CHAPTER 3

Cisco SME Cluster Management

The Cisco Fabric Manager provides a web-browser interface that displays real-time views of your network fabrics and lets you configure Cisco Storage Media Encryption with easy-to-use wizards. This chapter contains information about Cisco SME initial configuration and the tasks that are used to manage Cisco SME clusters using Cisco Fabric Manager.

About SME Cluster Management

An SME cluster consists of a group of MDS switches running the SME application in a single fabric environment where each switch is a member or node. The cluster infrastructure enables the SME application to offer high availability and load balancing by providing the ability to communicate and coordinate with the other members to maintain a consistent and distributed view of the application's configuration and operational state.

This chapter contains the following sections:

- Creating a Cisco SME Cluster Using the Cisco SME Wizard, page 3-2
- Archiving and Purging a Cisco SME Cluster, page 3-23
- Viewing Cisco SME Cluster Details, page 3-26
- Viewing Cluster States, page 3-26
- Viewing Members in a Cluster, page 3-27
- Viewing Cluster Information Using Fabric Manager Client, page 3-28
- Viewing Cluster Information Using Device Manager, page 3-29
- Cluster Quorum and Master Switch Election Overview, page 3-30
- In-Service Software Upgrade (ISSU) In a Two-Node Cluster, page 3-33

The process of configuring Cisco SME on an MDS switch with an installed MSM-18/4 module or on a Cisco MDS 9222i switch involves a number of configuration tasks that should be followed in chronological order. See the topics in the Before You Begin online help in Fabric Manager Web Server. Refer to Chapter 2, "Getting Started" and Chapter 4, "Cisco SME Interface Configuration" for information about the tasks that must be completed before creating an Cisco SME cluster.

Creating a Cisco SME Cluster Using the Cisco SME Wizard

The Cisco SME Wizard is an easy-to-use interface that walks you through the process of creating a Cisco SME cluster. The following sections describe the steps in this process:

- Launching Cisco SME Wizard, page 3-2
- Choosing a Cluster Name, page 3-3
- Selecting Fabrics, page 3-4
- Selecting Interfaces, page 3-4
- Selecting Master Key Security Levels, page 3-5
- Selecting Media Key Settings, page 3-8
- Specifying the Key Management Center Server, page 3-9
- Selecting Transport Settings, page 3-10
- Confirming the Cluster Creation, page 3-11
- Downloading Key File and Storing Keyshares, page 3-12
- Viewing the Transport Settings in Cluster Detail Page, page 3-21

Launching Cisco SME Wizard

To launch the Cisco SME wizard, follow these steps:

Step 1 Open the web browser to the Fabric Manager Web Client. Log in with the user name and password.

For login information, refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*

Figure 3-1 Cisco Fabric Manager Login

.

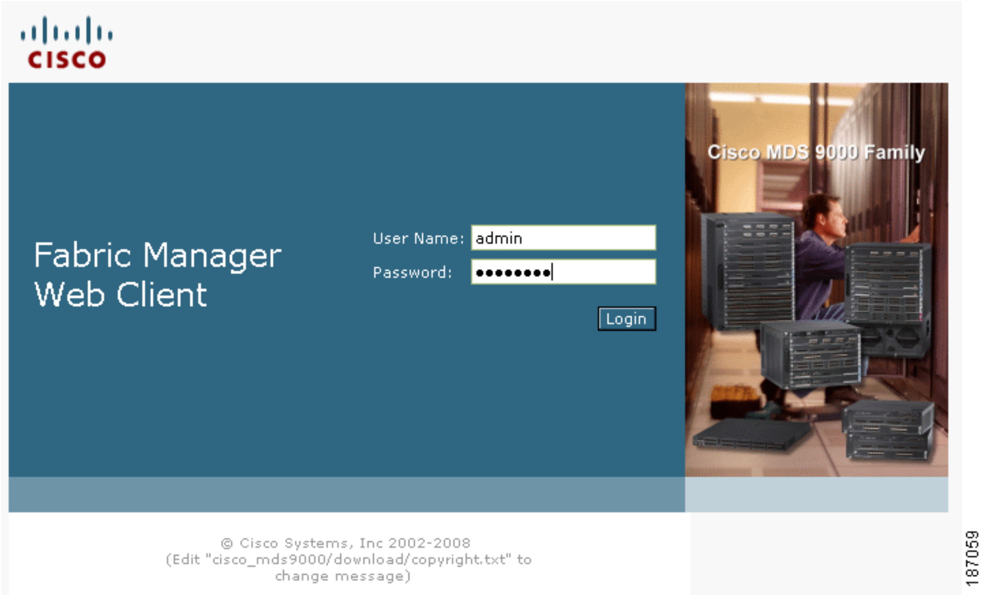

User Name: admin
Password: ••••••••

© Cisco Systems, Inc 2002-2008
(Edit "cisco_mds9000/download/copyright.txt" to change message)

Step 2 In the Fabric Manager Web Client, click the **SME** tab.

Step 3 Select **Clusters** in the navigation pane.

Figure 3-2 Fabric Manager Web Client SME Tab

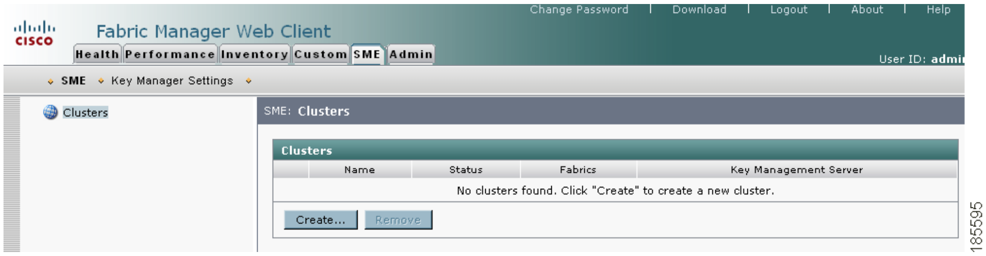

SME: Clusters

Clusters

Name	Status	Fabrics	Key Management Server

No clusters found. Click "Create" to create a new cluster.

Step 4 Click **Create** in the information pane.

The Cisco SME wizard launches to walk you through the easy configuration process.

Choosing a Cluster Name

In the Choose Name screen, enter a cluster name. Click **Next**.

Figure 3-3 Cisco SME Wizard - Create a Cluster Name

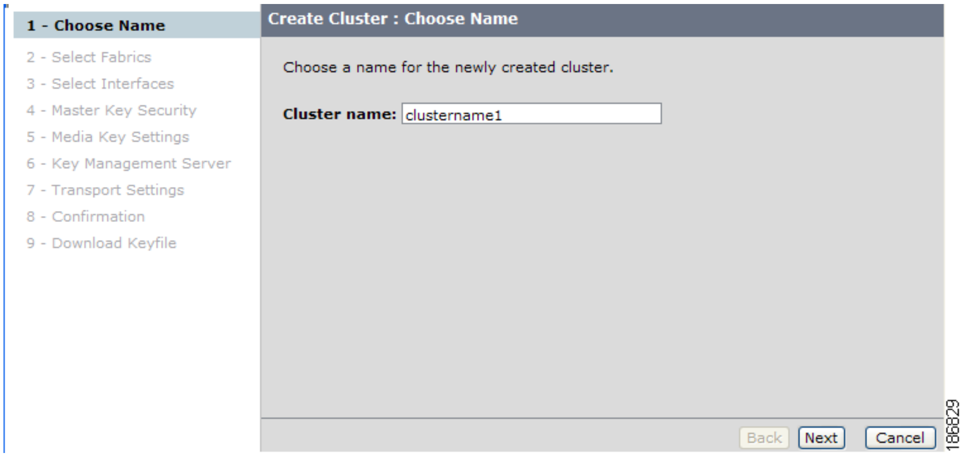

Create Cluster : Choose Name

Choose a name for the newly created cluster.

Cluster name: clustername1

Note Cluster names must not contain spaces or special characters.

Selecting Fabrics

In the Selecting Fabrics screen, highlight the fabric you want to include in the cluster. Click **Next**.

Figure 3-4 Cisco SME Wizard - Select Fabrics

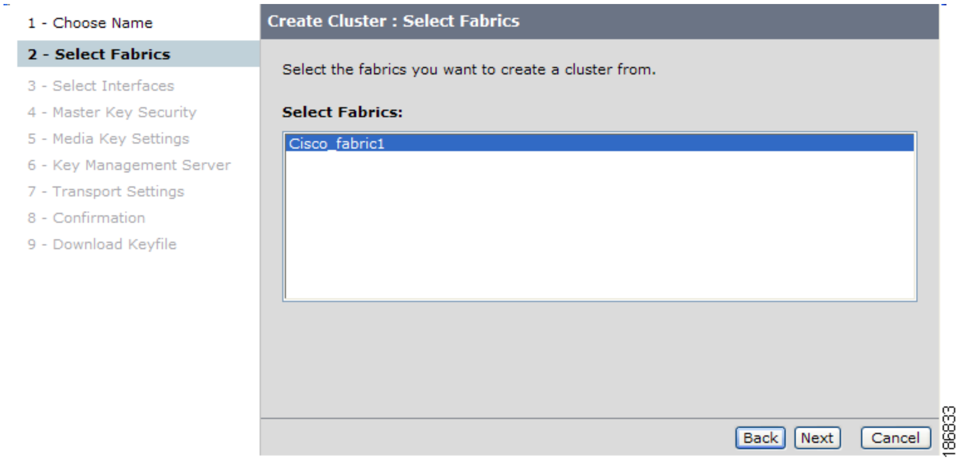

Select the fabrics you want to create a cluster from.

Select Fabrics:

Cisco_fabric1

Selecting Interfaces

In the Selecting Interfaces screen, highlight the SME interfaces you want to include in your cluster. Click **Next**. For information about adding interfaces, see Chapter 4, "Cisco SME Interface Configuration."

Note Cisco SAN-OS Release 3.2(2) supports one cluster per switch.

Figure 3-5 Select the Cisco SME Interfaces

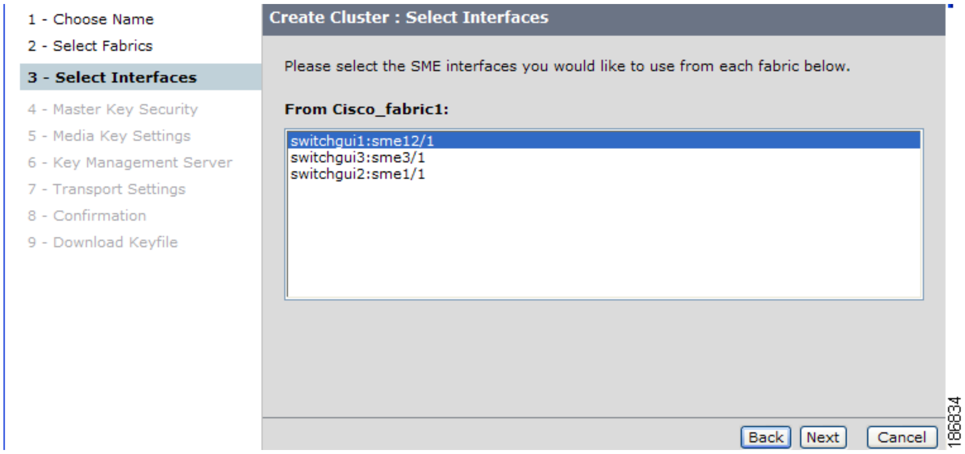

Please select the SME interfaces you would like to use from each fabric below.

