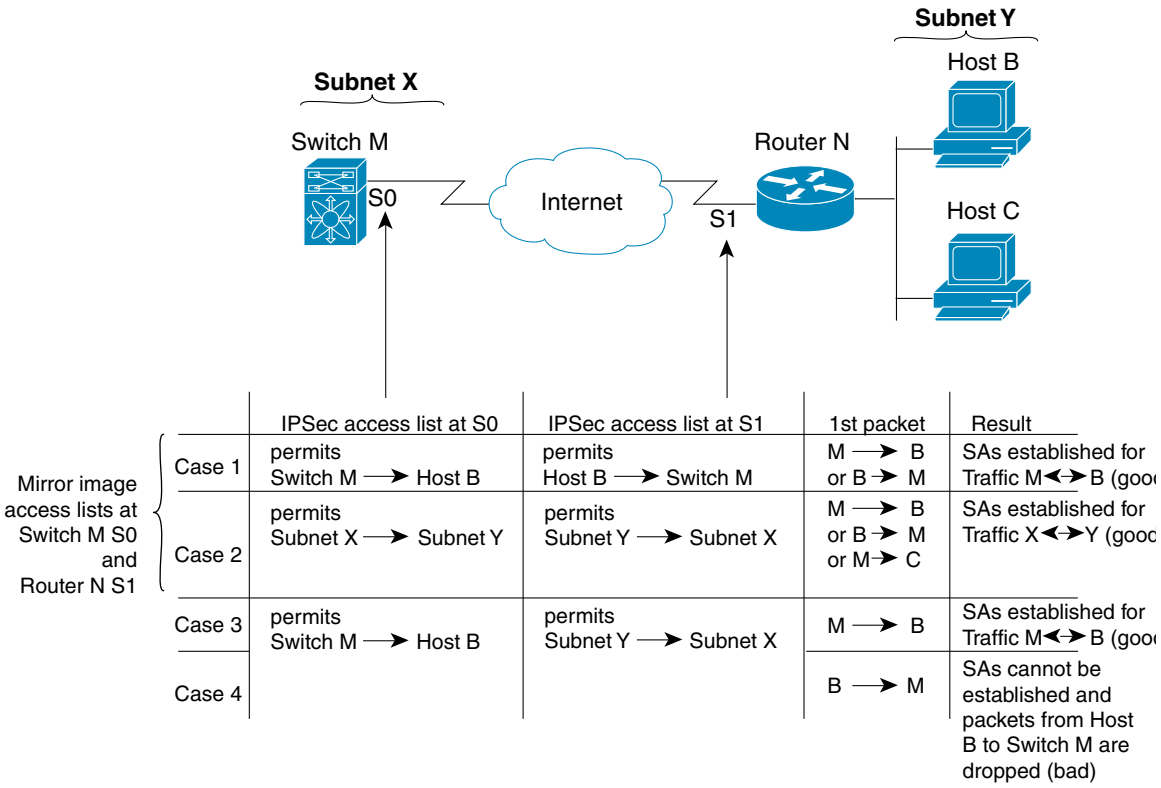

Send documentation comments to dcnm-san-docfeedback@cisco.com

As Figure 7-5 indicates, IPsec SAs can be established as expected whenever the two peers' crypto IPv4-ACLs are mirror images of each other. However, an IPsec SA can be established only some of the time when the IPv4-ACLs are not mirror images of each other. This can happen in the case when an entry in one peer's IPv4-ACL is a subset of an entry in the other peer's IPv4-ACL, such as shown in cases 3 and 4 of Figure 7-5. IPsec SA establishment is critical to IPsec. Without SAs, IPsec does not work, causing any packets matching the crypto IPv4-ACL criteria to be silently dropped instead of being forwarded with IPsec security.

In case 4, an SA cannot be established because SAs are always requested according to the crypto IPv4-ACLs at the initiating packet's end. In case 4, router N requests that all traffic between subnet X and subnet Y be protected, but this is a superset of the specific flows permitted by the crypto IPv4-ACL at switch M so the request is not permitted. Case 3 works because switch M's request is a subset of the specific flows permitted by the crypto IPv4-ACL at router N.

Because of the complexities introduced when crypto IPv4-ACLs are not configured as mirror images at peer IPsec devices, we strongly encourage you to use mirror image crypto IPv4-ACLs.

The any Keyword in Crypto IPv4-ACLs

Tip We recommend that you configure mirror image crypto IPv4-ACLs for use by IPsec and that you avoid using the **any** option.

The **any** keyword in a permit statement is discouraged when you have multicast traffic flowing through the IPsec interface. This configuration can cause multicast traffic to fail.

The **permit any** statement causes all outbound traffic to be protected (and all protected traffic sent to the peer specified in the corresponding crypto map entry) and requires protection for all inbound traffic. Then, all inbound packets that lack IPsec protection are silently dropped, including packets for routing protocols, NTP, echo, echo response, and so forth.

You need to be sure you define which packets to protect. If you must use **any** in a permit statement, you must preface that statement with a series of deny statements to filter out any traffic (that would otherwise fall within that permit statement) that you do not want to be protected.

About Transform Sets in IPsec

A transform set represents a certain combination of security protocols and algorithms. During the IPsec security association negotiation, the peers agree to use a particular transform set for protecting a particular data flow.

You can specify multiple transform sets, and then specify one or more of these transform sets in a crypto map entry. The transform set defined in the crypto map entry is used in the IPsec security association negotiation to protect the data flows specified by that crypto map entry's access list.

During IPsec security association negotiations with IKE, the peers search for a transform set that is the same at both peers. When such a transform set is found, it is selected and applied to the protected traffic as part of both peers' IPsec security associations.

Tip If you change a transform set definition, the change is only applied to crypto map entries that reference the transform set. The change is not applied to existing security associations, but used in subsequent negotiations to establish new security associations. If you want the new settings to take effect sooner, you can clear all or part of the security association database.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Note When you enable IPsec, the Cisco NX-OS software automatically creates a default transform set (ipsec_default_tranform_set) using AES-128 encryption and SHA-1 authentication algorithms.

Table 7-2 provides a list of allowed transform combinations for IPsec.

Table 7-2 IPsec Transform Configuration Parameters

Parameter	Accepted Values	Keyword
encryption algorithm	56-bit DES-CBC	**esp-des**
	168-bit DES	**esp-3des**
	128-bit AES-CBC	**esp-aes 128**
	128-bit AES-CTR[1]	**esp-aes 128 ctr**
	256-bit AES-CBC	**esp-aes 256**
	256-bit AES-CTR[1]	**esp-aes 256 ctr**
hash/authentication algorithm[1] (optional)	SHA-1 (HMAC variant)	**esp-sha1-hmac**
	MD5 (HMAC variant)	**esp-md5-hmac**
	AES-XCBC-MAC	**esp-aes-xcbc-mac**

1. If you configure the AES counter (CTR) mode, you must also configure the authentication algorithm.

The following table lists the supported and verified settings for IPsec and IKE encryption authentication algorithms on the Microsoft Windows and Linux platforms:

Platform	IKE	IPsec
Microsoft iSCSI initiator, Microsoft IPsec implementation on Microsoft Windows 2000 platform	3DES, SHA-1 or MD5, DH group 2	3DES, SHA-1
Cisco iSCSI initiator, Free Swan IPsec implementation on Linux platform	3DES, MD5, DH group 1	3DES, MD5

About Crypto Map Entries

Once you have created the crypto IPv4-ACLs and transform sets, you can create crypto map entries that combine the various parts of the IPsec SA, including the following:

- The traffic to be protected by IPsec (per the crypto IPv4-ACL). A crypto map set can contain multiple entries, each with a different IPv4-ACL.
- The granularity of the flow to be protected by a set of SAs.
- The IPsec-protected traffic destination (who the remote IPsec peer is).
- The local address to be used for the IPsec traffic (applying to an interface).
- The IPsec security to be applied to this traffic (selecting from a list of one or more transform sets).
- Other parameters to define an IPsec SA.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Crypto map entries with the same crypto map name (but different map sequence numbers) are grouped into a crypto map set.

When you apply a crypto map set to an interface, the following events occur:

- A security policy database (SPD) is created for that interface.
- All IP traffic passing through the interface is evaluated against the SPD.

If a crypto map entry sees outbound IP traffic that requires protection, an SA is negotiated with the remote peer according to the parameters included in the crypto map entry.

The policy derived from the crypto map entries is used during the negotiation of SAs. If the local switch initiates the negotiation, it will use the policy specified in the crypto map entries to create the offer to be sent to the specified IPsec peer. If the IPsec peer initiates the negotiation, the local switch checks the policy from the crypto map entries and decides whether to accept or reject the peer's request (offer).

For IPsec to succeed between two IPsec peers, both peers' crypto map entries must contain compatible configuration statements.

SA Establishment Between Peers

When two peers try to establish an SA, they must each have at least one crypto map entry that is compatible with one of the other peer's crypto map entries.

For two crypto map entries to be compatible, they must at least meet the following criteria:

- The crypto map entries must contain compatible crypto IPv4-ACLs (for example, mirror image IPv4-ACLs). If the responding peer entry is in the local crypto, the IPv4-ACL must be permitted by the peer's crypto IPv4-ACL.
- The crypto map entries must each identify the other peer or must have auto peer configured.
- If you create more than one crypto map entry for a given interface, use the seq-num of each map entry to rank the map entries: the lower the seq-num, the higher the priority. At the interface that has the crypto map set, traffic is evaluated against higher priority map entries first.
- The crypto map entries must have at least one transform set in common, where IKE negotiations are carried out and SAs are established. During the IPsec SA negotiation, the peers agree to use a particular transform set when protecting a particular data flow.

When a packet matches a permit entry in a particular IPv4-ACL, the corresponding crypto map entry is tagged, and the connections are established.

About SA Lifetime Negotiation

You can override the global lifetime values (size and time) by configuring an SA-specific lifetime value.

To specify SA lifetime negotiation values, you can optionally configure the lifetime value for a specified crypto map. If you do, this value overrides the globally set values. If you do not specify the crypto map specific lifetime, the global value (or global default) is used.

See the "Global Lifetime Values" section on page 7-18 for more information on global lifetime values.

Send documentation comments to dcnm-san-docfeedback@cisco.com

About the AutoPeer Option

Setting the peer address as **auto-peer** in the crypto map indicates that the destination endpoint of the traffic should be used as the peer address for the SA. Using the same crypto map, a unique SA can be set up at each of the endpoints in the subnet specified by the crypto map's IPv4-ACL entry. Auto-peer simplifies configuration when traffic endpoints are IPsec capable. It is particularly useful for iSCSI, where the iSCSI hosts in the same subnet do not require separate configuration.

Figure 7-6 shows a scenario where the auto-peer option can simplify configuration. Using the auto-peer option, only one crypto map entry is needed for all the hosts from subnet X to set up SAs with the switch. Each host will set up its own SA, but will share the crypto map entry. Without the auto-peer option, each host needs one crypto map entry.

Figure 7-6 iSCSI with End-to-End IPsec Using the auto-peer Option

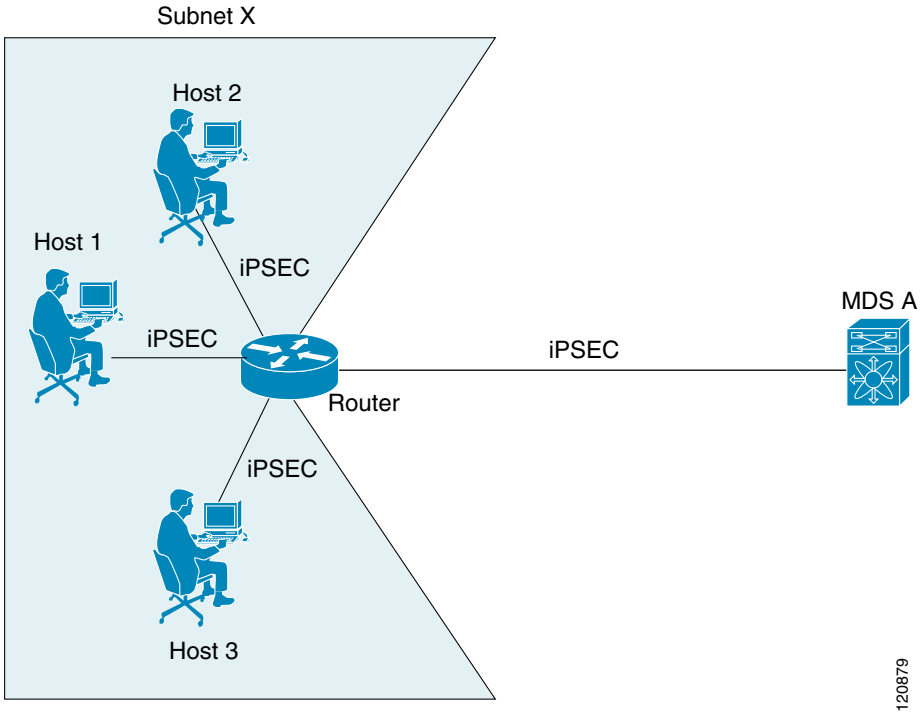

About Perfect Forward Secrecy

To specify SA lifetime negotiation values, you can also optionally configure the perfect forward secrecy (PFS) value in the crypto map.

The PFS feature is disabled by default. If you set the PFS group, you can set one of the DH groups: 1, 2, 5, or 14. If you do not specify a DH group, the software uses group 1 by default.

Send documentation comments to dcnm-san-docfeedback@cisco.com

About Crypto Map Set Interface Application

You need to apply a crypto map set to each interface through which IPsec traffic will flow. Applying the crypto map set to an interface instructs the switch to evaluate all the interface's traffic against the crypto map set and to use the specified policy during connection or SA negotiation on behalf of the traffic to be protected by crypto.

You can apply only one crypto map set to an interface. You can apply the same crypto map to multiple interfaces. However, you cannot apply more than one crypto map set to each interface.

IPsec Maintenance

Certain configuration changes will only take effect when negotiating subsequent security associations. If you want the new settings to take immediate effect, you must clear the existing security associations so that they will be reestablished with the changed configuration. If the switch is actively processing IPsec traffic, it is desirable to clear only the portion of the security association database that would be affected by the configuration changes (that is, clear only the security associations established by a given crypto map set). Clearing the full security association database should be reserved for large-scale changes, or when the router is processing very little other IPsec traffic.

Global Lifetime Values

If you have not configured a lifetime in the crypto map entry, the global lifetime values are used when negotiating new IPsec SAs.

You can configure two lifetimes: timed or traffic-volume. An SA expires after the first of these lifetimes is reached. The default lifetimes are 3,600 seconds (one hour) and 450 GB.

If you change a global lifetime, the new lifetime value will not be applied to currently existing SAs, but will be used in the negotiation of subsequently established SAs. If you wish to use the new values immediately, you can clear all or part of the SA database.

Assuming that the particular crypto map entry does not have lifetime values configured, when the switch requests new SAs it will specify its global lifetime values in the request to the peer; it will use this value as the lifetime of the new SAs. When the switch receives a negotiation request from the peer, it uses the value determined by the IKE version in use:

- If you use IKEv1 to set up IPsec SAs, the SA lifetime values are chosen to be the smaller of the two proposals. The same values are programmed on both the ends of the tunnel.
- If you use IKEv2 to set up IPsec SAs, the SAs on each end have their own set up of lifetime values and thus the SAs on both sides expire independently.

The SA (and corresponding keys) will expire according to whichever comes sooner, either after the specified amount of time (in seconds) has passed or after the specified amount of traffic (in bytes) has passed.

A new SA is negotiated before the lifetime threshold of the existing SA is reached to ensure that negotiation completes before the existing SA expires.

The new SA is negotiated when one of the following thresholds is reached (whichever comes first):

- 30 seconds before the lifetime expires or
- Approximately 10% of the lifetime in bytes remain

Send documentation comments to dcnm-san-docfeedback@cisco.com

If no traffic has passed through when the lifetime expires, a new SA is not negotiated. Instead, a new SA will be negotiated only when IPsec sees another packet that should be protected.

Prerequisites for IPsec

To use the IPsec feature, you need to perform the following tasks:

- Obtain the ENTERPRISE_PKG license (see the *Cisco MDS 9000 Family NX-OS Licensing Guide*).
- Configure IKE as described in the "About IKE Initialization" section on page 7-9.

Guidelines and Limitations

The following are the guidelines and limitations for IPsec network security:

- Crypto IPv4-ACL Guidelines, page 7-19
- Crypto Map Configuration Guidelines, page 7-20

Crypto IPv4-ACL Guidelines

Follow these guidelines when configuring IPv4-ACLs for the IPsec feature:

- The Cisco NX-OS software only allows name-based IPv4-ACLs.
- When an IPv4-ACL is applied to a crypto map, the following options apply:
  - Permit—Applies the IPsec feature to the traffic.
  - Deny—Allows clear text (default).

Note IKE traffic (UDP port 500) is implicitly transmitted in clear text.

- The IPsec feature only considers the source and destination IPv4 addresses and subnet masks, protocol, and single port number. There is no support for IPv6 in IPsec.

Note The IPsec feature does not support port number ranges and ignores higher port number field, if specified.

- The permit option causes all IP traffic that matches the specified conditions to be protected by crypto, using the policy described by the corresponding crypto map entry.
- The deny option prevents traffic from being protected by crypto. The first deny statement causes the traffic to be in clear text.
- The crypto IPv4-ACL you define is applied to an interface after you define the corresponding crypto map entry and apply the crypto map set to the interface.
- Different IPv4-ACLs must be used in different entries of the same crypto map set.
- Inbound and outbound traffic is evaluated against the same outbound IPv4-ACL. Therefore, the IPv4-ACL's criteria is applied in the forward direction to traffic exiting your switch, and the reverse direction to traffic entering your switch.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Each IPv4-ACL filter assigned to the crypto map entry is equivalent to one security policy entry. The IPsec feature supports up to 120 security policy entries for each MPS-14/2 module and Cisco MDS 9216i Switch.
- IPsec protection (see Figure 7-7) is applied to traffic between switch interface S0 (IPv4 address 10.0.0.1) and switch interface S1 (IPv4 address 20.0.0.2) as the data exits switch A's S0 interface enroute to switch interface S1. For traffic from 10.0.0.1 to 20.0.0.2, the IPv4-ACL entry on switch A is evaluated as follows:
  - source = IPv4 address 10.0.0.1
  - dest = IPv4 address 20.0.0.2

  For traffic from 20.0.0.2 to 10.0.0.1, that same IPv4-ACL entry on switch A is evaluated as follows:
  - source = IPv4 address 20.0.0.2
  - dest = IPv4 address 10.0.0.1

Figure 7-7 IPsec Processing of Crypto IPv4-ACLs

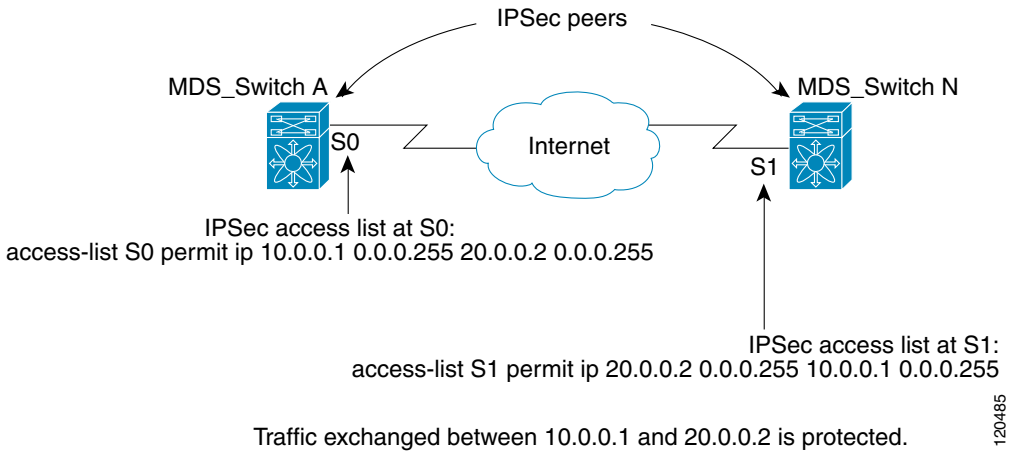

- If you configure multiple statements for a given crypto IPv4-ACL that is used for IPsec, the first permit statement that is matched is used to determine the scope of the IPsec SA. Later, if traffic matches a different permit statement of the crypto IPv4-ACL, a new, separate IPsec SA is negotiated to protect traffic matching the newly matched IPv4-ACL statement.
- Unprotected inbound traffic that matches a permit entry in the crypto IPv4-ACL for a crypto map entry flagged as IPsec is dropped, because this traffic was expected to be protected by IPsec.
- For IPsec to interoperate effectively with Microsoft iSCSI initiators, specify the TCP protocol and the local iSCSI TCP port number (default 3260) in the IPv4-ACL. This configuration ensures the speedy recovery of encrypted iSCSI sessions following disruptions such as Gigabit Ethernet interfaces shutdowns, VRRP switchovers, and port failures.

Crypto Map Configuration Guidelines

When configuring crypto map entries, follow these guidelines:

- The sequence number for each crypto map decides the order in which the policies are applied. A lower sequence number is assigned a higher priority.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Only one IPv4-ACL is allowed for each crypto map entry (the IPv4-ACL itself can have multiple permit or deny entries).
- When the tunnel endpoint is the same as the destination address, you can use the auto-peer option to dynamically configure the peer.
- For IPsec to interoperate effectively with Microsoft iSCSI initiators, specify the TCP protocol and the local iSCSI TCP port number (default 3260) in the IPv4-ACL. This configuration ensures the speedy recovery of encrypted iSCSI sessions following disruptions such as Gigabit Ethernet interfaces shutdowns, VRRP switchovers, and port failures.

Default Settings

Table 7-3 lists the default settings for IKE parameters.

Table 7-3 Default IKE Parameters

Parameters	Default
IKE	Disabled.
IKE version	IKE version 2.
IKE encryption algorithm	3DES.
IKE hash algorithm	SHA.
IKE authentication method	Not configurable (uses preshared keys).
IKE DH group identifier	Group 1.
IKE lifetime association	86,400 00 seconds (equals 24 hours).
IKE keepalive time for each peer (v2)	3,600 seconds (equals 1 hour).

Table 7-4 lists the default settings for IPsec parameters.

Table 7-4 Default IPsec Parameters

Parameters	Default
IPsec	Disabled.
Applying IPsec to the traffic.	Deny—allowing clear text.
IPsec PFS	Disabled.
IPsec global lifetime (traffic-volume)	450 Gigabytes.
IPsec global lifetime (time)	3,600 seconds (one hour).

Enabling IPsec Using FCIP Wizard

DCNM-SAN simplifies the configuration of IPsec and IKE by enabling and configuring these features as part of the FCIP configuration using the FCIP Wizard.

Detailed Steps

To enable IPsec using the FCIP Wizard, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 1 Click the FCIP Wizard icon in the toolbar.

Figure 7-8 FCIP Wizard

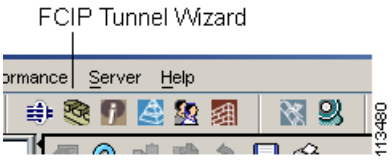

Step 2 Choose the switches that act as end points for the FCIP link and click **Next**.

Note These switches must have MPS-14/2 modules installed to configure IPsec on this FCIP link.

Step 3 Choose the Gigabit Ethernet ports on each MPS-14/2 module that will form the FCIP link.

Step 4 Check the **Enforce IPSEC Security** check box and set IKE Auth Key (see Figure 7-9).

Figure 7-9 Enabling IPsec on an FCIP Link

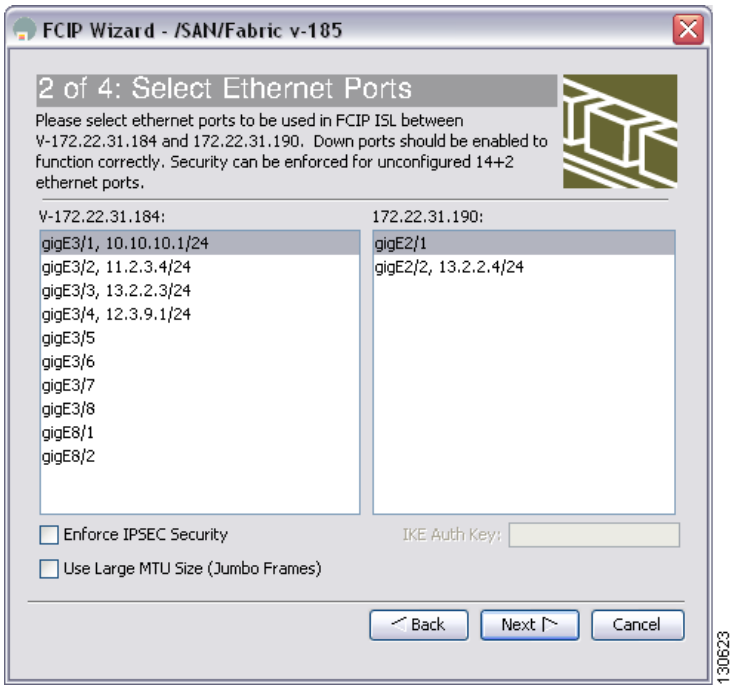

Step 5 Click **Next**. In the Specify Tunnel Properties dialog box, you see the TCP connection characteristics.

Step 6 Set the minimum and maximum bandwidth settings and round-trip time for the TCP connections on this FCIP link. Click the **Measure** button to measure the round-trip time between the Gigabit Ethernet endpoints.

Step 7 Check the **Enable Write Acceleration** check box to enable FCIP write acceleration on this FCIP link.

Step 8 Check the **Enable Optimum Compression** check box to enable IP compression on this FCIP link.

Step 9 Click **Next** to configure the FCIP tunnel parameters.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 10 Set the Port VSAN for nontrunk/auto and allowed VSAN list for the trunk tunnel. Choose a **Trunk Mode** for this FCIP link. See the *IP Services Configuration Guide*, *Cisco DCNM for SAN*.

Step 11 Click **Finish** to create this FCIP link or click **Cancel** to exit the FCIP Wizard without creating an FCIP link.

Detailed Steps

To verify that IPsec and IKE are enabled, follow these steps:

Step 1 Expand **Switches > Security** and then select **IPSEC** in the Physical Attributes pane.

Step 2 The **Control** tab is the default. Verify that the switches you want to modify for IPSec are enabled in the Status column.

Step 3 Expand **Switches > Security,** and then select **IKE** in the Physical Attributes pane.

Step 4 The **Control** tab is the default. Verify that the switches you want to modify for IKE are enabled in the Status column.

Configuring IPsec and IKE Manually

This section describes how to manually configure IPsec and IKE.

If you are not using the FCIP Wizard, see "Enabling IPsec Using FCIP Wizard" section on page 7-21.

IPsec provides secure data flows between participating peers. Multiple IPsec data flows can exist between two peers to secure different data flows, with each tunnel using a separate set of SAs.

Prerequisites

After you have completed IKE configuration, configure IPsec.

Detailed Steps

To configure IPsec in each participating IPsec peer, follow these steps:

Step 1 Identify the peers for the traffic to which secure tunnels should be established.