From Cisco_fabric1:

switchgui1:sme12/1
switchgui3:sme3/1
switchgui2:sme1/1

Selecting Master Key Security Levels

There are 3 master key security levels: Basic, Standard, and Advanced. Standard and Advanced security levels require smart cards. Table 3-1 describes the master key security levels.

Caution You can not modify the cluster security level after a cluster is created. Before confirming the cluster creation, you will be prompted to review the cluster details. At that time, you can return to modify the security level.

Note For information on cluster security, see Cisco Storage Media Encryption Security Overview, page 1-9 and Master Key Security Modes, page 6-3.

Table 3-1 Master Key Security Levels

Security Level	Definition
Basic	The master key is stored in a file and encrypted with a password. To retrieve the master key, you need access to the file and the password.
Standard	Standard security requires one smart card. When you create a cluster and the master key is generated, you are prompted to insert the smart card into the smart card reader. The master key is then written to the smart card. To retrieve the master key, you need the smart card and the smart card pin.
Advanced	Advanced security requires 5 smart cards. When you create a cluster and select Advanced security mode, you designate the number of smart cards (2 or 3 of 5 smart cards or 2 of 3 smart cards) that are required to recover the master key when data needs to be retrieved. For example, if you specify 2 of 5 smart cards, then you will need 2 of the 5 smart cards to recover the master key. Each smart card is owned by a Cisco SME Recovery Officer. **Note** The greater the number of required smart cards to recover the master key, the greater the security. However, if smart cards are lost or if they are damaged, this reduces the number of available smart cards that could be used to recover the master key.

Note For Basic and Standard security modes, one user should hold the Cisco SME Administrator and the Cisco SME Recovery Officer roles.

In the Master Key Security screen, select the cluster security type you wish to use. You can choose any of the following security levels:

- Selecting Basic Security, page 3-6
- Selecting Standard Security, page 3-6
- Selecting Advanced Security, page 3-7

Send documentation comments to mdsfeedback-doc@cisco.com

Selecting Basic Security

In the Master Key Security screen, select **Basic.** Click **Next**.

Figure 3-6 Selecting Basic Security

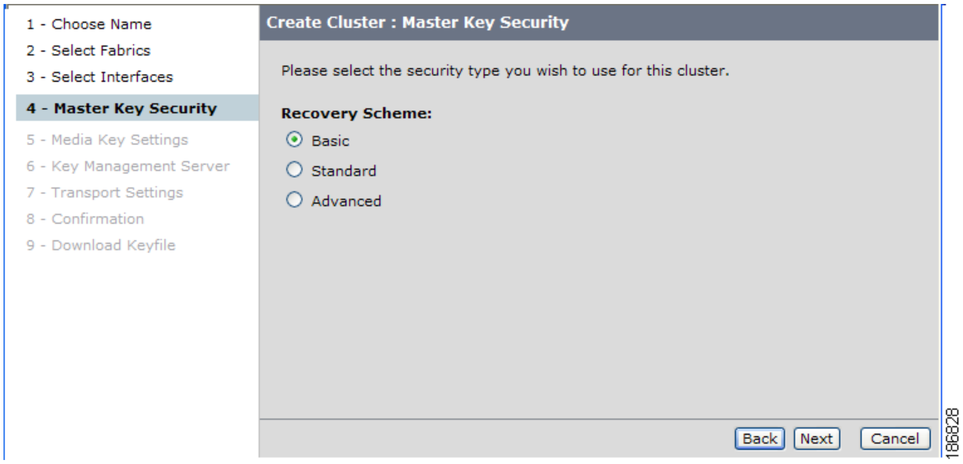

Create Cluster : Master Key Security

Please select the security type you wish to use for this cluster.

Recovery Scheme:

● Basic
○ Standard
○ Advanced

For the Basic security level, after the cluster is created the switch generates the master key file and you are prompted for a password to protect the file.

Selecting Standard Security

In the Master Key Security screen, select **Standard** and click **Next**. For Standard security, one Cisco SME Recovery Officer must be present to login and enter the smart card PIN.

Figure 3-7 Selecting Standard Security

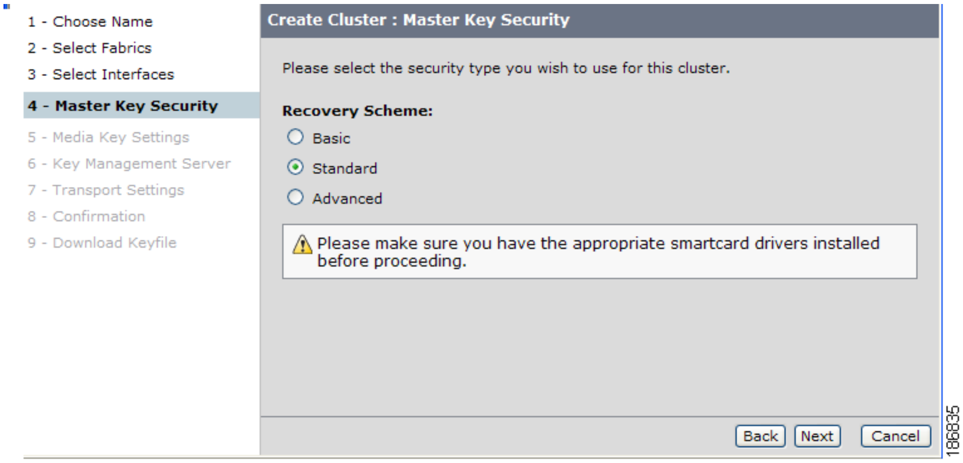

Create Cluster : Master Key Security

Please select the security type you wish to use for this cluster.

Recovery Scheme:

○ Basic
● Standard
○ Advanced

⚠ Please make sure you have the appropriate smartcard drivers installed before proceeding.

Selecting Advanced Security

When Advanced security is selected, you need to designate the number of cards that are required to recover the master key. This can be 2 or 3 of 5 smart cards or 2 of 3 smart cards. You will need to configure all 5 smart cards during the cluster creations process; however, you will only need the quorum number (that you designated in this step), to recover the master key.

In the Master Key Security screen, select **Advanced**. Enter the number of required smart cards for the quorum (2 of 3 or 2 of 5 or 3 of 5). Click **Next**.

- For Advanced security, 5 Cisco SME Recovery Officers must be present to login and enter the smart card PIN for each of the 5 smart cards.
- Be sure that the smart card reader is connected using the USB port (see Installing Smart Card Drivers, page 2-21 in Chapter 2, "Getting Started").
- When you insert a smart card into the reader, the card will be verified. You will be prompted to initialize the card if the card has not been previously initialized.

Figure 3-8 Selecting Advanced Security

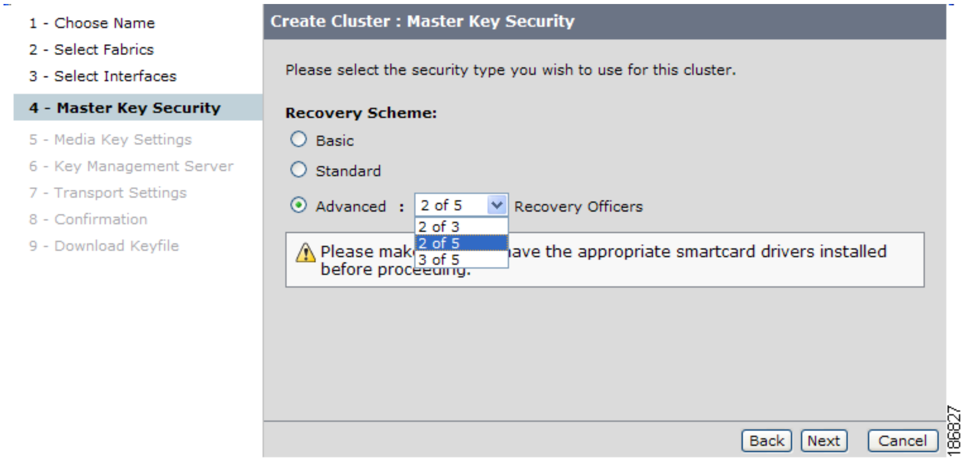

Create Cluster : Master Key Security

Please select the security type you wish to use for this cluster.

Recovery Scheme:

○ Basic

○ Standard

● Advanced : 2 of 5 Recovery Officers

2 of 3
2 of 5
3 of 5

⚠ Please mak ave the appropriate smartcard drivers installed before proceeding.

Selecting Media Key Settings

Caution You can not modify the media key settings after a cluster is created.

In the Media Key Settings screen, select the media key settings.

Figure 3-9 Media Key Settings

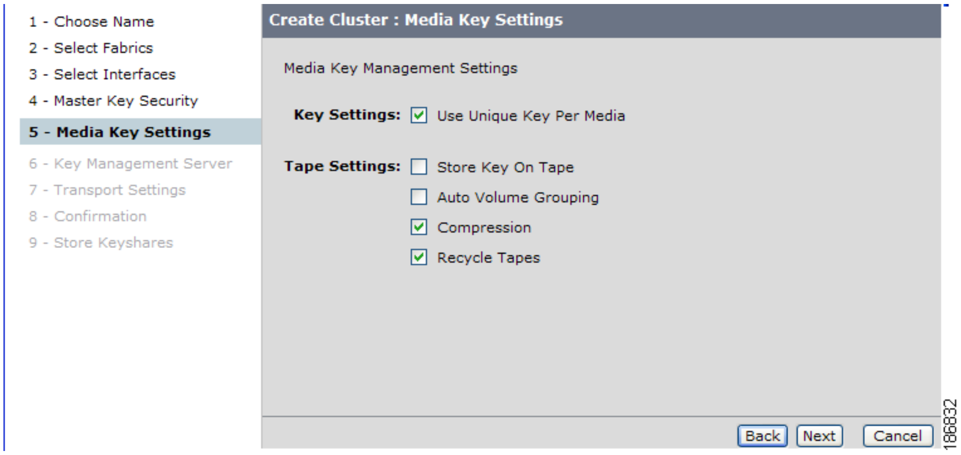

Create Cluster : Media Key Settings

Media Key Management Settings

Key Settings: ☑ Use Unique Key Per Media

Tape Settings: □ Store Key On Tape

□ Auto Volume Grouping

☑ Compression

☑ Recycle Tapes

Table 3-2 lists the media key settings and definitions.

For additional information on media key settings, see Key Management Settings, page 6-4.

Table 3-2 Media Key Settings

Media Key Setting	Definition
Use unique key per media	In unique key mode, a unique key is issued for each tape volume. The default is unique key mode.
Store key on tape	If you choose unique key mode (see above), this mode allows you to store the encrypted media key on the tape volume not in the Cisco KMC. This provides better scaling when your backup environment includes a large number of tapes. This is recommended for managing a large number of tape volume keys. Key-on-tape mode is disabled by default.
Auto-volume grouping	Cisco SME automatically creates a volume group and categorizes the appropriate tape volumes encrypted under this group based on the backup application's volume pool configuration. Auto-volume grouping is disabled by default.
Compression	Cisco SME can perform compression followed by encryption if this option is selected. Compression is enabled by default. **Note** Compression will be enabled for a tape drive in one of two ways: (a) configuration or (b) if the compression is not enabled through configuration and the tape drive is enabled for compression, compression is implicitly enabled for this tape drive.
Recycle Tapes	Select this option to enable purging of the keys upon tape recycling. When a tape is recycled or re-labeled, a new key is generated and used for encryption. Enabling this option purges the key that was used to encrypt data before the tape was recycled. **Note** This option must be disabled if the tapes are cloned offline without the involvement of the backup application itself. Tape recycling is enabled by default.

Specifying the Key Management Center Server

In the Key Management Server screen, specify the key management center server from the drop down menu. Click **Next**.

For information about key management, see Chapter 6, “Cisco SME Key Management.”

Note The Cisco Key Management Center should be the same location as the Fabric Manager Server.

Figure 3-10 Specify Key Management Server

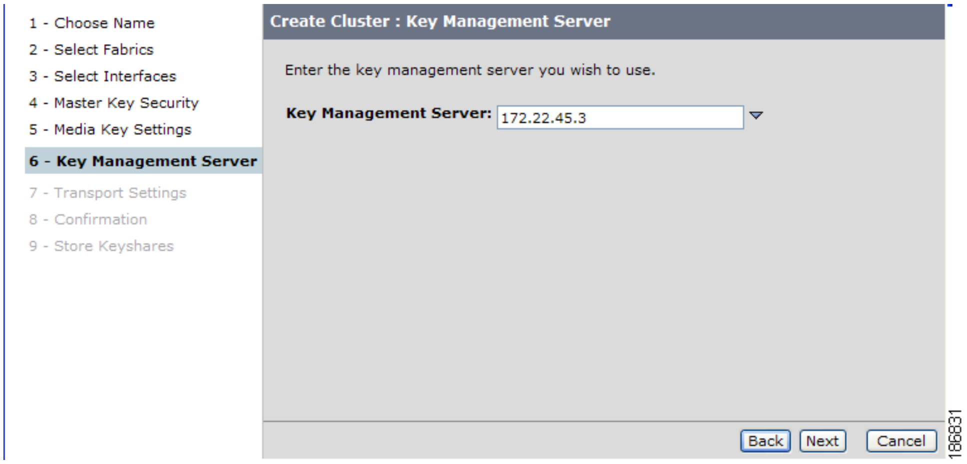

1 - Choose Name
2 - Select Fabrics
3 - Select Interfaces
4 - Master Key Security
5 - Media Key Settings
6 - Key Management Server
7 - Transport Settings
8 - Confirmation
9 - Store Keyshares

Create Cluster : Key Management Server

Enter the key management server you wish to use.

Key Management Server: 172.22.45.3

Selecting Transport Settings

In the Transport Settings screen, to enable Transport Settings, select **On**. If enabled, specify the Trust Point from the dropdown menu.

For more information about Trust Points, see *Cisco MDS 9000 Family CLI Configuration Guide*.

To enable Transport settings, select **On** as shown in Figure 3-11.

Figure 3-11 Transport Settings - On

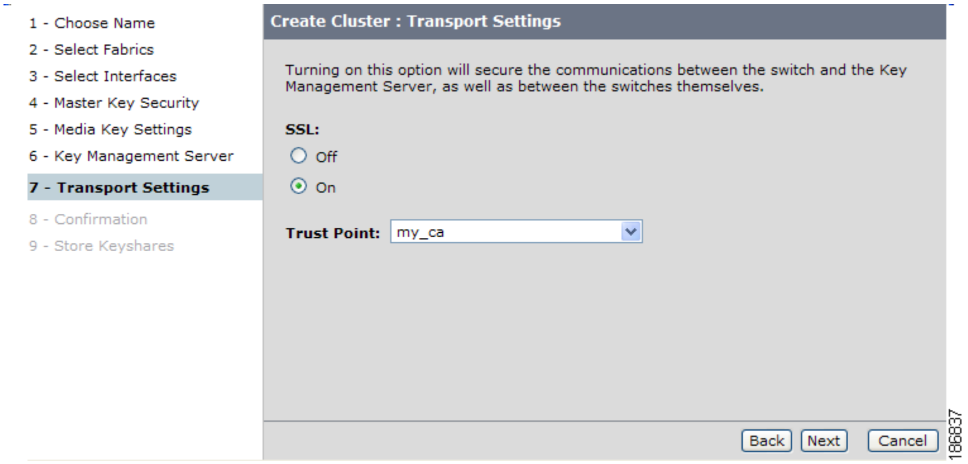

1 - Choose Name
2 - Select Fabrics
3 - Select Interfaces
4 - Master Key Security
5 - Media Key Settings
6 - Key Management Server
7 - Transport Settings
8 - Confirmation
9 - Store Keyshares

Create Cluster : Transport Settings

Turning on this option will secure the communications between the switch and the Key Management Server, as well as between the switches themselves.

SSL:
○ Off
● On

Trust Point: my_ca

If **On** is selected in the Transport Setting, SSL is enabled on KMC with the following results:

- New clusters are created. If **Off** is selected, cluster creation fails.
- Previously created clusters are updated by enabling SSL with trustpoint on the switches. KMC server connection state remains as ‘none’ till the cluster is updated.

To disable Transport Settings, select **Off** as shown in Figure 3-12.

Figure 3-12 Transport Settings - Off

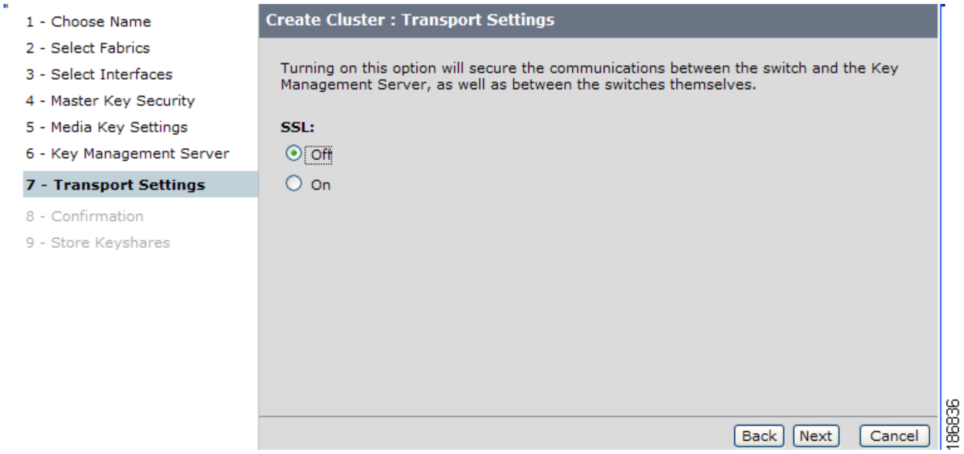

Create Cluster : Transport Settings

Turning on this option will secure the communications between the switch and the Key Management Server, as well as between the switches themselves.

SSL:

● Off
○ On

For more information on viewing or editing the transport settings in the cluster details page, see Viewing the Transport Settings in Cluster Detail Page, page 3-21

Confirming the Cluster Creation

In the Confirmation screen, review the cluster configuration information. Click **Back** to change any settings. Click **Confirm** to create the cluster.

Figure 3-13 Cluster Confirmation

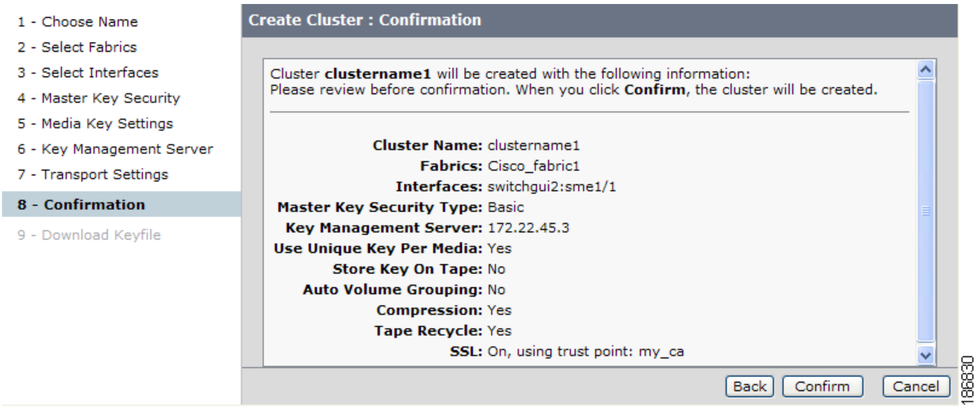

Create Cluster : Confirmation

Cluster **clustername1** will be created with the following information:
Please review before confirmation. When you click **Confirm**, the cluster will be created.

Cluster Name: clustername1
Fabrics: Cisco_fabric1
Interfaces: switchgui2:sme1/1
Master Key Security Type: Basic
Key Management Server: 172.22.45.3
Use Unique Key Per Media: Yes
Store Key On Tape: No
Auto Volume Grouping: No
Compression: Yes
Tape Recycle: Yes
SSL: On, using trust point: my_ca

You will see an indication that the operation is in progress until the entire configuration is applied.

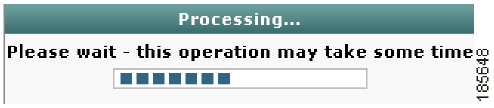

Processing...

Please wait - this operation may take some time

Downloading Key File and Storing Keyshares

This section describes the downloading of key file for basic security level and storing keyshares for the standard and advanced security level.

- Basic Security Download Key File, page 3-12
- Standard Security Confirmation and Stored Keyshares, page 3-13
- Advanced Security Confirmation and Stored Keyshares, page 3-16

Basic Security Download Key File

For the basic security level, follow these steps:

Step 1 Enter the password to encrypt the master key file. Retype the password to confirm it. Click **Download**.