Step 2 Configure the transform set with the required protocols and algorithms.

Step 3 Create the crypto map and apply access control lists (IPv4-ACLs), transform sets, peers, and lifetime values as applicable.

Step 4 Apply the crypto map to the required interface.

This section includes the following topics:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Using IPsec

To use the IPsec feature, follow these steps:

Step 1 Obtain the ENTERPRISE_PKG license to enable IPsec for iSCSI and to enable IPsec for FCIP. See the *Cisco MDS 9000 Family NX-OS Licensing Guide* for more information.

Step 2 Configure IKE as described in the "Configuring IPsec and IKE Manually" section on page 7-23.

Note The IPsec feature inserts new headers in existing packets (see the *IP Services Configuration Guide, Cisco DCNM for SAN* for more information).

Configuring an IKE Policy

Detailed Steps

To configure the IKE policy negotiation parameters, follow these steps:

Step 1 Expand **Switches > Security,** and then select **IKE**.

Step 2 Click the **Policies** tab.

You see the existing IKE polices in the Information pane.

Step 3 Click **Create Row** to create an IKE policy.

Step 4 Enter the **Priority** for this switch. You can enter a value from one through 255, one being the highest.

Step 5 Select appropriate values for the encryption, hash, authentication, and DHGroup fields.

Step 6 Enter the lifetime for the policy. You can enter a lifetime from 600 to 86400 seconds.

Step 7 Click **Create** to create this policy, or click **Close** to discard any unsaved changes.

Configuring the Keepalive Time for a Peer

Detailed Steps

To configure the keepalive time for each peer, follow these steps:

Step 1 Expand **Switches > Security**, and then select **IKE**.

Step 2 Click the **Global** tab.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 3 Enter a value (in seconds) in the **KeepAliveInterval (sec)**. The keepalive interval in seconds is used by the IKE entity on the managed device with all the peers for the DOI corresponding to this conceptual row.

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

Configuring the Initiator Version

Detailed Steps

To configure the initiator version, follow these steps:

Step 1 Expand **Switches > Security,** and then select **IKE**.

Step 2 Click the **Initiator Version** tab.

You see the existing initiator versions for the peers in the Information pane.

Step 3 Click **Create Row** to create an initiator version.

Step 4 Select the Switches for the remote peer for which this IKE protocol initiator is configured.

Step 5 Enter the IP address of the remote peer.

IKEv1 represents the IKE protocol version used when connecting to a remote peer.

Step 6 Click **Create** to create this initiator version or click **Close** to discard any unsaved changes.

Clearing IKE Tunnels or Domains

If an IKE tunnel ID is not specified for the IKE configuration, you can clear all existing IKE domain connections by issuing the **clear crypto ike domain ipsec sa** command in EXEC mode.

Detailed Steps

To clear all the IKE tunnels or domains, follow these steps:

Step 1 Expand **Switches > Security,** and then select **IKE** in the Physical Attributes pane.

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

You see the IKE tunnels.

Step 3 Click the **Action** column and select **Clear** to clear the tunnel.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Refreshing SAs

Detailed Steps

To refresh the SAs after changing the IKEv2 configuration, follow these steps:

Step 1 Expand **Switches > Security,** and then select **IKE** in the Physical Attributes pane.

Step 2 Click the **Pre-Shared AuthKey** tab in the Information pane.

Step 3 Click **Refresh Values**.

Configuring Crypto

This sections includes the following topics:

Configuring Transform Sets

Detailed Steps

To configure transform sets, follow these steps:

Step 1 Expand **Switches > Security,** and then select **IPSec** in the Physical Attributes pane.

Step 2 Click the **Transform Set** tab in the Information pane.

Step 3 Click **Create Row**.

You see the Create IPSEC dialog box shown in Figure 7-10.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 7-10 Create IPSEC

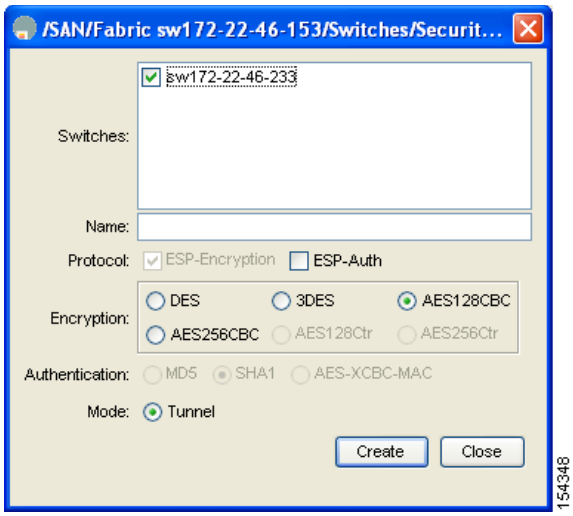

Step 4 Select the switches that you want to create a transform set for in the Create Transform Set dialog box.

Step 5 Assign a name and protocol for the transform set.

Step 6 Select the encryption and authentication algorithm. See Table 7-2 to verify the allowed transform combinations.

Step 7 Click **Create** to create the transform set or you click **Close**.

Creating Crypto Map Entries

Detailed Steps

To create mandatory crypto map entries, follow these steps:

Step 1 Expand **Switches > Security,** and then select **IPSEC** in the Physical Attributes pane.

Step 2 Click the **CryptoMap Set Entry** tab.

Step 3 (Optional) Click **Create Row** to create a crypto map entry.

Step 4 Select the switch that you want to configure or modify. If you are creating a crypto map, set the setName and priority for this crypto map.

Step 5 Select the IPv4-ACL Profile and TransformSetIdList from the drop-down list for this crypto map.

Step 6 (Optional) Check the **AutoPeer** check box or set the peer address if you are creating a crypto map. See the "About the AutoPeer Option" section on page 7-17.

Step 7 Choose the appropriate PFS selection. See the "About Perfect Forward Secrecy" section on page 7-17.

Step 8 Supply the Lifetime and LifeSize. See the "About SA Lifetime Negotiation" section on page 7-16.

Step 9 Click **Create** if you are creating a crypto map, or click **Apply Changes** if you are modifying an existing crypto map.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Setting the SA Lifetime

Detailed Steps

To set the SA lifetime for a specified crypto map entry, follow these steps:

Step 1 Expand **Switches > Security** and then select **IPSEC** in the Physical Attributes pane.

Step 2 Click the **CryptoMap Set Entry** tab.

Step 3 Scroll to the right half of the dialog box.

Step 4 Double-click and modify the value in the **Life Time(sec)** column.

Step 5 Click **Apply Changes** to save your changes.

Configuring Perfect Forward Secrecy

Detailed Steps

To configure the PFS value, follow these steps:

Step 1 Expand **Switches > Security** and then select **IPSEC** in the Physical Attributes pane.

Step 2 Click the **CryptoMap Set Entry** tab.

Step 3 From the drop-down list in the PFS column select the appropriate value.

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

Applying a Crypto Map Set

Detailed Steps

To apply a crypto map set to an interface, follow these steps:

Step 1 Expand **Switches > Security,** and then select **IPSEC** in the Physical Attributes pane.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 2 Click the **Interfaces** tab.

Step 3 Select the switch and interface you want to configure.

Step 4 Enter the name of the crypto map that you want to apply to this interface in the CryptomapSetName field.

Step 5 Click **Create** to apply the crypto map to the selected interface or click **Close** to exit the dialog box without applying the crypto map.

Configuring Global Lifetime Values

Detailed Steps

To configure global SA lifetimes, follow these steps:

Step 1 Choose **Switches > Security,** and then select **IPSEC** in the Physical Attributes pane.

Step 2 You see the IPsec configuration in the Information pane.

Step 3 Click the **Global** tab.

Step 4 Double-click and edit the value in the **Life Time(sec)** column.

Step 5 Click **Apply Changes** to save your changes.

Field Descriptions for IPsec

The following are the field descriptions for IPsec.

IPsec

Field	Description
Interface, CryptomapName	The binding of cryptomap sets to the interfaces of the managed entity.

IKE Global

Field	Description
RemIdentity	Displays the keepalive interval in seconds used by the IKE entity on the managed device with all the peers for the DOI corresponding to this conceptual row.
Key	Displays the type of keepalives to be used by the IKE entity on the managed device with all the peers for the DOI corresponding to this conceptual row.

Send documentation comments to dcnm-san-docfeedback@cisco.com

IKE Pre-Shared AuthKey

Field	Description
KeepAliveInterval (sec)	The Phase 1 ID identity of the peer for which this pre-shared key is configured on the local entity.
IdentityType	The pre-shared authorization key used in authenticating the peer corresponding to this conceptual row.

IKE Policies

Field	Description
Priority	The priority of this ISAKMP policy entry. The policy with lower value would take precedence over the policy with higher value in the same DOI.
Encr	The encryption transform specified by this ISAKMP policy specification. The Internet Key Exchange (IKE) tunnels setup using this policy item would use the specified encryption transform to protect the ISAKMP PDUs.
Hash	The hash transform specified by this ISAKMP policy specification. The IKE tunnels setup using this policy item would use the specified hash transform to protect the ISAKMP PDUs.
Auth	The peer authentication method specified by this ISAKMP policy specification. If this policy entity is selected for negotiation with a peer, the local entity would authenticate the peer using the method specified by this object.
DHGroup	Specifies the Oakley group used for Diffie Hellman exchange in the Main Mode. If this policy item is selected to negotiate Main Mode with an IKE peer, the local entity chooses the group specified by this object to perform Diffie Hellman exchange with the peer.
Lifetime (sec)	Specifies the lifetime in seconds of the IKE tunnels generated using this policy specification.

IKE Initiator Version

Field	Description
Address	The address of the remote peer corresponding to this conceptual row. This object cannot be modified while the corresponding value of cicIkeCfgInitiatorStatus is equal to active.
Version	The IKE protocol version used when connecting to a remote peer specified in cicIkeCfgInitiatorPAddr. This object cannot be modified while the corresponding value of cicIkeCfgInitiatorStatus is equal to active.

Send documentation comments to dcnm-san-docfeedback@cisco.com

IKE Tunnels

Field	Description
LocalAddress	The address of the local endpoint for the Phase-1 tunnel.
RemoteAddresss	The address of the remote endpoint of the Phase-1 tunnel.
AuthMethod	The authentication method used in Phase-1 negotiations on the control tunnel corresponding to this conceptual row.
Action	The action to be taken on this tunnel. If clear, then this tunnel is cleared. If re-key, then re-keying is forced on this tunnel. The value none would be returned on doing read of this object.

IPSEC Global

Field	Description
Lifetime (sec)	The default lifetime (in seconds) assigned to an IPsec tunnel as a global policy (maybe overridden in specific cryptomap definitions).
Lifesize (KB)	The default life size in KB assigned to an IPsec tunnel as a global policy (unless overridden in cryptomap definition).

IPSEC Transform Set

Field	Description
Id	This is the sequence number of the transform set that uniquely identifies the transform set. Distinct transform sets must have distinct sequence numbers.
Protocol	Represents the suite of Phase-2 security protocols of this transform set.
ESP Encryption	Represents the transform used for ESP encryption.
ESP Authentication	Represents the transform used to implement integrity check with ESP protocol.
Mode	Represents the encapsulation mode of the transform set.

Send documentation comments to dcnm-san-docfeedback@cisco.com

IPSEC CryptoMap Set Entry

Field	Description
IpFilter	Specifies an IP protocol filter to be secured using this cryptomap entry. When it has a value of zero-length string, it is not valid/applicable.
TransformSetIdList	The list of cipsXformSetId that are members of this CipsStaticCryptomapEntry. The value of this object is a concatenation of zero ormore 4-octet strings, where each 4-octet string contains a 32-bit cipsXformSetId value in network byte order. A zero length string value means this list has no members.
AutoPeer	If true the destination address is taken as the peer address, while creating the tunnel.
Peer Address	The IP address of the peer to which this cryptomap entry is currently connected.
PFS	Identifies whether the tunnels instantiated due to this policy item should use Perfect Forward Secrecy (PFS) and if so, what group of Oakley they should use.
LifeTime	Specifies the lifetime of the IPsec Security Associations (SA) created using this IPsec policy entry.
Lifesize Value	Identifies the life size (maximum traffic in bytes that may be carried) of the IPSec SAs created using this IPSec policy entry. When a Security Association (SA) is created using this IPsec policy entry, its life size takes the value of this object.

IPSEC Interfaces

Field	Description
CryptomapName	The index of the static cryptomap table. The value of the string is the name string assigned by the NMS when defining a cryptomap set.
InterfaceList	Interfaces belong to the cryptomap.

IPSEC Tunnels

Field	Description
Local Address	The IP address of the local endpoint for the IPsec Phase-2 tunnel.
RemoteAdddress	The type of the IP address of the remote endpoint for the IPsec Phase-2 tunnel.
ESP Encryption	The encryption algorithm used by the outbound security association of the IPsec Phase-2 tunnel.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
ESP Encryption KeySize	The key size in bits of the negotiated key to be used with the algorithm denoted by ceipSecTunOutSaEncryptAlgo. For DES and 3DES the key size is respectively 56 and 168. For AES, this will denote the negotiated key size.
ESP Authentication	The authentication algorithm used by the inbound encapsulation security protocol (ESP) security association of the IPsec Phase-2 tunnel.
LifeSize (KB)	The negotiated life size of the IPSEC Phase-2 tunnel in kilobytes.
LifeTime (sec)	The negotiated lifetime of the IPSEC Phase-2 tunnel in seconds. If the tunnel was setup manually, the value of this MIB element should be 0.
Action	The status of the MIB table row.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

CHAPTER 8

Configuring FC-SP and DHCHAP

This chapter includes the following topics:

Information About Fabric Authentication

Fibre Channel Security Protocol (FC-SP) capabilities provide switch-switch and host-switch authentication to overcome security challenges for enterprise-wide fabrics. Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) is an FC-SP protocol that provides authentication between Cisco MDS 9000 Family switches and other devices. DHCHAP consists of the CHAP protocol combined with the Diffie-Hellman exchange.