Figure 3-14 Entering the Password for the Master Key File

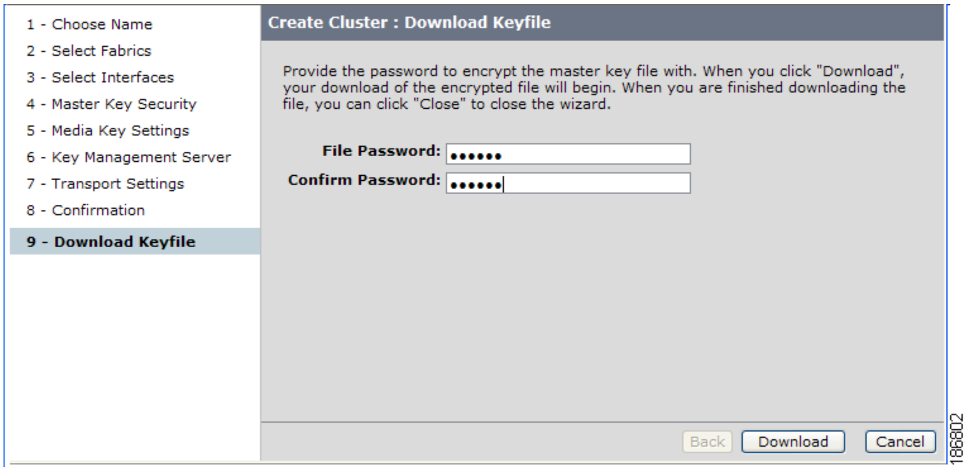

Create Cluster : Download Keyfile

Provide the password to encrypt the master key file with. When you click "Download", your download of the encrypted file will begin. When you are finished downloading the file, you can click "Close" to close the wizard.

File Password: ••••••
Confirm Password: ••••••

Step 2 A File Download screen prompts you to open or save the encrypted file.

Figure 3-15 Saving the Master Key File

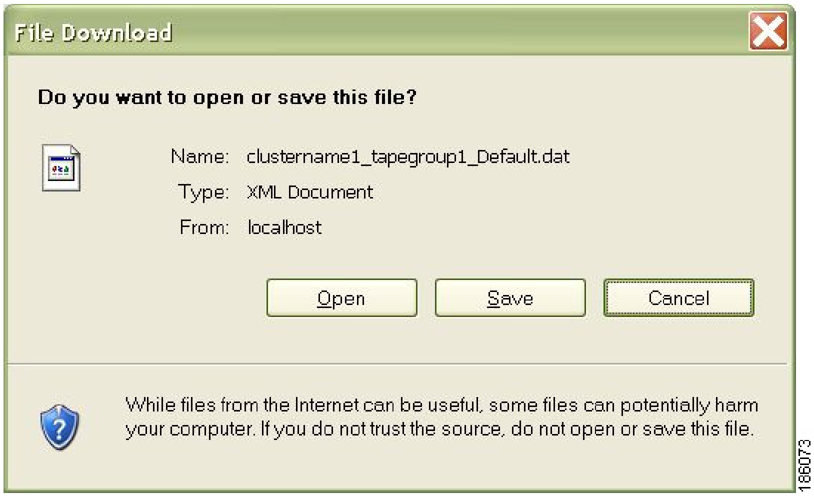

File Download

Do you want to open or save this file?

Name: clustername1_tapegroup1_Default.dat
Type: XML Document
From: localhost

While files from the Internet can be useful, some files can potentially harm your computer. If you do not trust the source, do not open or save this file.

Standard Security Confirmation and Stored Keyshares

For the standard security level, follow these steps:

Step 1 In the Confirmation screen, click **Confirm** to create the cluster.

Figure 3-16 Standard Security Confirmation

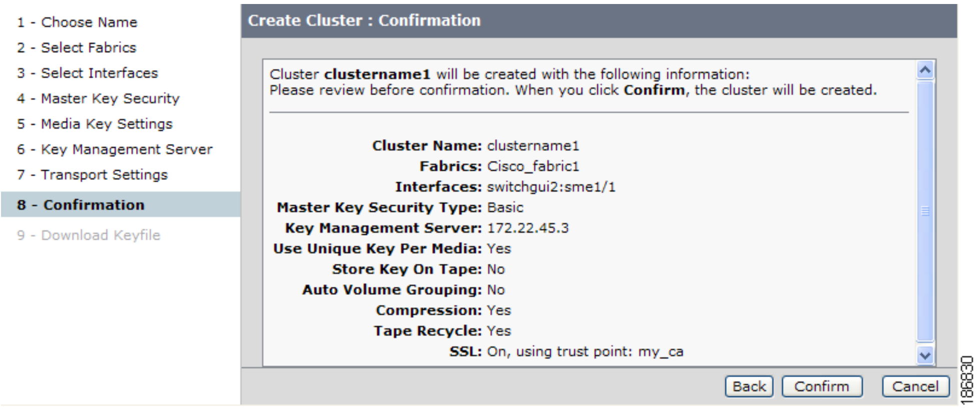

Create Cluster : Confirmation

Cluster **clustername1** will be created with the following information:
Please review before confirmation. When you click **Confirm**, the cluster will be created.

Cluster Name: clustername1
Fabrics: Cisco_fabric1
Interfaces: switchgui2:sme1/1
Master Key Security Type: Basic
Key Management Server: 172.22.45.3
Use Unique Key Per Media: Yes
Store Key On Tape: No
Auto Volume Grouping: No
Compression: Yes
Tape Recycle: Yes
SSL: On, using trust point: my_ca

Step 2 A Store Keyshares screen opens. After the smart card applet finishes loading, click **Next**.

Figure 3-17 Standard Mode - Applet loading

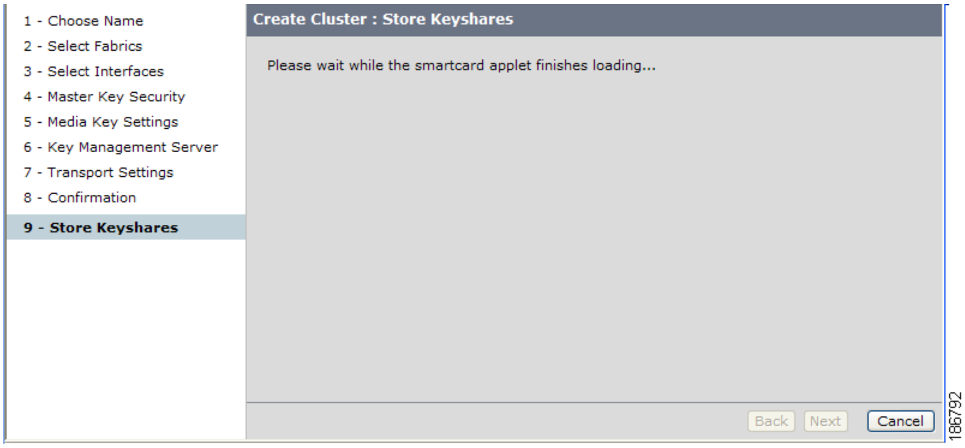

Create Cluster : Store Keyshares

Please wait while the smartcard applet finishes loading...

When entering smart card information, note the following:

- Be sure that the smart card reader is connected using the USB port (see Installing Smart Card Drivers, page 2-21 in Chapter 2, "Getting Started").
- When you insert a smart card into the reader, the card will be verified. You will be prompted to initialize the card if the card has not been previously initialized.
- Make sure that you have the appropriate smart card drivers installed before proceeding.

Step 3 Enter the switch login information (username and password used to log in to Fabric Manager), the PIN number for the smart card, and a label that will identify the smart card. The PIN number and label were defined during the smart card initialization. Click **Next**.

Figure 3-18 Entering Switch Credentials and Pin Information

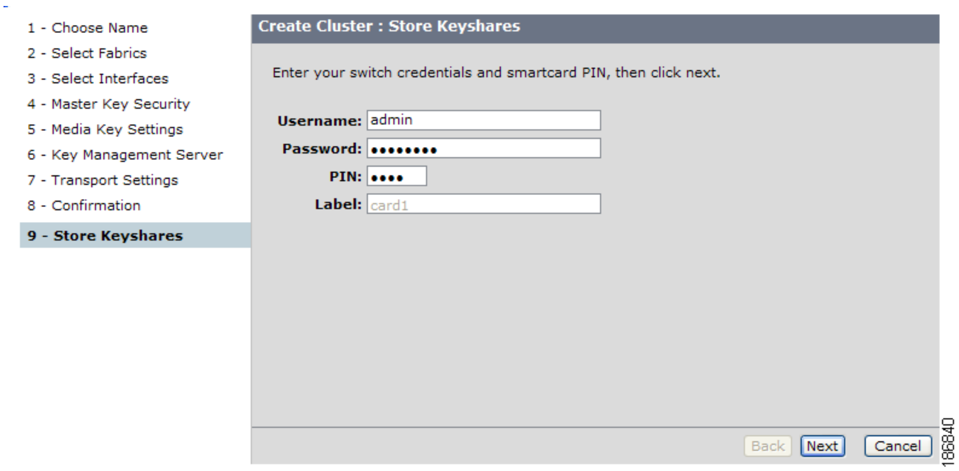

Create Cluster : Store Keyshares

Enter your switch credentials and smartcard PIN, then click next.

Username: admin
Password: ••••••••
PIN: ••••
Label: card1

Figure 3-19 Recovery Shares Successfully Stored

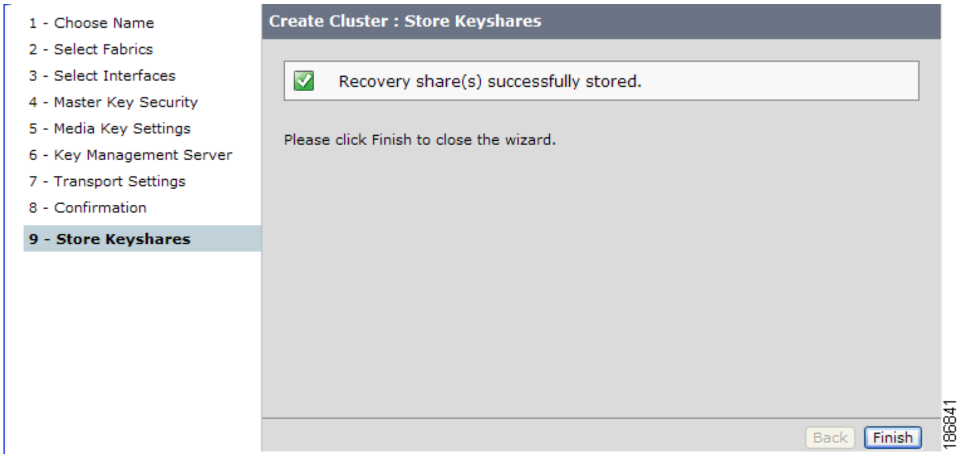

Create Cluster : Store Keyshares

☑ Recovery share(s) successfully stored.

Please click Finish to close the wizard.

Step 4 Click **Finish** to create a cluster.