To authenticate through VFC ports, FC-SP peers use the port VSAN for communication. Hence, the port VSAN needs to be the same and active on both the peers to send and receive authentication messages.

All switches in the Cisco MDS 9000 Family enable fabric-wide authentication from one switch to another switch, or from a switch to a host. These switch and host authentications are performed locally or remotely in each fabric. As storage islands are consolidated and migrated to enterprise-wide fabrics new security challenges arise. The approach of securing storage islands cannot always be guaranteed in enterprise-wide fabrics.

For example, in a campus environment with geographically distributed switches someone could maliciously interconnect incompatible switches or you could accidentally do so, resulting in Inter-Switch Link (ISL) isolation and link disruption. This need for physical security is addressed by switches in the Cisco MDS 9000 Family (see Figure 8-1).

Send documentation comments to dcnm-san-docfeedback@cisco.com

Figure 8-1 Switch and Host Authentication

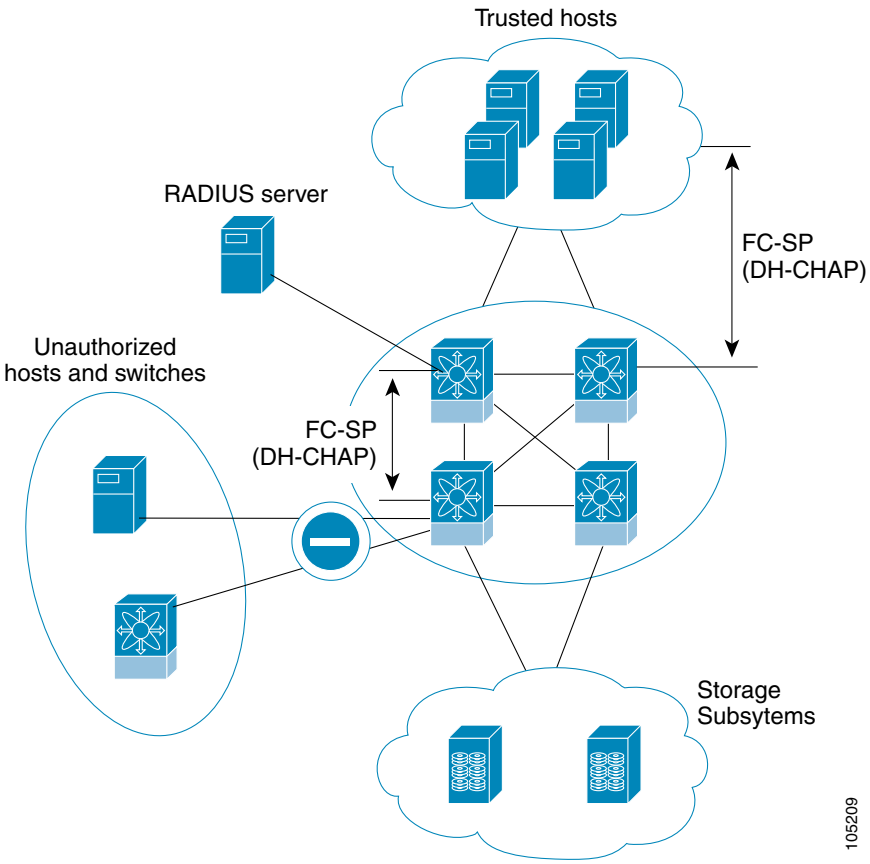

Note Fibre Channel (FC) host bus adapters (HBAs) with appropriate firmware and drivers are required for host-switch authentication.

DHCHAP

DHCHAP is an authentication protocol that authenticates the devices connecting to a switch. Fibre Channel authentication allows only trusted devices to be added to a fabric, which prevents unauthorized devices from accessing the switch.

Note The terms FC-SP and DHCHAP are used interchangeably in this chapter.

DHCHAP is a mandatory password-based, key-exchange authentication protocol that supports both switch-to-switch and host-to-switch authentication. DHCHAP negotiates hash algorithms and DH groups before performing authentication. It supports MD5 and SHA-1 algorithm-based authentication.

Configuring the DHCHAP feature requires the ENTERPRISE_PKG license (see the *Cisco MDS 9000 Family NX-OS Licensing Guide*).

Send documentation comments to dcnm-san-docfeedback@cisco.com

DHCHAP Compatibility with Existing Cisco MDS Features

This section identifies the impact of configuring the DHCHAP feature along with existing Cisco MDS features:

- PortChannel interfaces—If DHCHAP is enabled for ports belonging to a PortChannel, DHCHAP authentication is performed at the physical interface level, not at the PortChannel level.
- FCIP interfaces—The DHCHAP protocol works with the FCIP interface just as it would with a physical interface.
- Port security or fabric binding—Fabric binding policies are enforced based on identities authenticated by DHCHAP.
- VSANs—DHCHAP authentication is not done on a per-VSAN basis.
- High availability—DHCHAP authentication works transparently with existing HA features.

About Enabling DHCHAP

By default, the DHCHAP feature is disabled in all switches in the Cisco MDS 9000 Family.

You must explicitly enable the DHCHAP feature to access the configuration and verification commands for fabric authentication. When you disable this feature, all related configurations are automatically discarded.

About DHCHAP Authentication Modes

The DHCHAP authentication status for each interface depends on the configured DHCHAP port mode.

When the DHCHAP feature is enabled in a switch, each Fibre Channel interface or FCIP interface may be configured to be in one of four DHCHAP port modes:

- On—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software moves the link to an isolated state.
- Auto-Active—During switch initialization, if the connecting device supports DHCHAP authentication, the software performs the authentication sequence. If the connecting device does not support DHCHAP authentication, the software continues with the rest of the initialization sequence.
- Auto-Passive (default)—The switch does not initiate DHCHAP authentication, but participates in DHCHAP authentication if the connecting device initiates DHCHAP authentication.
- Off—The switch does not support DHCHAP authentication. Authentication messages sent to such ports return error messages to the initiating switch.

Note Whenever DHCHAP port mode is changed to a mode other than the Off mode, reauthentication is performed.

Table 8-1 identifies the switch-to-switch authentication behavior between two Cisco MDS switches in various modes.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 8-1 DHCHAP Authentication Status Between Two MDS Switches

Switch N DHCHAP Modes	Switch 1 DHCHAP Modes			
	on	**auto-active**	**auto-passive**	**off**
on	FC-SP authentication is performed.	FC-SP authentication is performed.	FC-SP authentication is performed.	Link is brought down.
auto-Active				FC-SP authentication is *not* performed.
auto-Passive			FC-SP authentication is *not* performed.	
off	Link is brought down.	FC-SP authentication is *not* performed.		

About the DHCHAP Hash Algorithm

Cisco MDS switches support a default hash algorithm priority list of MD5 followed by SHA-1 for DHCHAP authentication.

Tip If you change the hash algorithm configuration, then change it globally for all switches in the fabric.

Caution RADIUS and TACACS+ protocols always use MD5 for CHAP authentication. Using SHA-1 as the hash algorithm may prevent RADIUS and TACACS+ usage—even if these AAA protocols are enabled for DHCHAP authentication.

About the DHCHAP Group Settings

All switches in the Cisco MDS Family support all DHCHAP groups specified in the standard: 0 (null DH group, which does not perform the Diffie-Hellman exchange), 1, 2, 3, or 4.

Tip If you change the DH group configuration, change it globally for all switches in the fabric.

About the DHCHAP Password

DHCHAP authentication in each direction requires a shared secret password between the connected devices. To do this, you can use one of three approaches to manage passwords for all switches in the fabric that participate in DHCHAP.

- Approach 1—Use the same password for all switches in the fabric. This is the simplest approach. When you add a new switch, you use the same password to authenticate that switch in this fabric. It is also the most vulnerable approach if someone from the outside maliciously attempts to access any one switch in the fabric.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- Approach 2—Use a different password for each switch and maintain that password list in each switch in the fabric. When you add a new switch, you create a new password list and update all switches with the new list. Accessing one switch yields the password list for all switches in that fabric.
- Approach 3—Use different passwords for different switches in the fabric. When you add a new switch, multiple new passwords corresponding to each switch in the fabric must be generated and configured in each switch. Even if one switch is compromised, the password of other switches are still protected. This approach requires considerable password maintenance by the user.

Note All passwords are restricted to 64 alphanumeric characters and can be changed, but not deleted.

Tip We recommend using RADIUS or TACACS+ for fabrics with more than five switches. If you need to use a local password database, you can continue to do so using Approach 3 and using the Cisco MDS 9000 Family DCNM-SAN to manage the password database.

About Password Configuration for Remote Devices

You can configure passwords in the local authentication database for other devices in a fabric. The other devices are identified by their device name, which is also known as the switch WWN or device WWN. The password is restricted to 64 characters and can be specified in clear text (0) or in encrypted text (7).

Note The switch WWN identifies the physical switch. This WWN is used to authenticate the switch and is different from the VSAN node WWN.

About the DHCHAP Timeout Value

During the DHCHAP protocol exchange, if the MDS switch does not receive the expected DHCHAP message within a specified time interval, authentication failure is assumed. The time ranges from 20 (no authentication is performed) to 1000 seconds. The default is 30 seconds.

When changing the timeout value, consider the following factors:

- The existing RADIUS and TACACS+ timeout values.

The same value must also be configured on all switches in the fabric.

Enabling FC-SP on ISLs

There is an ISL pop-up menu in DCNM-SAN called Enable FC-SP that enables FC-SP on switches at either end of the ISL. You are prompted for an FC-SP generic password, then asked to set FC-SP interface mode to ON for affected ports. Right-click an ISL and click **Enable FC-SP** to access this feature.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Default Settings

Table 8-2 lists the default settings for all fabric security features in any switch.

Table 8-2 Default Fabric Security Settings

Parameters	Default
DHCHAP feature	Disabled
DHCHAP hash algorithm	A priority list of MD5 followed by SHA-1 for DHCHAP authentication
DHCHAP authentication mode	Auto-passive
DHCHAP group default priority exchange order	0, 4, 1, 2, and 3 respectively
DHCHAP timeout value	30 seconds

Configuring DHCAP

To configure DHCHAP authentication using the local password database, follow these steps:

Step 1 Enable DHCHAP.

Step 2 Identify and configure the DHCHAP authentication modes.

Step 3 Configure the hash algorithm and DH group.

Step 4 Configure the DHCHAP password for the local switch and other switches in the fabric.

Step 5 Configure the DHCHAP timeout value for reauthentication.

Step 6 Verify the DHCHAP configuration.

Enabling DHCHAP

To enable DHCHAP for a Cisco MDS switch , follow these steps:

Step 1 Expand **Switches**, expand **Security** and then select **FC-SP**.

The **Control** tab is the default. You see the FC-SP enable state for all switches in the fabric.

Step 2 Set the Command drop-down menu to enable for all switches that you want to enable FC-SP on.

Step 3 Click the **Apply Changes** icon to enable FC-SP and DHCHAP on the selected switches.

Configuring the DHCHAP Mode

To configure the DHCHAP mode for a particular interface, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 1 Expand **Switches**, expand **Interfaces,** and then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **FC-SP** tab.

Step 3 Set the **Mode** drop-down menu to the DHCHAP authentication mode you want to configure for that interface.

Step 4 Click the **Apply Changes** icon to save these DHCHAP port mode settings.

Configuring the DHCHAP Hash Algorithm

To configure the hash algorithm, follow these steps:

Step 1 Choose **Switches > Security,** and then select **FC-SP**.

Step 2 Click the **General/Password** tab.

You see the DHCHAP general settings mode for each switch.

Step 3 Change the DHCHAP HashList for each switch in the fabric.

Step 4 Click the **Apply Changes** icon to save the updated hash algorithm priority list.

Configuring the DHCHAP Group Settings

To change the DH group settings, follow these steps:

Step 1 Expand **Switches > Security,** and then select **FC-SP**.

Step 2 Click the **General/Password** tab.

Step 3 Change the DHCHAP GroupList for each switch in the fabric.

Step 4 Click the **Apply Changes** icon to save the updated hash algorithm priority list.

Configuring DHCHAP Passwords for the Local Switch

To configure the DHCHAP password for the local switch, follow these steps:

Step 1 Expand **Switches > Security,** and then select **FC-SP**.

You see the FC-SP configuration in the Information pane.

Step 2 Click the **Local Passwords** tab.

Step 3 Click the **Create Row** icon to create a new local password.

Send documentation comments to dcnm-san-docfeedback@cisco.com

You see the Create Local Passwords dialog box.

Step 4 (Optional) Check the switches that you want to configure the same local password on.

Step 5 Select the switch WNN and fill in the Password field.

Step 6 Click **Create** to save the updated password.

Configuring DHCHAP Passwords for Remote Devices

To locally configure the remote DHCHAP password for another switch in the fabric, follow these steps:

Step 1 Right-click an ISL and select **Enable FC-SP** from the drop-down list (see Figure 8-2).

Figure 8-2 Enable FC-SP

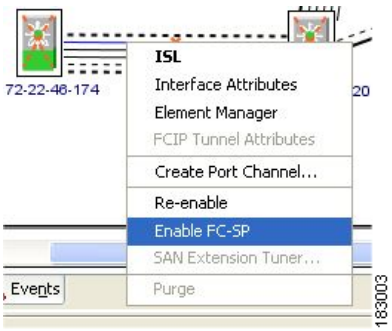

You see the Enable FC-SP dialog box.