Figure 3-20 Standard Security Cluster Created

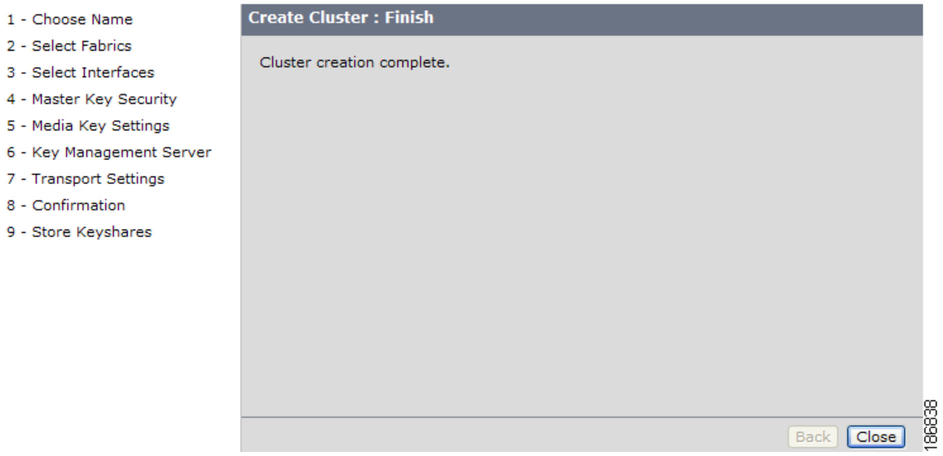

Create Cluster : Finish

Cluster creation complete.

Step 5 After the cluster creation is completed, click **Close** to return to the Fabric Manager Web Client and to view the smart card information.

Step 6 View the smart card information.

Figure 3-21 Viewing Standard Security Smart Card Information

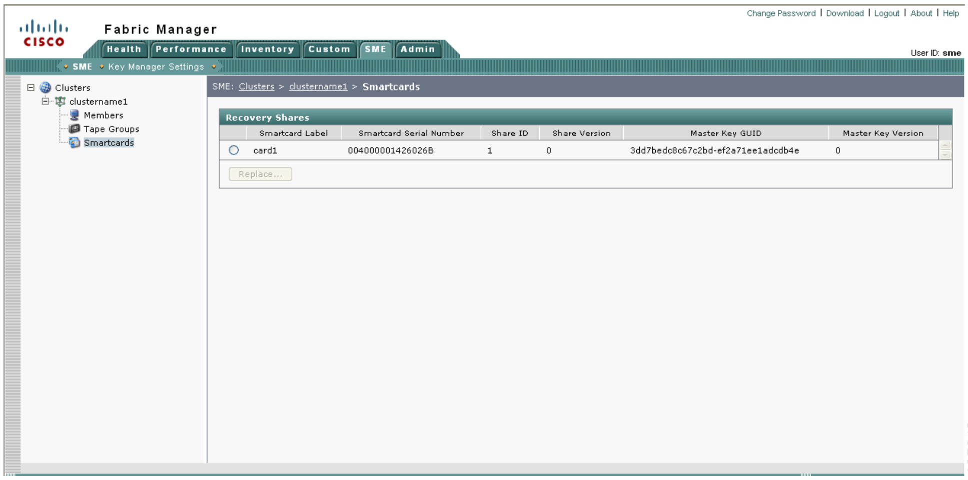

SME: Clusters > clustername1 > Smartcards

Recovery Shares

	Smartcard Label	Smartcard Serial Number	Share ID	Share Version	Master Key GUID	Master Key Version
○	card1	004000001426026B	1	0	3dd7bedc8c67c2bd-ef2a71ee1adcdb4e	0

Advanced Security Confirmation and Stored Keyshares

In the advanced security level, follow these steps:

Step 1 In the Confirmation screen, click **Confirm** to create the cluster.

Figure 3-22 Advanced Security Confirmation

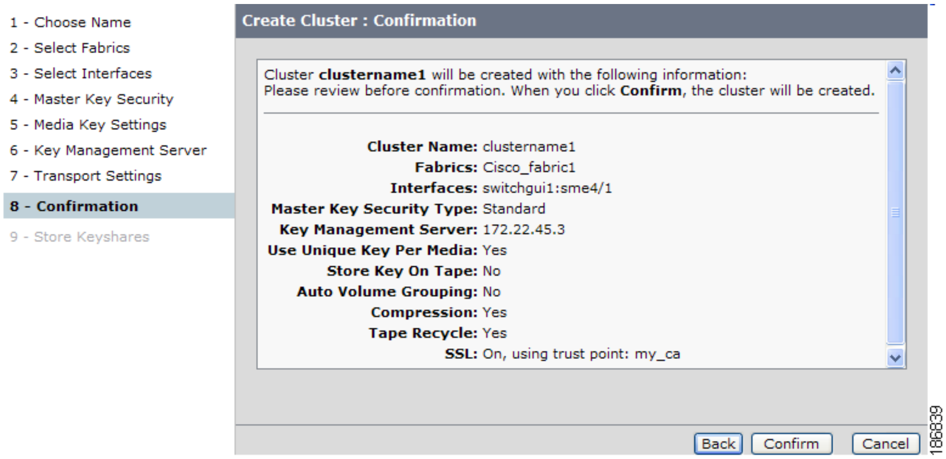

Create Cluster : Confirmation

Cluster **clustername1** will be created with the following information:
Please review before confirmation. When you click **Confirm**, the cluster will be created.

Cluster Name: clustername1
Fabrics: Cisco_fabric1
Interfaces: switchgui1:sme4/1
Master Key Security Type: Standard
Key Management Server: 172.22.45.3
Use Unique Key Per Media: Yes
Store Key On Tape: No
Auto Volume Grouping: No
Compression: Yes
Tape Recycle: Yes
SSL: On, using trust point: my_ca

Step 2 A Store Keyshares screen opens. After the smart card applet finishes loading, click **Next**.

Figure 3-23 Loading the Smart Card Applet

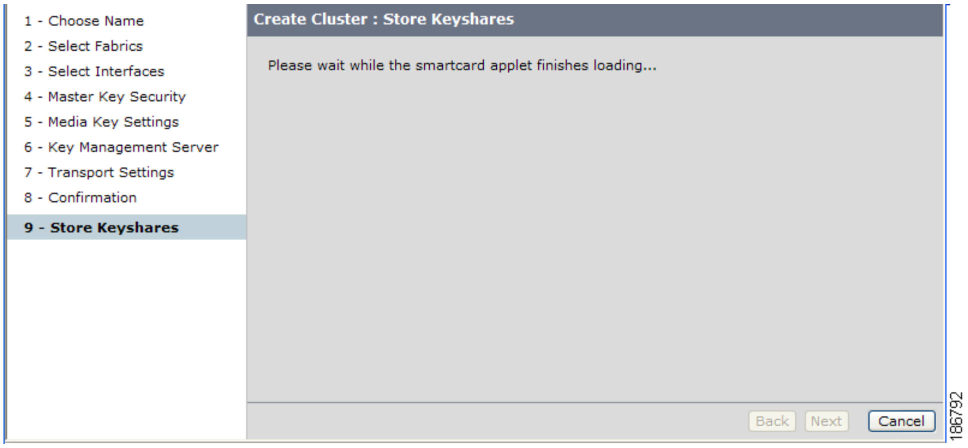

Create Cluster : Store Keyshares

Please wait while the smartcard applet finishes loading...

When entering smart card information, note the following:

- Be sure that the smart card reader is connected using the USB port (see Installing Smart Card Drivers, page 2-21 in Chapter 2, "Getting Started").
- When you insert a smart card into the reader, the card will be verified. You will be prompted to initialize the card if the card has not been previously initialized.
- For each smart card, each Cisco SME Recovery Officer must login and enter the smart card pin.
- Make sure that you have the correct smart card drivers installed before proceeding.

Step 3 Enter the switch login information (username and password used to log in to Fabric Manager), the PIN number for the smart card, and a label that will identify the smart card. The PIN number and label were defined during the smart card initialization. Click **Next**.

Figure 3-24 Entering Switch Credentials and Pin Information For the First Recovery Officer

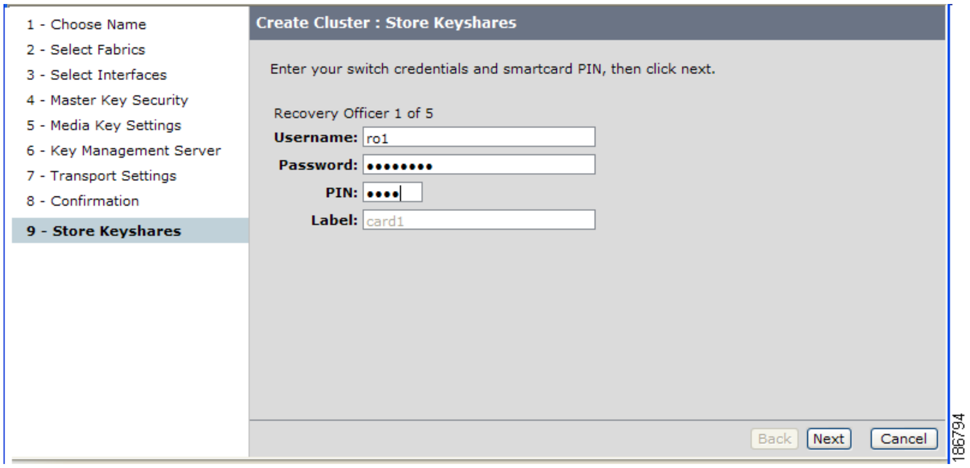

Create Cluster : Store Keyshares

Enter your switch credentials and smartcard PIN, then click next.

Recovery Officer 1 of 5

Username: ro1
Password: ••••••••
PIN: ••••
Label: card1

You will see a notification that the keyshare is being stored. This notification will be shown after each keyshare is stored.

Step 4 Click **Next**.

Figure 3-25 Storing the Keyshare For the First Recovery Officer

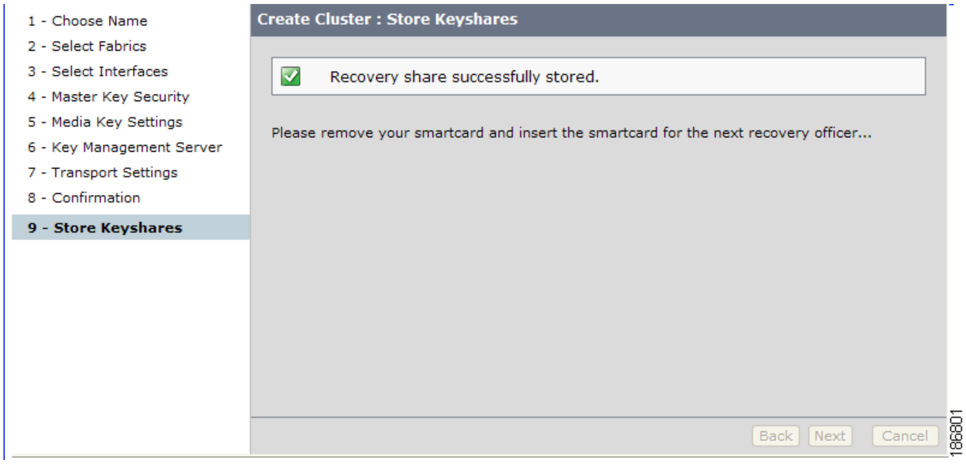

Create Cluster : Store Keyshares

☑ Recovery share successfully stored.

Please remove your smartcard and insert the smartcard for the next recovery officer...

Step 5 Enter the switch credentials and PIN information for the second recovery officer. Click **Next**.

Figure 3-26 Entering Switch Credentials and Pin Information For the Second Recovery Officer

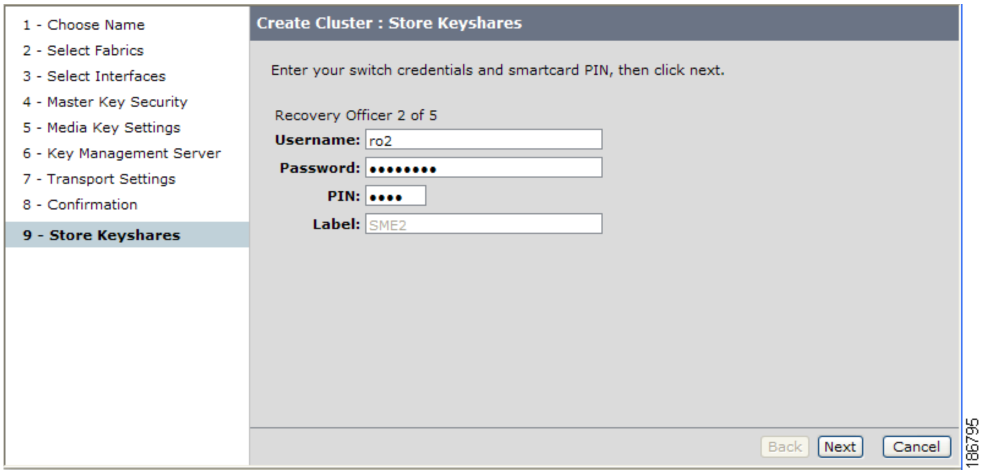

Create Cluster : Store Keyshares

Enter your switch credentials and smartcard PIN, then click next.

Recovery Officer 2 of 5

Username: ro2

Password: ••••••••

PIN: ••••

Label: SME2

Step 6 Enter the switch credentials and PIN information for the third recovery officer. Click **Next**.

Figure 3-27 Entering Switch Credentials and Pin Information For the Third Recovery Officer

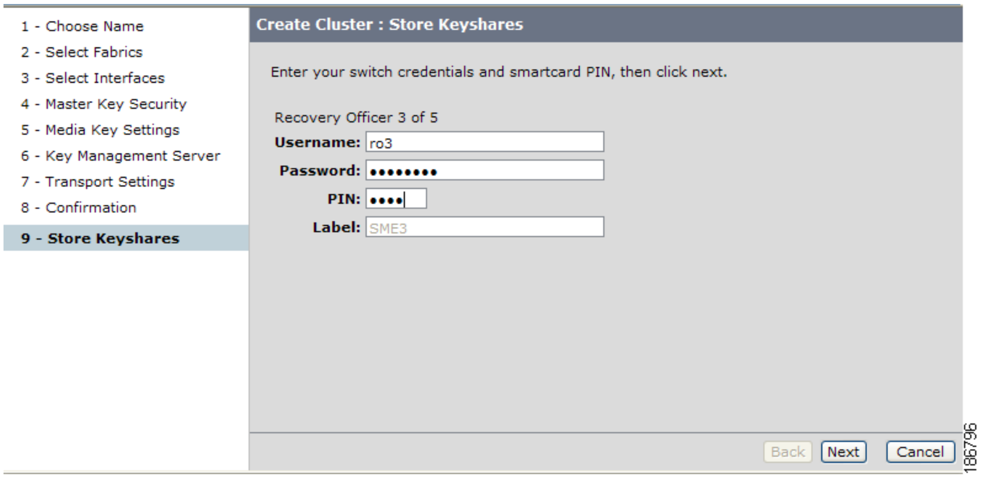

Create Cluster : Store Keyshares

Enter your switch credentials and smartcard PIN, then click next.

Recovery Officer 3 of 5

Username: ro3

Password: ••••••••

PIN: ••••

Label: SME3

Step 7 Enter the switch credentials and PIN information for the fourth recovery officer. Click **Next**.

Figure 3-28 Entering Switch Credentials and Pin Information For the Fourth Recovery Officer

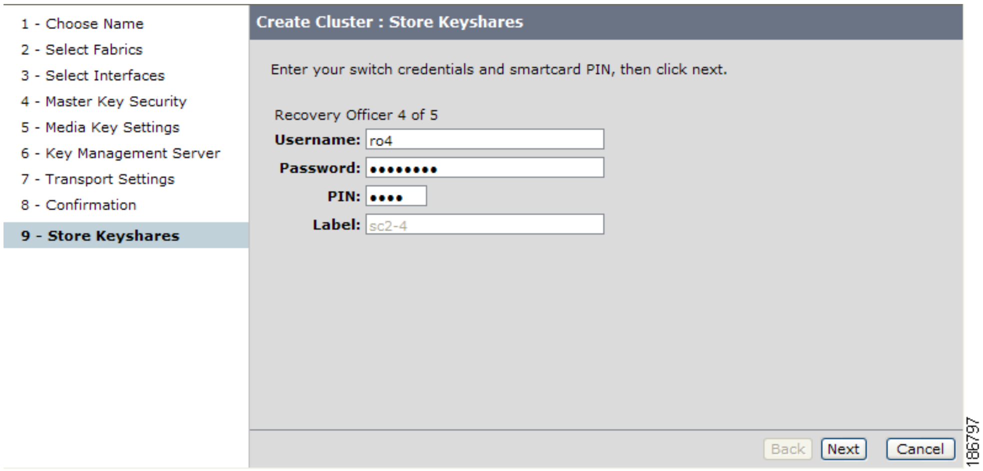

Create Cluster : Store Keyshares

Enter your switch credentials and smartcard PIN, then click next.

Recovery Officer 4 of 5

Username: ro4

Password: ••••••••

PIN: ••••

Label: sc2-4

Step 8 Enter the switch credentials and PIN information for the fifth recovery officer. Click **Next**.

Figure 3-29 Entering Switch Credentials and Pin Information For the Fifth Recovery Officer

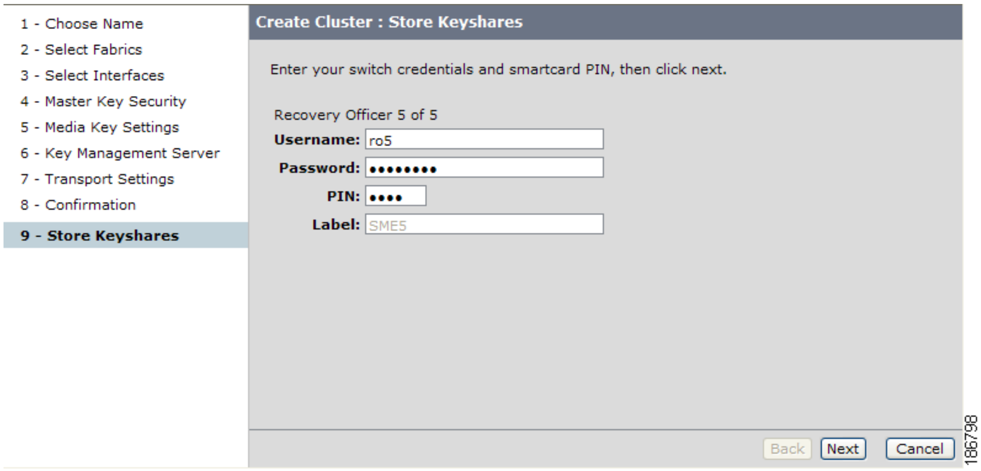

Create Cluster : Store Keyshares

Enter your switch credentials and smartcard PIN, then click next.

Recovery Officer 5 of 5

Username: ro5
Password: ••••••••
PIN: ••••
Label: SME5

Step 9 Click **Finish** to return to the Fabric Manager Web Client to view the smart card information.

Figure 3-30 Recovery Shares Successfully Stored

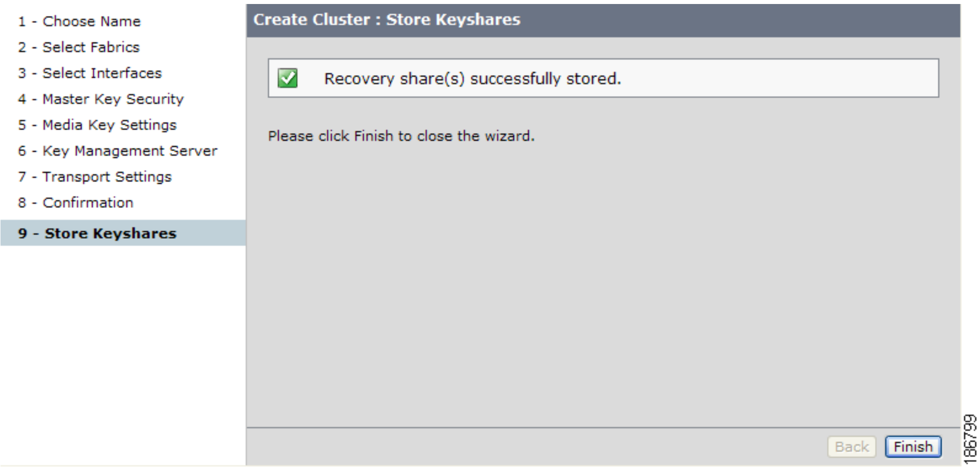

Create Cluster : Store Keyshares

✓ Recovery share(s) successfully stored.

Please click Finish to close the wizard.

Step 10 View the smart card information by selecting Smartcards.