Figure 8-3 Enable FC-SP Dialog Box

Step 2 Click **Apply** to save the updated password.

Configuring the DHCHAP Timeout Value

To configure the DHCHAP timeout value, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 1 Expand **Switches > Security,** and then select **FC-SP**.

You see the FC-SP configuration in the Information pane.

Step 2 Click the **General/Password** tab.

You see the DHCHAP general settings mode for each switch.

Step 3 Change the DHCHAP timeout value for each switch in the fabric.

Step 4 Click the **Apply Changes** icon to save the updated information.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

CHAPTER 9

Configuring Port Security

All switches in the Cisco MDS 9000 Family provide port security features that reject intrusion attempts and report these intrusions to the administrator.

This chapter includes the following topics:

Information About Port Security

All switches in the Cisco MDS 9000 Family provide port security features that reject intrusion attempts and report these intrusions to the administrator.

Typically, any Fibre Channel device in a SAN can attach to any SAN switch port and access SAN services based on zone membership. Port security features prevent unauthorized access to a switch port in the Cisco MDS 9000 Family in the following ways:

- Login requests from unauthorized Fibre Channel devices (Nx ports) and switches (xE ports) are rejected.
- All intrusion attempts are reported to the SAN administrator through system messages.
- Configuration distribution uses the CFS infrastructure, and is limited to those switches that are CFS capable. Distribution is disabled by default.
- Configuring the port security policy requires the ENTERPRISE_PKG license (see the Cisco MDS 9000 Family NX-OS Licensing Guide).

This section includes the following topics:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Port Security Enforcement

To enforce port security, configure the devices and switch port interfaces through which each device or switch is connected, and activate the configuration.

- Use the port world wide name (pWWN) or the node world wide name (nWWN) to specify the Nx port connection for each device.
- Use the switch world wide name (sWWN) to specify the xE port connection for each switch.

Each Nx and xE port can be configured to restrict a single port or a range of ports.

Enforcement of port security policies are done on every activation and when the port tries to come up.

The port security feature uses two databases to accept and implement configuration changes.

- Configuration database—All configuration changes are stored in the configuration database.
- Active database—The database currently enforced by the fabric. The port security feature requires all devices connecting to a switch to be part of the port security active database. The software uses this active database to enforce authorization.

About Auto-Learning

You can instruct the switch to automatically learn (auto-learn) the port security configurations over a specified period. This feature allows any switch in the Cisco MDS 9000 Family to automatically learn about devices and switches that connect to it. Use this feature when you activate the port security feature for the first time as it saves tedious manual configuration for each port. You must configure auto-learning on a per-VSAN basis. If enabled, devices and switches that are allowed to connect to the switch are automatically learned, even if you have not configured any port access.

When auto-learning is enabled, learning happens only for the devices or interfaces that were not already logged into the switch. Learned entries on a port are cleaned up after you shut down that port if auto-learning is still enabled.

Learning does not override the existing configured port security policies. So, for example, if an interface is configured to allow a specific pWWN, then auto-learning will not add a new entry to allow any other pWWN on that interface. All other pWWNs will be blocked even in auto-learning mode.

No entries are learned for a port in the shutdown state.

Send documentation comments to dcnm-san-docfeedback@cisco.com

When you activate the port security feature, auto-learning is also automatically enabled.

Note If you enable auto-learning before activating port security, you cannot activate until auto-learning is disabled.

Port Security Activation

By default, the port security feature is not activated in any switch in the Cisco MDS 9000 Family.

By activating the port security feature, the following apply:

- Auto-learning is also automatically enabled, which means:
  - From this point, auto-learning happens only for the devices or interfaces that were not logged into the switch.
  - You cannot activate the database until you disable auto-learning.
- All the devices that are already logged in are learned and are added to the active database.
- All entries in the configured database are copied to the active database.

After the database is activated, subsequent device login is subject to the activated port bound WWN pairs, excluding the auto-learned entries. You must disable auto-learning before the auto-learned entries become activated.

When you activate the port security feature, auto-learning is also automatically enabled. You can choose to activate the port security feature and disable auto-learning.

Tip If a port is shut down because of a denied login attempt, and you subsequently configure the database to allow that login, the port does not come up automatically. You must explicitly issue a **no shutdown** CLI command to bring that port back online.

Database Activation Rejection

Database activation is rejected in the following cases:

- Missing or conflicting entries exist in the configuration database but not in the active database.
- The auto-learning feature was enabled before the activation. To reactivate a database in this state, disable auto-learning.
- The exact security is not configured for each PortChannel member.
- The configured database is empty but the active database is not.

If the database activation is rejected due to one or more conflicts listed in the previous section, you may decide to proceed by forcing the port security activation.

About Enabling Auto-learning

The state of the auto-learning configuration depends on the state of the port security feature:

- If the port security feature is not activated, auto-learning is disabled by default.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- If the port security feature is activated, auto-learning is enabled by default (unless you explicitly disabled this option).

Tip If auto-learning is enabled on a VSAN, you can only activate the database for that VSAN by using the **force** option.

Auto-learning Device Authorization

Table 9-1 summarizes the authorized connection conditions for device requests.

Table 9-1 Authorized Auto-learning Device Requests

Condition	Device (pWWN, nWWN, sWWN)	Requests Connection to	Authorization
1	Configured with one or more switch ports	A configured switch port	Permitted
2		Any other switch port	Denied
3	Not configured	A switch port that is not configured	Permitted if auto-learning enabled
4			Denied if auto-learning disabled
5	Configured or not configured	A switch port that allows any device	Permitted
6	Configured to log in to any switch port	Any port on the switch	Permitted
7	Not configured	A port configured with some other device	Denied

Authorization Scenarios

Assume that the port security feature is activated and the following conditions are specified in the active database:

- A pWWN (P1) is allowed access through interface fc1/1 (F1).
- A pWWN (P2) is allowed access through interface fc1/1 (F1).
- A nWWN (N1) is allowed access through interface fc1/2 (F2).
- Any WWN is allowed access through interface fc1/3 (F3).
- A nWWN (N3) is allowed access through any interface.
- A pWWN (P3) is allowed access through interface fc1/4 (F4).
- A sWWN (S1) is allowed access through interface fc1/10-13 (F10 to F13).
- A pWWN (P10) is allowed access through interface fc1/11 (F11).

Table 9-2 summarizes the port security authorization results for this active database. The conditions listed refer to the conditions from Table 9-1.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 9-2 Authorization Results for Scenario

Device Connection Request	Authorization	Condition	Reason
P1, N2, F1	Permitted	1	No conflict.
P2, N2, F1	Permitted	1	No conflict.
P3, N2, F1	Denied	2	F1 is bound to P1/P2.
P1, N3, F1	Permitted	6	Wildcard match for N3.
P1, N1, F3	Permitted	5	Wildcard match for F3.
P1, N4, F5	Denied	2	P1 is bound to F1.
P5, N1, F5	Denied	2	N1 is only allowed on F2.
P3, N3, F4	Permitted	1	No conflict.
S1, F10	Permitted	1	No conflict.
S2, F11	Denied	7	P10 is bound to F11.
P4, N4, F5 (auto-learning on)	Permitted	3	No conflict.
P4, N4, F5(auto-learning off)	Denied	4	No match.
S3, F5 (auto-learning on)	Permitted	3	No conflict.
S3, F5 (auto-learning off)	Denied	4	No match.
P1, N1, F6 (auto-learning on)	Denied	2	P1 is bound to F1.
P5, N5, F1 (auto-learning on)	Denied	7	Only P1 and P2 bound to F1.
S3, F4 (auto-learning on)	Denied	7	P3 paired with F4.
S1, F3 (auto-learning on)	Permitted	5	No conflict.
P5, N3, F3	Permitted	6	Wildcard (*) match for F3 and N3.
P7, N3, F9	Permitted	6	Wildcard (*) match for N3.

About WWN Identification

If you decide to manually configure port security, be sure to adhere to the following guidelines:

- Identify switch ports by the interface or by the fWWN.
- Identify devices by the pWWN or by the nWWN.
- If an Nx port is allowed to log in to SAN switch port Fx, then that Nx port can only log in through the specified Fx port.
- If an Nx port's nWWN is bound to an Fx port WWN, then all pWWNs in the Nx port are implicitly paired with the Fx port.
- TE port checking is done on each VSAN in the allowed VSAN list of the trunk port.
- All PortChannel xE ports must be configured with the same set of WWNs in the same PortChannel.
- E port security is implemented in the port VSAN of the E port. In this case the sWWN is used to secure authorization checks.
- Once activated, the config database can be modified without any effect on the active database.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- By saving the running configuration, you save the configuration database and activated entries in the active database. Learned entries in the active database are not saved.

Activation and Auto-learning Configuration Distribution

Activation and auto-learning configurations in distributed mode are remembered as actions to be performed when you commit the changes in the pending database.

Learned entries are temporary and do not have any role in determining if a login is authorized or not. As such, learned entries do not participate in distribution. When you disable learning and commit the changes in the pending database, the learned entries become static entries in the active database and are distributed to all switches in the fabric. After the commit, the active database on all switches are identical and learning can be disabled.

If the pending database contains more than one activation and auto-learning configuration when you commit the changes, then the activation and auto-learning changes are consolidated and the behavior may change (see Table 9-3).

Table 9-3 Scenarios for Activation and Auto- learning Configurations in Distributed Mode

Scenario	Actions	Distribution = OFF	Distribution = ON
A and B exist in the configuration database, activation is not done and devices C,D are logged in.	1. You activate the port security database and enable auto-learning.	configuration database = {A,B} active database = {A,B, C[1], D*}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled}
	2. A new entry E is added to the configuration database.	configuration database = {A,B, E} active database = {A,B, C*, D*}	configuration database = {A,B} active database = {null} pending database = {A,B, E + activation to be enabled}
	3. You issue a commit.	Not applicable	configuration database = {A,B, E} active database = {A,B, E, C*, D*} pending database = empty

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 9-3 Scenarios for Activation and Auto- learning Configurations in Distributed Mode (continued)

Scenario	Actions	Distribution = OFF	Distribution = ON
A and B exist in the configuration database, activation is not done and devices C,D are logged in.	1. You activate the port security database and enable auto-learning.	configuration database = {A,B} active database = {A,B, C*, D*}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled}
	2. You disable learning.	configuration database = {A,B} active database = {A,B, C, D}	configuration database = {A,B} active database = {null} pending database = {A,B + activation to be enabled + learning to be disabled}
	3. You issue a commit.	Not applicable	configuration database = {A,B} active database = {A,B} and devices C and D are logged out. This is equal to an activation with auto-learning disabled. pending database = empty

1. The * (asterisk) indicates learned entries.

Tip In this case, we recommend that you perform a commit at the end of each operation: after you activate port security and after you enable auto-learning.

Database Interaction

Table 9-4 lists the differences and interaction between the active and configuration databases.

Table 9-4 Active and Configuration Port Security Databases

Active Database	Configuration Database
Read-only.	Read-write.
Saving the configuration only saves the activated entries. Learned entries are not saved.	Saving the configuration saves all the entries in the configuration database.
Once activated, all devices that have already logged into the VSAN are also learned and added to the active database.	Once activated, the configuration database can be modified without any effect on the active database.
You can overwrite the active database with the configured database by activating the port security database. Forcing an activation may violate the entries already configured in the active database.	You can overwrite the configuration database with the active database.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Database Scenarios

Figure 9-1 depicts various scenarios to depict the active database and the configuration database status based on port security configurations.