Figure 3-31 Viewing Advanced Security Smart Card Information

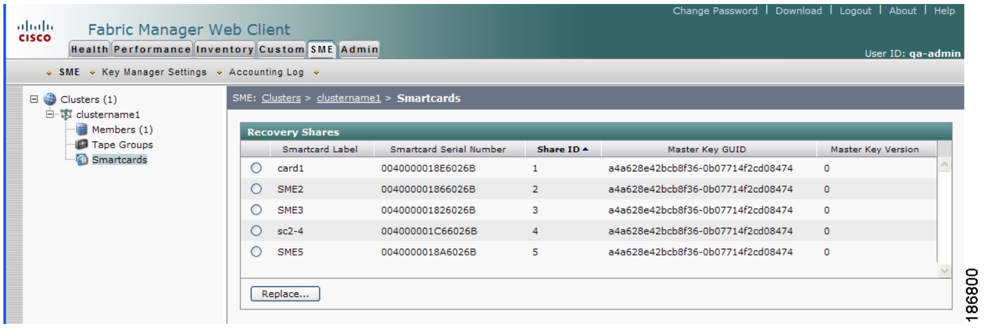
Fabric Manager Web Client
Change Password | Download | Logout | About | Help
Health Performance Inventory Custom SME Admin
User ID: qa-admin
SME Key Manager Settings Accounting Log
Clusters (1)
clustername1
Members (1)
Tape Groups
Smartcards
Replace...
186800

SME: Clusters > clustername1 > Smartcards

Recovery Shares

	Smartcard Label	Smartcard Serial Number	Share ID ▲	Master Key GUID	Master Key Version
○	card1	0040000018E6026B	1	a4a628e42bcb8f36-0b07714f2cd08474	0
○	SME2	004000001866026B	2	a4a628e42bcb8f36-0b07714f2cd08474	0
○	SME3	004000001826026B	3	a4a628e42bcb8f36-0b07714f2cd08474	0
○	sc2-4	004000001C66026B	4	a4a628e42bcb8f36-0b07714f2cd08474	0
○	SME5	0040000018A6026B	5	a4a628e42bcb8f36-0b07714f2cd08474	0

Viewing the Transport Settings in Cluster Detail Page

To view the transport settings, select the newly created cluster in the navigation pane to display the cluster detail page.

Figure 3-32 Transport Settings - SSL On

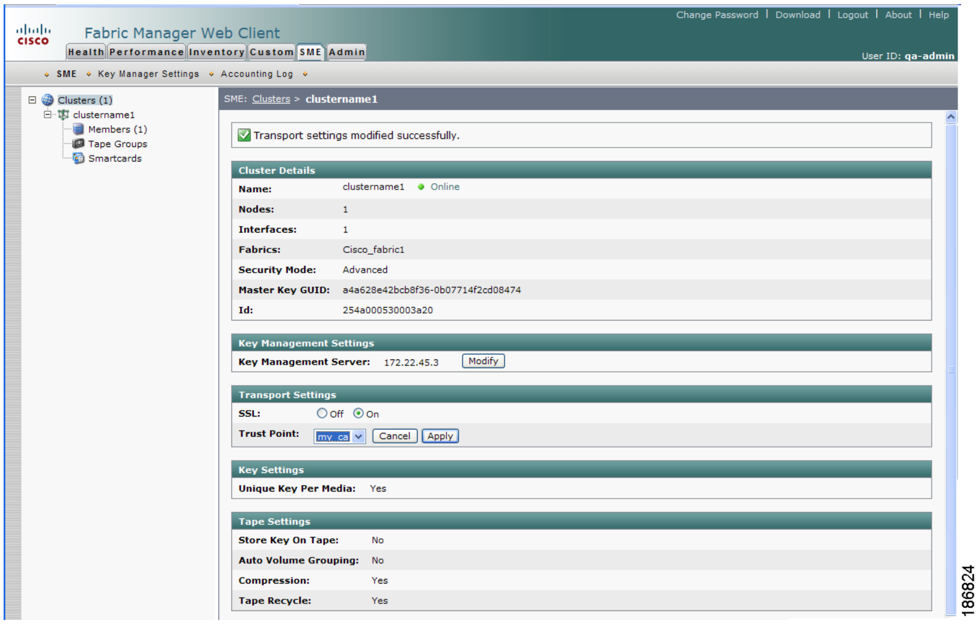
Fabric Manager Web Client
Change Password | Download | Logout | About | Help
Health Performance Inventory Custom SME Admin
User ID: qa-admin
SME Key Manager Settings Accounting Log
Clusters (1)
clustername1
Members (1)
Tape Groups
Smartcards
Modify
Cancel
Apply
186824

SME: Clusters > clustername1

☑ Transport settings modified successfully.

Cluster Details

Name:	clustername1 ● Online
Nodes:	1
Interfaces:	1
Fabrics:	Cisco_fabric1
Security Mode:	Advanced
Master Key GUID:	a4a628e42bcb8f36-0b07714f2cd08474
Id:	254a000530003a20

Key Management Settings

Key Management Server: 172.22.45.3

Transport Settings

SSL: ○ Off ● On

Trust Point: my_ca

Key Settings

Unique Key Per Media: Yes

Tape Settings

Store Key On Tape:	No
Auto Volume Grouping:	No
Compression:	Yes
Tape Recycle:	Yes

Figure 3-33 Transport Settings - SSL Off

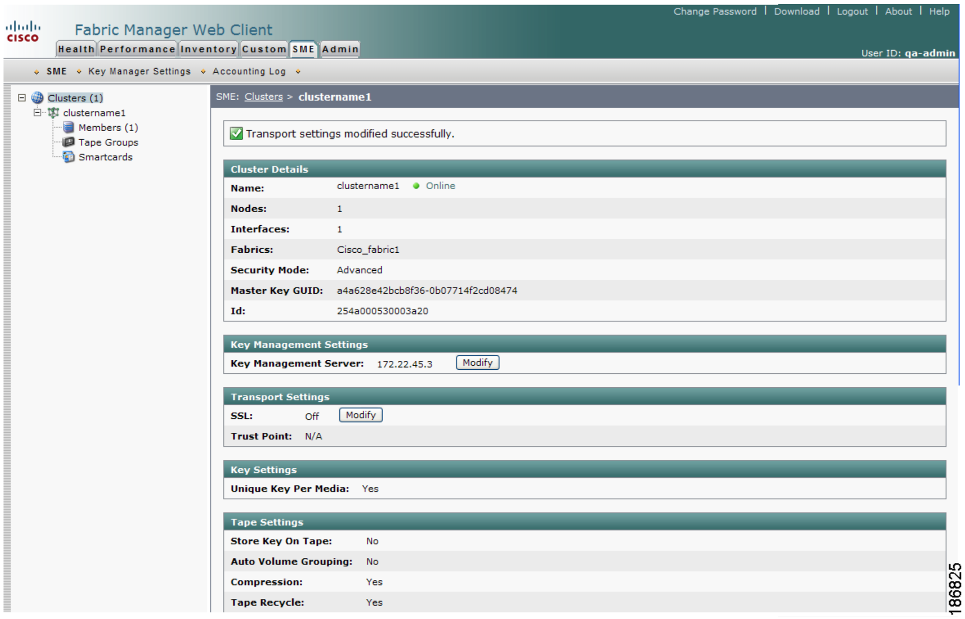
Fabric Manager Web Client
Change Password | Download | Logout | About | Help
Health Performance Inventory Custom SME Admin
User ID: qa-admin
SME Key Manager Settings Accounting Log
Clusters (1)
clustername1
Members (1)
Tape Groups
Smartcards
SME: Clusters > clustername1
Modify
186825

☑ Transport settings modified successfully.

Cluster Details

Name:	clustername1 ● Online
Nodes:	1
Interfaces:	1
Fabrics:	Cisco_fabric1
Security Mode:	Advanced
Master Key GUID:	a4a628e42bcb8f36-0b07714f2cd08474
Id:	254a000530003a20

Key Management Settings

Key Management Server: 172.22.45.3

Transport Settings

SSL: Off

Trust Point: N/A

Key Settings

Unique Key Per Media: Yes

Tape Settings

Store Key On Tape:	No
Auto Volume Grouping:	No
Compression:	Yes
Tape Recycle:	Yes

Figure 3-34 Transport Settings - SSL Modify

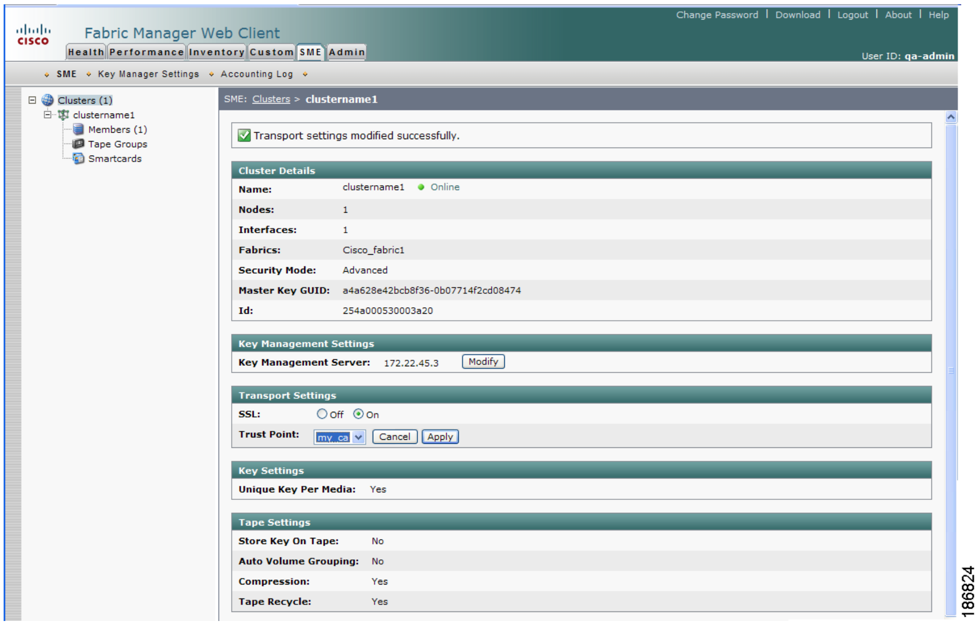
Fabric Manager Web Client
Change Password | Download | Logout | About | Help
Health Performance Inventory Custom SME Admin
User ID: qa-admin
SME Key Manager Settings Accounting Log
Clusters (1)
clustername1
Members (1)
Tape Groups
Smartcards
SME: Clusters > clustername1
Modify
Cancel
Apply
186824

☑ Transport settings modified successfully.

Cluster Details

Name:	clustername1 ● Online
Nodes:	1
Interfaces:	1
Fabrics:	Cisco_fabric1
Security Mode:	Advanced
Master Key GUID:	a4a628e42bcb8f36-0b07714f2cd08474
Id:	254a000530003a20

Key Management Settings

Key Management Server: 172.22.45.3

Transport Settings

SSL: ○ Off ● On

Trust Point: my_ca

Key Settings

Unique Key Per Media: Yes

Tape Settings

Store Key On Tape:	No
Auto Volume Grouping:	No
Compression:	Yes
Tape Recycle:	Yes

You can also modify the transport settings in the cluster detail page. Select SSL and choose a Trust Point from the drop down menu. Click **Apply** to save the settings.

Archiving and Purging a Cisco SME Cluster

You can archive clusters that are Online, Pending, or Deprecated. For information on cluster states, see Viewing Cluster States, page 3-26.

Archiving and then purging a Cisco SME cluster involves the following:

- Delete all tape groups, tape devices, and tape volume groups (see Chapter 5, "Cisco SME Tape Configuration")
- Delete all switches and Cisco SME interfaces from the cluster (see Chapter 4, "Cisco SME Interface Configuration")
- Change the cluster state to Archived (see Archiving a Cisco SME Cluster, page 3-23)
- Purge (permanently delete) a Cisco SME cluster (see Purging a Cisco SME Cluster, page 3-24)

Note You can only purge a cluster that is in the Archived state.