Figure 9-1 Port Security Database Scenarios

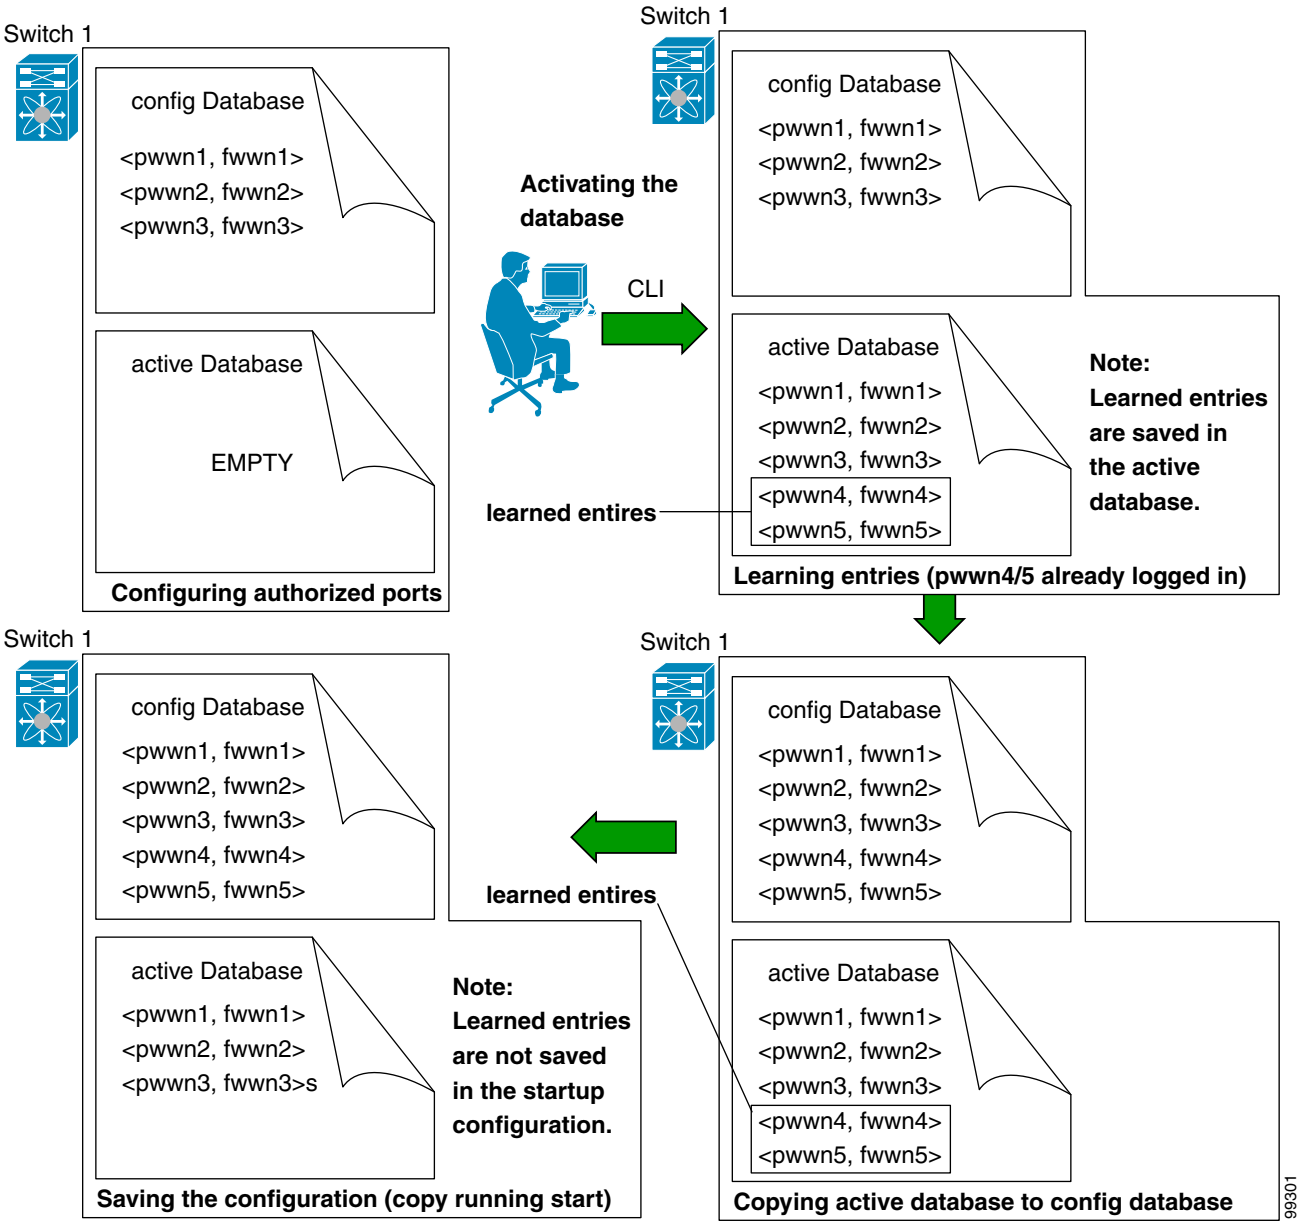

Guidelines and Limitations

- Port security is only supported for Fibre Channel ports.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Database Merge Guidelines

A database merge refers to a union of the configuration database and static (unlearned) entries in the active database.

When merging the database between two fabrics, follow these guidelines:

- Verify that the activation status and the auto-learning status is the same in both fabrics.
- Verify that the combined number of configurations for each VSAN in both databases does not exceed 2 K.

Caution If you do not follow these two conditions, the merge will fail. The next distribution will forcefully synchronize the databases and the activation states in the fabric.

Default Settings

Table 9-5 lists the default settings for all port security features in any switch.

Table 9-5 Default Security Settings

Parameters	Default
Auto-learn	Enabled if port security is enabled.
Port security	Disabled
Distribution	Disabled. **Note** Enabling distribution enables it on all VSANs in the switch.

Configuring Port Security

The steps to configure port security depend on which features you are using. Auto-learning works differently if you are using CFS distribution.

This section includes the following topics:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring Port Security with Auto-Learning and CFS Distribution

Detailed Steps

To configure port security, using auto-learning and CFS distribution, follow these steps:

Step 1 Enable port security. See the "Enabling Port Security" section on page 9-13.

Step 2 Enable CFS distribution. See the "Enabling Distribution" section on page 9-18.

Step 3 Activate port security on each VSAN. This turns on auto-learning by default. See the "Activating Port Security" section on page 9-13.

Step 4 Issue a CFS commit to copy this configuration to all switches in the fabric. See the "Committing the Changes" section on page 9-19. At this point, all switches are activated, and auto-learning.

Step 5 Wait until all switches and all hosts are automatically learned.

Step 6 Disable auto-learn on each VSAN. See the"Disabling Auto-learning" section on page 9-16.

Step 7 Issue a CFS commit to copy this configuration to all switches in the fabric. See the "Committing the Changes" section on page 9-19. At this point, the auto-learned entries from every switch are combined into a static active database that is distributed to all switches.

Step 8 Copy the active database to the configure database on each VSAN. See the "Copying the Port Security Database" section on page 9-19.

Step 9 Issue a CFS commit to copy this configuration to all switches in the fabric. See the "Committing the Changes" section on page 9-19. This ensures that the configure database is the same on all switches in the fabric.

Step 10 Copy the running configuration to the startup configuration, using the fabric option. This saves the port security configure database to the startup configuration on all switches in the fabric.

Configuring Port Security with Auto-Learning without CFS

Detailed Steps

To configure port security using auto-learning without CFS, follow these steps:

Step 1 Enable port security. See the "Enabling Port Security" section on page 9-13.

Step 2 Activate port security on each VSAN. This turns on auto-learning by default. See the "Activating Port Security" section on page 9-13.

Step 3 Wait until all switches and all hosts are automatically learned.

Step 4 Disable auto-learn on each VSAN. See the "Disabling Auto-learning" section on page 9-16.

Step 5 Copy the active database to the configure database on each VSAN. See the "Copying the Port Security Database" section on page 9-19.

Step 6 Copy the running configuration to the startup configuration This saves the port security configure database to the startup configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 7 Repeat Step 1 through Step 6 for all switches in the fabric.

Configuring Port Security with Manual Database Configuration

Detailed Steps

To configure port security and manually configure the port security database, follow these steps:

Step 1 Enable port security. See the "Enabling Port Security" section on page 9-13.

Step 2 Manually configure all port security entries into the configure database on each VSAN. See the "Configuring Port Security with Manual Database Configuration" section on page 9-11.

Step 3 Activate port security on each VSAN. This turns on auto-learning by default. See the "Activating Port Security" section on page 9-13.

Step 4 Disable auto-learn on each VSAN. See the "Disabling Auto-learning" section on page 9-16.

Step 5 Copy the running configuration to the startup configuration This saves the port security configure database to the startup configuration.

Step 6 Repeat Step 1 through Step 5 for all switches in the fabric.

Configuring Port Security Using the Configuration Wizard

The Port Security Configuration wizard provides step-by-step procedures for setting up the Port Security Policy for a selected VSAN. The Port Security Configuration wizard also supports the central management through CFS, making it possible to complete the entire configuration at one place.

The wizard automatically conducts few essential operations. For example, if you want central management, the wizard conducts operations to check CFS capability, enable CFS, and issue CFS commit at the proper stages.

To manage security at a particular port, you do not need to run through the wizard to configure the port security policy from the VSAN wide, but you can directly edit accesses on the port itself. This operation can be done through the Port Binding dialog box. If the port's belonging switch has not enabled port security yet, the dialog box enables security first. If the port security is enabled, the dialog box will edit the policy database based on user operations.

CFS should be enabled on all switches in the VSAN. A CFS master switch is selected to do all configurations. All changes will be distributed to the VSAN through the CFS **commit** command.

Prerequisites

- Enable port security on the switch.
- Define port security policy either manually by editing bound devices or switches or ports or by using autolearning.
- Activate port security policy.
- Ensure that activated and configured databases are synchronized through copy.
- Copy the activated database to be the startup configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Detailed Steps

To configure port security, follow these steps:

Step 1 Click the **Port Security** button on the toolbar.

Before launching the Port Security Setup Wizard, DCNM-SAN checks the CFS capability of the switches in the VSAN.

If VSAN context is not available, the wizard prompts to select VSAN as shown in Figure 9-2.

Figure 9-2 Select VSAN Window

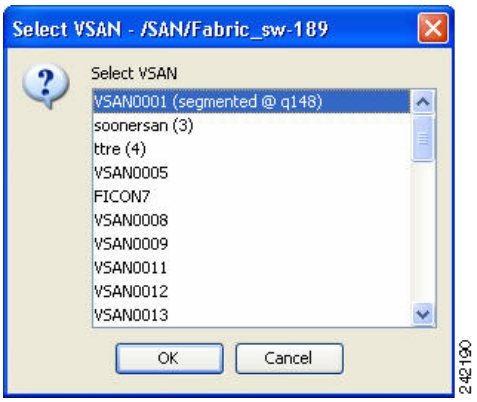

Step 2 Select the VSAN from the list and click **OK**.

Step 3 Do the following in the Select Master Switch page:

- Select the required master switch.
- Select **Automatically learn all logged in ports in VSAN** to Autolearn port configuration.

Step 4 Click **Next** to proceed.

You see the Edit and Activate Configuration page.

Note From Cisco NX-OS Release 5.2, devices can bind to vFC interfaces.

Step 5 Click **Insert** to create port binding.

Note When interfaces are inserted for binding, vFC ports can be selected.

Step 6 Two types of port binding can be created using the Insert Port Security Devices dialog box:

- Port WWN-pWWN bound to an interface WWN.
- Switch-Switch WWN bound to an interface. (Mainly useful for ISL binding).

Step 7 Select the type of port binding by clicking the radio buttons and enter the supporting values.

Step 8 Click **OK.**

Step 9 Click **Close** to exit the Insert Port Security window.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Note To delete an entry in the Edit and Activate Configuration page of the wizard, select the entry and click the **Delete** button.

Step 10 Click **Finish** to complete the Port Security Configuration for the selected switch.

Enabling Port Security

By default, the port security feature is disabled in all switches in the Cisco MDS 9000 Family.

Detailed Steps

To enable port security, follow these steps:

Step 1 Expand a **VSAN,** and then select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **CFS** tab.

Step 3 Enable CFS on all participating switches in the VSAN by clicking each entry in the Global column and selecting **enable**.

Step 4 Click **Apply Changes** to enable CFS distribution for the port security feature.

Step 5 Click the **Control** tab.

You see the port security enable state for all switches in the selected VSAN.

Step 6 Set the Command column to **enable** for each switch in the VSAN.

Step 7 Click the **CFS** tab and set the Command column to **commit** on all participating switches in the VSAN.

Step 8 Click **Apply Changes** to distribute the enabled port security to all switches in the VSAN.

Activating Port Security

Detailed Steps

To activate port security, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Actions** tab.

Step 3 Click in the Action column under Activation, next to the switch or VSAN on which you want to activate port security. You see a drop-down menu with the following options:

- **activate**—Valid port security settings are activated.
- **activate (TurnLearningOff)**—Valid port security settings are activated and auto-learn turned off.

Send documentation comments to dcnm-san-docfeedback@cisco.com

- **forceActivate**—Activation is forced.
- **forceActivate(TurnLearningOff)**—Activation is forced and auto-learn is turned off.
- **deactivate**—All currently active port security settings are deactivated.
- **NoSelection**— No action is taken.

Step 4 Set the Action field you want for that switch.

Step 5 Uncheck the **AutoLearn** check box for each switch in the VSAN to disable auto-learning.

Step 6 Click the **CFS** tab and set the command column to **commit** on all participating switches in the VSAN.

Step 7 Click **Apply Changes** in DCNM-SAN or **Apply** in Device Manager to save these changes.

Note If required, you can disable auto-learning (see the "Disabling Auto-learning" section on page 9-16).

Activating the Port Security Forcefully

If the port security activation request is rejected, you can force the activation.

Note An activation using the **force** option can log out existing devices if they violate the active database.

Detailed Steps

To forcefully activate the port security database, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Actions** tab.

Step 3 Click in the **Action** column under Activation, next to the switch or VSAN on which you want to activate port security and select the **forceactivate** option.

Step 4 Set the Action field you want for that switch.

Step 5 Click the **CFS** tab and set the command column to **commit** on all participating switches in the VSAN.

Step 6 Click **Apply Changes** in DCNM-SAN or **Apply** in Device Manager to save these changes.

Reactivating the Database

Detailed Steps

Tip If auto-learning is enabled, and you cannot activate the database, you will not be allowed to proceed .

To reactivate the port security database, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 1 Disable auto-learning.

Step 2 Copy the active database to the configured database.

Tip If the active database is empty, you cannot perform this step.

Step 3 Make the required changes to the configuration database.

Step 4 Activate the database.

Copying an Active Database to the Config Database

Detailed Steps

To copy the active database to the config database, follow these steps:

Step 1 Expand a **VSAN** and select **Port Security** in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Actions** tab.

You see the switches for that VSAN.

Step 3 Check the **CopyActive ToConfig** check box next to the switch for which you want to copy the database.

The active database is copied to the config database when the security setting is activated.

Step 4 Uncheck the **CopyActive ToConfig** check box if you do not want the database copied when the security setting is activated.

Step 5 Click the **CFS** tab and set the command column to **commit** on all participating switches in the VSAN.

Step 6 Click **Apply Changes** to save these changes or click **Undo Changes** to discard any unsaved changes.

Configuring Auto-learning

This section contains the following topics:

Enabling Auto-learning

Detailed Steps

To enable auto-learning, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Actions** tab.

Step 3 Click in the Action column under Activation, next to the switch or VSAN on which you want to activate port security. You see a drop-down menu with the following options:

- **activate**—Valid port security settings are activated.
- **activate (TurnLearningOff)**—Valid port security settings are activated and auto-learn turned off.
- **forceActivate**—Activation is forced.
- **forceActivate(TurnLearningOff)**—Activation is forced and auto-learn is turned off.
- **deactivate**—All currently active port security settings are deactivated.
- **NoSelection**— No action is taken.

Step 4 Select one of the port security options for that switch.

Step 5 Check the **AutoLearn** check box for each switch in the VSAN to enable auto-learning.

Step 6 Click the **Apply Changes** icon to save these changes.

Disabling Auto-learning

Detailed Steps

To disable auto-learning, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Actions** tab.

You see the switches for that VSAN.

Step 3 Uncheck the **AutoLearn** check box next to the switch if you want to disable auto-learning.

Step 4 Click the **Apply Changes** icon to save these changes.

Configuring Port Security Manually

This section includes the following topics:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Task Flow for Configuring Port Security

Follow these steps to configure port security on any switch in the Cisco MDS 9000 Family:

Step 1 Identify the WWN of the ports that need to be secured.

Step 2 Secure the fWWN to an authorized nWWN or pWWN.

Step 3 Activate the port security database.

Step 4 Verify your configuration.

Adding Authorized Port Pairs

Detailed Steps

After identifying the WWN pairs that need to be bound, add those pairs to the port security database.

Tip Remote switch binding can be specified at the local switch. To specify the remote interfaces, you can use either the fWWN or sWWN-interface combination.

To add authorized port pairs for port security, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

Step 2 Click the **Config Database** tab.

Step 3 Click **Create Row** to add an authorized port pair.

You see the Create Port Security dialog box.

Step 4 Double-click the device from the available list for which you want to create the port security setting.

Step 5 Double-click the port from the available list to which you want to bind the device.

Step 6 Click **Create** to create the port security setting.

Step 7 Click the **Apply Changes** icon to save these changes.

Deleting Port Security Setting

Detailed Steps

To delete a port security setting from the configured database on a switch, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

Step 2 Click the **Config Database** tab.

You see the configured port security settings for that VSAN.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 3 Click the row you want to delete.

Step 4 Click **Delete Row**.

You see the confirmation dialog box.

Step 5 Click **Yes** to delete the row, or click **No** to close the confirmation dialog box without deleting the row.

Step 6 Click the **Apply Changes** icon to save these changes.

Configuring Port Security Distribution

The port security feature uses the Cisco Fabric Services (CFS) infrastructure to enable efficient database management, provide a single point of configuration for the entire fabric in the VSAN, and enforce the port security policies throughout the fabric (see Chapter 7, "Configuring IPsec Network Security").

This section includes the following topics:

Enabling Distribution

Detailed Steps

All the configurations performed in distributed mode are stored in a pending (temporary) database. If you modify the configuration, you need to commit or discard the pending database changes to the configurations. The fabric remains locked during this period. Changes to the pending database are not reflected in the configurations until you commit the changes.

Note Port activation or deactivation and auto-learning enable or disable do not take effect until after a CFS commit if CFS distribution is enabled. Always follow any one of these operations with a CFS commit to ensure proper configuration. See the "Activation and Auto-learning Configuration Distribution" section on page 9-6.

Tip In this case, we recommend that you perform a commit at the end of each operation: after you activate port security and after you enable auto learning.

To enable distribution, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Control** tab.

You see the switches for that VSAN.

Step 3 In the Command column, select **enable** or **disable** from the drop-down menu.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 4 Click the **Apply Changes** icon to save the changes.

Locking the Fabric

The first action that modifies the existing configuration creates the pending database and locks the feature in the VSAN. Once you lock the fabric, the following situations apply:

- No other user can make any configuration changes to this feature.
- A copy of the configuration database becomes the pending database.

Committing the Changes

If you commit the changes made to the configurations, the configurations in the pending database are distributed to other switches. On a successful commit, the configuration change is applied throughout the fabric and the lock is released.

Interacting with the Database

This section includes the following topics:

Copying the Port Security Database

Detailed Steps

Tip We recommend that you copy the active database to the configuration database after disabling auto-learning. This action ensures that the configuration database is in sync with the active database. If distribution is enabled, this command creates a temporary copy (and consequently a fabric lock) of the configuration database. If you lock the fabric, you need to commit the changes to the configuration databases in all the switches.

To copy the active database to the configuration database, follow these steps:

Step 1 Expand a **Fabric**, expand a **VSAN,** and then select **Port Security** in the Logical Domains pane.

Step 2 Click the **Actions** tab. You see all the configuration databases.

Step 3 Select the appropriate configuration database and check the **Copy Active to Config** check box.

Step 4 Click the **Apply Changes** icon to save your changes.

Send documentation comments to dcnm-san-docfeedback@cisco.com

To view the differences between the active database and the configuration database, follow these steps:

Detailed Steps

Step 1 Expand a **Fabric**, expand a **VSAN,** and then select **Port Security** in the Logical Domains pane.

You see the Port Security information in the Information pane.

Step 2 Click the **Database Differences** tab. You see all the configuration databases.

Step 3 Select the appropriate configuration database. Select the **Active** or **Config** option to compare the differences between the selected database and the active or configuration database.

Step 4 Click the **Apply Changes** icon to save your changes.

Deleting the Port Security Database

Tip If the distribution is enabled, the deletion creates a copy of the database. An explicit deletion is required to actually delete the database.

Detailed Steps

To delete a port security database, follow these steps:

Step 1 Expand a **Fabric**, expand a **VSAN,** and then select **Port Security** in the Logical Domains pane.

You see the Port Security information in the Information pane.

Step 2 Click the **Config Database** tab. You see all the configuration databases.

Step 3 Select the appropriate configuration database and click the **Delete Row** button.

Step 4 Click **Yes** if you want to delete the configuration database.

Cleaning the Port Security Database

Detailed Steps

To clear all existing statistics from the port security database for a specified VSAN, follow these steps:

Step 1 Expand a **Fabric**, expand a **VSAN,** and then select **Port Security** in the Logical Domains pane.

You see the Port Security information in the Information pane.

Step 2 Click the **Statistics** tab.

You see all the configuration databases.

Step 3 Select the appropriate configuration database and check the **Clear** option.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 4 Click the **Apply Changes** icon to save your changes.

To clear any learned entries in the active database for a specified interface within a VSAN, follow these steps:

Step 1 Expand a **Fabric**, expand a **VSAN,** and then select **Port Security** in the Logical Domains pane.

You see the Port Security information in the Information pane.

Step 2 Select the **Actions** tab. You see all the configuration databases.

Step 3 Select the appropriate configuration database and check the **AutoLearn** option.

Step 4 Click the **Apply Changes** icon to save your changes.

Note You can clear the Statistics and the AutoLearn option only for switches that are local and do not acquire locks. Also, learned entries are only local to the switch and do not participate in distribution.

Verifying Port Security Configuration

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

- Displaying Activated Port Security Settings, page 9-21
- Displaying Port Security Statistics, page 9-21
- Displaying Port Security Violations, page 9-22

Displaying Activated Port Security Settings

To display active port security settings, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Active Database** tab.

You see the active port security settings for that VSAN.

Displaying Port Security Statistics

To display port security statistics, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

Send documentation comments to dcnm-san-docfeedback@cisco.com

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Statistics** tab.

You see the port security statistics for that VSAN.

Displaying Port Security Violations

Port violations are invalid login attempts (for example, login requests from unauthorized Fibre Channel devices). You can display a list of these attempts on a per-VSAN basis.

To display port security violations, follow these steps:

Step 1 Expand a **VSAN** and select **Port Securit**y in the Logical Domains pane.

You see the port security configuration for that VSAN in the Information pane.

Step 2 Click the **Violations** tab. You see the port security violations for that VSAN.

Field Descriptions for Port Security

The following are the field descriptions for port security.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Port Security Actions

Field	Description
Activation	
Action	• activate— Results in the valid port bindings on this VSAN/VLAN being activated. • activate (Turn LearningOff)— Results in the valid port bindings on this VSAN/VLAN being activated and copied to the active database and will also result in auto learn being turned off on this VSAN/VLAN, once the activation is complete. • force activate— Results in forced activation, even if there are errors during activation and the activated port bindings will be copied to the active database. • force activate (Turn Learning Off)—Results in forced activation along with turning auto learn off after activation and the activated port bindings will be copied to the active database. • deactivate— Results in deactivation of currently activated valid port bindings (if any), on this VSAN/VLAN. Currently active entries (if any), which would have been present in the active database, will be removed. • Activation will not be allowed on a VSAN if auto-learn is enabled on that VSAN
Enabled	The state of activation on this VSAN/VLAN. If true, then an activation has been attempted as the most recent operation on this VSAN/VLAN. If false, then an activation has not been attempted as the most recent operation on this VSAN/VLAN.
Result	Indicates the outcome of the most recent activation/deactivation.
Last Change	When the valid port bindings on this VSAN/VLAN were last activated. If the last activation took place prior to the last re-initialization of the agent, then this value will be N/A.
CopyActiveToConfig	If enabled, results in the active port binding database to be copied on to the configuration database on this VSAN/VLAN. Note that the learned entries are also copied.
AutoLearn	Helps to learn the valid port binding configuration of devices/ports logged into the local device on all its ports and populate the above active database with the same. This mechanism of learning the configuration of devices/ports logged into the local device over a period of time and populating the configuration is a convenience mechanism for users. If enabled on a particular VSAN, all subsequent logins (FLOGIs) on that VSAN will be populated in the enforced port binding database, provided it is not in conflict with existing enforced port bindings on that VSAN. When disabled, the mechanism of learning is stopped. The learned entries will however be in the active database.
Clear AutoLearned	

Send documentation comments to dcnm-san-docfeedback@cisco.com

Field	Description
Action	• Clear VSAN results in port bind auto-learned entries being cleared on this VSAN. • Clear Interface(s) results in port bind auto-learned entries being cleared on the interface specified on this VSAN.
Interface	Specifies the interface(s) on which the port bind auto-learned entries need to be cleared.

Port Security Config Database

Field	Description
Interface or fWWN	Represents the address of the port on the local device through which the device specified can FLOGI. • If fwwn, then the value is the fabric WWN of a port on the local device. • If intfIndex, then a port on the local device is being represented by its interface. • If wildCard, then it represents a wild-card entry. The wild-card represents any port on the local device.
Type	The mechanism to identify a switch port.
WWN	Represents the logging-in device address.
Available Interface	Displays the available interface. The interfaces available are: • Fibre Channel • PortChannel • Ethernet PortChannel • VFC

Port Security Active Database

Field	Description
Interface or fWWN	The address of a port on the local device.
Type	The mechanism to identify a switch port. • fwwn— The local switch port is identified by Fabric WWN(fWWN). • intfIndex— The local switch port is identified by ifIndex. • wildCard—Wild card (any switch port on local device).
WWN	Represents the logging-in device address.
IsLearnt	Indicates if this entry is a learned entry or not.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Port Security Database Differences

Field	Description
CompareWith	Specifies the database for the comparison. • configDb— Compares the configuration database with respect to active database on this VSAN/VLAN. So, the active database will be the reference database and the results of the difference operation will be with respect to the active database. • activeDb— Compares the active database with respect to configuration database on this VSAN/VLAN. So, the configuration database will be the reference database and the results of the difference operation will be with respect to the configuration database.
VSANId	The ID of the VSAN to compare against.
Interface/fWWN	The address of a port on the local device.
Type	The mechanism to identify a switch port. • fwwn— The local switch port is identified by Fabric WWN(fWWN). • intfIndex— The local switch port is identified by ifIndex. • wildCard— Wild card (any switch port on local device).
WWN	Represents the logging in device address.
Reason	Indicates the reason for the difference between the databases being compared, for this entry.

Port Security Violations

Field	Description
Interface	The fWWN of the port on the local device where the login was denied.
End Device	The pWWN of the device that was denied FLOGI on one of the local device's ports.
Or Switch	The sWWN of the device (if the device happens to be a switch), that was denied entry on one of the local device's ports.
Time	When the login denial took place.
Count	The number of times this particular pWWN/nWWN or sWWN has been denied login on this particular local interface.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Port Security Statistics

Field	Description
AllowedLogins	The number of FLOGI requests that have been allowed on this VSAN/VLAN.
DeniedLogins	The number of FLOGI requests that have been denied on this VSAN/VLAN.
Clear	When set to clear, it results in port bind statistic counters being cleared on this VSAN/VLAN.

Feature History for Port Security

Table 9-6 lists the release history for this feature. Only features that were introduced or modified in Cisco NX-OS Release 5.x or a later release appear in the table.

Table 9-6 Feature History for Port Security

Feature Name	Releases	Feature Information
vFC Interfaces	5.2	Devices can be bound to vFC interfaces.

Send documentation comments to dcnm-san-docfeedback@cisco.com

CHAPTER 10

Configuring Fabric Binding

This chapter describes the fabric binding feature provided in the Cisco MDS 9000 Family of directors and switches. It includes the following sections:

This chapter includes the following topics:

Information About Fabric Binding

The fabric binding feature ensures ISLs are only enabled between specified switches in the fabric binding configuration. Fabric binding is configured on a per-VSAN basis.

This feature helps prevent unauthorized switches from joining the fabric or disrupting current fabric operations. It uses the Exchange Fabric Membership Data (EFMD) protocol to ensure that the list of authorized switches is identical in all switches in the fabric.

This section inludes the following topics:

Port Security Versus Fabric Binding

Port security and fabric binding are two independent features that can be configured to complement each other. Table 10-1 compares the two features.

Table 10-1 Fabric Binding and Port Security Comparison

Fabric Binding	Port Security
Uses a set of sWWNs and a persistent domain ID.	Uses pWWNs and nWWNs or fWWNs and sWWNs.
Binds the fabric at the switch level.	Binds devices at the interface level.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Table 10-1 Fabric Binding and Port Security Comparison (continued)

Fabric Binding	Port Security
Authorizes only the configured sWWN stored in the fabric binding database to participate in the fabric.	Allows a preconfigured set of Fibre Channel devices to logically connect to a SAN ports. The switch port, identified by a WWN or interface number, connects to a Fibre Channel device (a host or another switch), also identified by a WWN. By binding these two devices, you lock these two ports into a group (or list).
Requires activation on a per VSAN basis.	Requires activation on a per VSAN basis.
Allows specific user-defined switches that are allowed to connect to the fabric, regardless of the physical port to which the peer switch is connected.	Allows specific user-defined physical ports to which another device can connect.
Does not learn about switches that are logging in.	Learns about switches or devices that are logging in if learning mode is enabled.
Cannot be distributed by CFS and must be configured manually on each switch in the fabric.	Can be distributed by CFS.

Port-level checking for xE ports is as follows:

- The switch login uses both port security binding and fabric binding for a given VSAN.
- Binding checks are performed on the port VSAN as follows:
  - E port security binding check on port VSAN
  - TE port security binding check on each allowed VSAN

While port security complements fabric binding, they are independent features and can be enabled or disabled separately.

Fabric Binding Enforcement

To enforce fabric binding, configure the switch world wide name (sWWN) to specify the xE port connection for each switch. Enforcement of fabric binding policies are done on every activation and when the port tries to come up. In a FICON VSAN, the fabric binding feature requires all sWWNs connected to a switch and their persistent domain IDs to be part of the fabric binding active database. In a Fibre Channel VSAN, only the sWWN is required; the domain ID is optional.

Note All switches in a Fibre Channel VSAN using fabric binding must be running Cisco MDS SAN-OS Release 3.0(1) and NX-OS Release 4.1(1b) or later.

Licensing Requirements for Fabric Binding

Fabric binding requires that you install either the MAINFRAME_PKG license or the ENTERPRISE_PKG license on your switch.

Send documentation comments to dcnm-san-docfeedback@cisco.com

See the *Cisco MDS 9000 Family NX-OS Licensing Guide* for more information on license feature support and installation.

Default Settings

Table 10-2 lists the default settings for the fabric binding feature.

Table 10-2 Default Fabric Binding Settings

Parameters	Default
Fabric binding	Disabled

Configuring Fabric Binding

Detailed Steps

To configure fabric binding in each switch in the fabric, follow these steps:

Step 1 Enable the fabric configuration feature.

Step 2 Configure a list of sWWNs and their corresponding domain IDs for devices that are allowed to access the fabric.

Step 3 Activate the fabric binding database.

Step 4 Copy the fabric binding active database to the fabric binding config database.

Step 5 Save the fabric binding configuration.

Step 6 Verify the fabric binding configuration.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

CHAPTER 11

Configuring Cisco TrustSec Fibre Channel Link Encryption

This chapter provides an overview of the Cisco TrustSec Fibre Channel (FC) Link Encryption feature and describes how to configure and set up link-level encryption between switches.

This chapter includes the following topics:

Information About Cisco TrustSec FC Link Encryption

Cisco TrustSec FC Link Encryption is an extension of the Fibre Channel-Security Protocol (FC-SP) feature and uses the existing FC-SP architecture to provide integrity and confidentiality of transactions. Encryption is now added to the peer authentication capability to provide security and prevent unwanted traffic interception. Peer authentication is implemented according to the FC-SP standard using the Diffie-Hellman Challenge Handshake Authentication Protocol (DHCHAP) protocol.

Note Cisco TrustSec FC Link Encryption is currently only supported between Cisco MDS switches. This feature is not supported when you downgrade to software versions which do not have the Encapsulating Security Protocol (ESP) support.

This section includes the following topics:

Supported Modules

The following modules are supported for the Cisco TrustSec FC Link Encryption feature:

Send documentation comments to dcnm-san-docfeedback@cisco.com

- 1/2/4/8 Gbps 24-Port Fibre Channel switching module (DS-X9224-96K9)
- 1/2/4/8 Gbps 48-Port Fibre Channel switching module (DS-X9248-96K9)
- 1/2/4/8 Gbps 4/44-Port Fibre Channel switching module (DS-X9248-48K9)

Cisco TrustSec FC Link Encryption Terminology

The following Cisco TrustSec FC Link Encryption-related terms are used in this chapter:

- Galois Counter Mode (GCM)—A block cipher mode of operation providing confidentiality and data-origin authentication.
- Galois Message Authentication Code (GMAC)—A block cipher mode of operation providing only data-origin authentication. It is the authentication-only variant of GCM.
- Security Association (SA)—A connection that handles the security credentials and controls how they propagate between switches. The SA includes parameters such as salt and keys.
- Key—A 128-bit hexadecimal string that is used for frame encryption and decryption. The default value is zero.
- Salt —A 32-bit hexadecimal number that is used during encryption and decryption. The same salt must be configured on both sides of the connection to ensure proper communication. The default value is zero.
- Security Parameters Index (SPI) number—A 32-bit number that identifies the SA to be configured to the hardware. The range is from 256 to 4,294,967,295.

Support for AES Encryption

The Advanced Encryption Standard (AES) is the symmetric cipher algorithm that provides a high-level of security, and can accept different key sizes.

The Cisco TrustSec FC Link Encryption feature supports the 128-bit AES for security encryption and enables either AES-GCM or AES-GMAC for an interface. The AES-GCM mode provides encryption and authentication of the frames and AES-GMAC provides only the authentication of the frames that are being passed between the two peers.

Guidelines and Limitations

This section lists the guidelines for Cisco TrustSec FC Link Encryption:

- Ensure that Cisco TrustSec FC Link Encryption is enabled only between MDS switches. This feature is supported only on E-ports or the ISLs, and errors will result if non-MDS switches are used.
- Ensure that the peers in the connection have the same configurations. If there are differences in the configurations, a “port re-init limit exceeded” error message is displayed.
- Before applying the SA to the ingress and egress hardware of a switch interface, ensure that the interface is in the admin shut mode.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Configuring Cisco TrustSec Fibre Channel Link Encryption

This section includes the following topics:

Enabling Cisco TrustSec FC Link Encryption

By default, the FC-SP feature and the Cisco TrustSec FC Link Encryption feature are disabled in all switches in the Cisco MDS 9000 Family.

You must explicitly enable the FC-SP feature to access the configuration and verification commands for fabric authentication and encryption. When you disable this feature, all related configurations are automatically discarded.

Setting Up Security Associations

To perform encryption between the switches, a security association (SA) needs to be set up. An administrator manually configures the SA before the encryption can take place. The SA includes parameters such as keys and salt, that are required for encryption. You can set up to 2000 SAs in a switch.

Note Cisco TrustSec FC Link Encryption is currently supported only on DHCHAP on and off modes.

Setting Up Security Association Parameters

Detailed Steps

To set up the SA parameters, such as keys and salt, using DCNM-SAN, follow these steps:

Step 1 Expand **Switches > Security,** and then select **FC-SP (DHCHAP)**.

You see the FC-SP configuration in the Information pane.

Step 2 Click the **SA** tab.

You see the SA parameters for each switch.

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

You see the Create SA Parameters dialog box.

Step 4 Select the switches on which you want to perform an encryption.

Step 5 Select a value for the SP. The range is from 256 to 65536.

Step 6 Enter a value for the salt. Alternatively, click **Salt Generator** to select a value

Step 7 Enter a value for the key. Alternatively, click **Key Generator** to select a value.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 8 Click **Create** to save the changes.

To set up the SA parameters, such as keys and salt, using Device Manager, follow these steps:

Step 1 Choose **Switches > Security,** and then select **FC-SP**.

You see the FC-SP configuration dialog box.

Step 2 Click the **SA** tab.

You see the SA parameters for each switch.

Step 3 Click **Create** to create new parameters.

You see the Create FC-SP SA dialog box.

Step 4 Select a value for the SP. The range is from 256 to 65536.

Step 5 Enter a value for the salt. Alternatively, click **Salt Generator** to select a value

Step 6 Enter a value for the key. Alternatively, click **Key Generator** to select a value.

Step 7 Click **Create** to save the changes.

Configuring ESP Settings

Note To apply the SA to the ingress and egress hardware of an interface, the interface needs to be in the admin shut mode.

Note The ESP modes are set only after a SA is configured to either the ingress or the egress hardware. If SA has not been configured, ESP is turned off and encapsulation does not occur.

Note An ESP mode change always needs a port flap because the change is not seamless if it is done after you configure the port; although the configurations are not rejected.

Detailed Steps

To configure ESP settings, follow these steps:

Step 1 Expand **Switches > Security,** and then select **FC-SP (DHCHAP)**.

You see the FC-SP configuration in the Information pane.

Step 2 Click the **ESP Interfaces** tab.

You see the Interface details for each switch.

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

You see the Create ESP Interfaces dialog box.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 4 Select the switches on which you want to perform an encryption.

Step 5 Enter an interface for the selected switch.

Step 6 Select the appropriate ESP mode for the encryption.

Step 7 Enter the appropriate egress port for the encryption.

Step 8 Enter the appropriate ingress port for the encryption.

Step 9 Click **Create** to save the changes.

To configure ESP settings using Device Manager, follow these steps:

Step 1 Expand **Switches > Security,** and then select **FC-SP**.

You see the FC-SP configuration dialog box.

Step 2 Click the **ESP Interfaces** tab.

You see the Interface details for each switch.

Step 3 Click **Create.**

You see the Create FC-SP ESP Interfaces dialog box.

Step 4 Enter an interace for any switch for encryption. Alternatively, you can select values from the available interfaces for the selected switch.

Step 5 Select the appropriate ESP mode for the encryption.

Step 6 Enter the appropriate egress port for the encryption.

Step 7 Enter the appropriate ingress port for the encryption.

Step 8 Click **Create** to save the changes.

Configuring ESP Using ESP Wizard

You can configure and set up link-level encryption between switches using ESP wizard. You can configure an existing Inter-Switch Link (ISL) as a secure ISL or edit an existing secure ingress SPI and egress SPI using this wizard.

Detailed Steps

To configure ESP using ESP wizard, follow these steps:

Step 1 Right-click **Tools > Security> FC-SP ESP Link Security** to launch the ESP wizard from DCNM-SAN.

Step 2 Select the appropriate ISL to secure or edit security.

Note Only ISLs with FC-SP port mode turned on and available on ESP- capable switches or blades are displayed.

Step 3 Create new Security Associations (SAs).

Send documentation comments to dcnm-san-docfeedback@cisco.com

You can create a new SA for each switch or use the existing SAs. You can click **View Existing SA** to view the existing SAs.

Note The existing list of SAs displays all existing SAs for a switch. The wizard runs only when a pair of switches have a common SA. The wizard checks for this requirement when you select **Next** and a warning message is displayed if a pair of switches do not have a common SA. You must create a common SA on the pair of the switches to run this wizard.

Step 4 Specify the Egress port, Ingress port, and ESP mode for the selected ISL.

The Egress and Ingress ports are auto populated with SPIs of the SAs common to a pair of switches incase of a secured ISL.

In this scenario, the mode is disabled and you cannot edit the modes for a secured ISL.

Note You can modify an existing ESP configuration provided the selected ISLs are enabled.

Step 5 Review your configuration.

Step 6 Click **Finish** to start the configuration for the ESP setup. You can view the status of the configuration in the status column.

Changing Keys for Switches

After the SA is applied to the ingress and egress ports, you should change the keys periodically in the configuration. The keys should be changed sequentially to avoid traffic disruption.

Verifying Cisco TrustSec Fibre Channel Link Encryption Configuration

You can view information about the Cisco TrustSec FC Link Encryption feature using the **show** commands in DCNM-SAN or Device Manager.

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

This section has the following topics:

Displaying FC-SP Interface Statistics

You can view the statistics data that displays the Encapsulating Security Protocol-ESP Security Parameter (SPI) mismatches and Interface-Encapsulating Security Protocol authentication failures information using DCNM-SAN.

To view the ESP statistics for an interface, follow these steps:

Send documentation comments to dcnm-san-docfeedback@cisco.com

Step 1 Expand **Interfaces > FC Physical,** and then select **FC-SP**.

You see the FC-SP configuration in the Information pane.

Step 2 Click the **FC-SP** tab.

You see view the FC-SP statistics data in the Information pane.

Step 3 Click **Refresh** to refresh the statistics data.

Displaying FC-SP Interface Statistics Using Device Manager

To view the ESP statistics for an interface using Device Manager, follow these steps:

Step 1 Choose **Security > FC Physical,** and then select **FC-SP**.

You see the FC-SP configuration in the Information pane.

Step 2 Click the **Statistics** tab.

You see the statistics in the Information pane.

Step 3 Click **Refresh** to refresh the statistics data.

Send documentation comments to dcnm-san-docfeedback@cisco.com

Send documentation comments to dcnm-san-docfeedback@cisco.com

INDEX

Symbols

Numerics

A

C

Send documentation comments to dcnm-san-docfeedback@cisco.com

D

Send documentation comments to dcnm-san-docfeedback@cisco.com

E

F

G

H

I

Send documentation comments to dcnm-san-docfeedback@cisco.com

L

M

N

Send documentation comments to dcnm-san-docfeedback@cisco.com

O

P

R

Send documentation comments to dcnm-san-docfeedback@cisco.com

S

T

Send documentation comments to dcnm-san-docfeedback@cisco.com

U

V

W

Send documentation comments to dcnm-san-docfeedback@cisco.com