Archiving a Cisco SME Cluster

Archiving deletes the cluster from the switch and it retains the keys in the Cisco KMC.

To change the cluster state to Archived, follow these steps:

Step 1 Click **Clusters** in the navigation pane to display the clusters.

Step 2 Select a cluster in the information pane and click **Remove**.

Note Remove changes the cluster state to Archived and retains the keys in the Cisco KMC.

Figure 3-35 Cisco SME Clusters

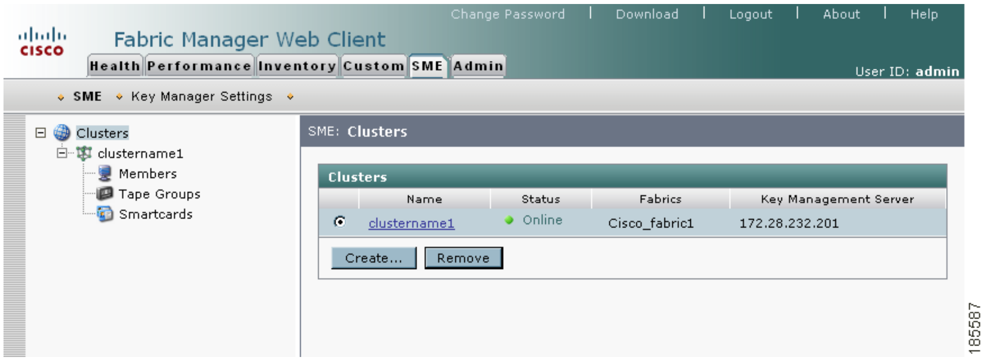

	Name	Status	Fabrics	Key Management Server
●	clustername1	Online	Cisco_fabric1	172.28.232.201

Caution Do not click Remove again unless you want to permanently delete the cluster configuration information and the master key from the Cisco KMC.

Step 3 Click **OK**.

Figure 3-36 Cluster Removal Confirmation

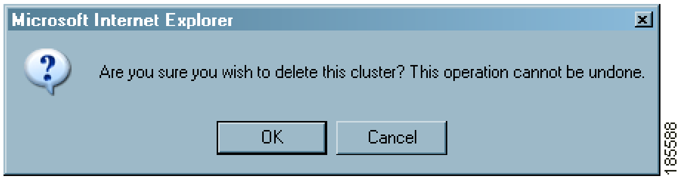

Microsoft Internet Explorer

Are you sure you wish to delete this cluster? This operation cannot be undone.

Step 4 Refresh Fabric Manager Web Client to view the notification that the cluster has been archived.

Figure 3-37 Cluster Archived

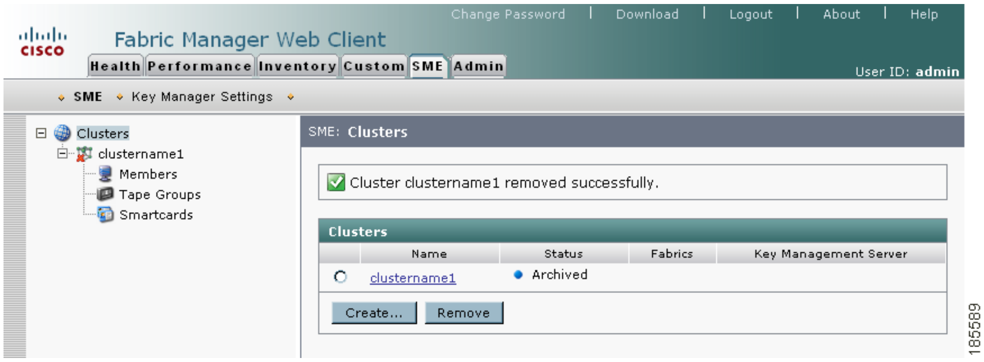

Change Password | Download | Logout | About | Help
Fabric Manager Web Client
Health Performance Inventory Custom SME Admin
User ID: admin
SME Key Manager Settings
Clusters
clustername1
Members
Tape Groups
Smartcards
SME: Clusters

Cluster clustername1 removed successfully.

Clusters

	Name	Status	Fabrics	Key Management Server
○	clustername1	● Archived		

Purging a Cisco SME Cluster

Purging a Cisco SME cluster includes the following:

- Delete all cluster elements (tape paths, tape devices, volume groups, tape groups, and switches)
- Delete (unbind) any Cisco SME interfaces that are configured in the cluster
- Change the cluster state to Archived
- Purge the cluster to remove the cluster and the master keys from the Cisco KMC.

Note Only clusters with a status of Archived can be purged.

To purge a Cisco SME cluster, follow these steps:

Step 1 Click **Clusters** in the navigation pane to display the Cisco SME clusters.

Step 2 Select an archived cluster in the information pane and click **Remove**.

Figure 3-38 Archived Cisco SME Cluster

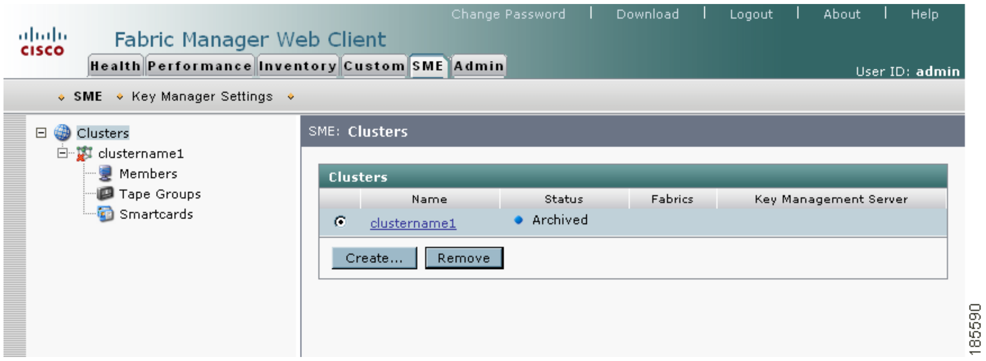

Clusters

	Name	Status	Fabrics	Key Management Server
●	clustername1	● Archived		

Step 3 Click **OK** to delete the cluster.

Caution Do not click OK unless you want to permanently delete the cluster configuration information and the master key from the Cisco KMC.

Figure 3-39 Purging Confirmation

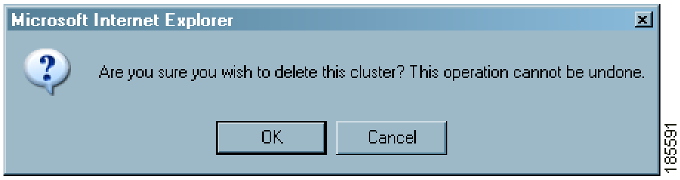

Are you sure you wish to delete this cluster? This operation cannot be undone.

Figure 3-40 Purged Cluster Confirmation

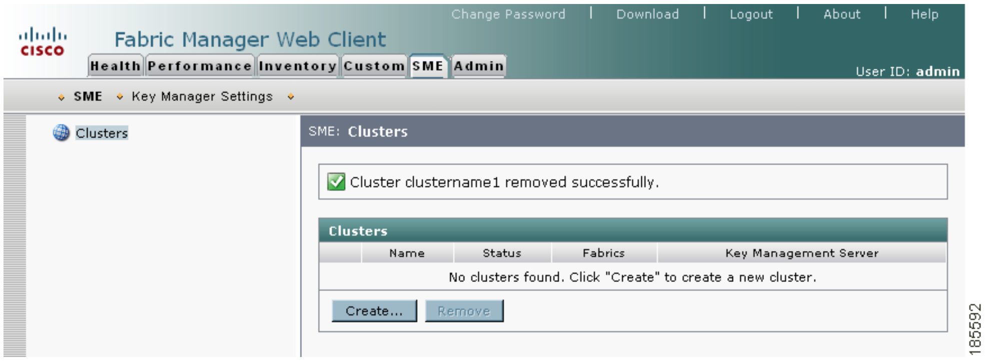

✓ Cluster clustername1 removed successfully.

Clusters

	Name	Status	Fabrics	Key Management Server
	No clusters found. Click "Create" to create a new cluster.			

Viewing Cisco SME Cluster Details

To view cluster details, click the cluster name.

Figure 3-41 Viewing Cluster Details

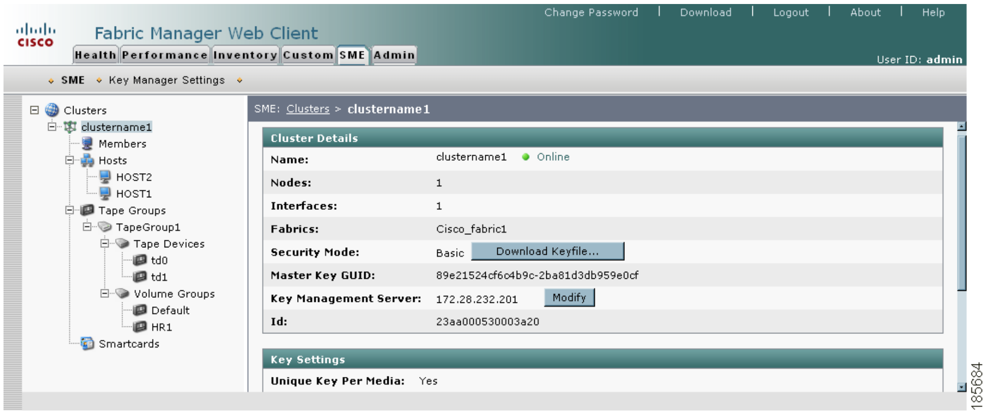

SME: Clusters > clustername1

Cluster Details

Name:	clustername1 ● Online
Nodes:	1
Interfaces:	1
Fabrics:	Cisco_fabric1
Security Mode:	Basic
Master Key GUID:	89e21524cf6c4b9c-2ba81d3db959e0cf
Key Management Server:	172.28.232.201
Id:	23aa000530003a20

Key Settings

Unique Key Per Media:	Yes

Note You can use the links across the top of the information pane to navigate within the cluster.

Viewing Cluster States

Cisco SME clusters can be in one of the following cluster states:

- Online—The Cisco SME cluster is available on the switches and is reachable from the Fabric Manager Server
- Archived—The Cisco SME cluster has been removed from the switches; however, the keys belonging to the cluster are archived in the Cisco KMC.
- Pending—The first Cisco SME interface has not been added to a cluster and it is not yet online
- Offline—The switches of the cluster are not reachable from Fabric Manager
- Deprecated—The Cisco SME cluster with all Cisco SME interfaces removed; the cluster is unusable

To view the cluster status, follow these steps:

Step 1 Click **Clusters** in the navigation pane to view a list of all Cisco SME clusters and their status.

Step 2 Alternately, click **Clusters** in the navigation pane and then click a cluster name to view the status of a specific cluster.

Figure 3-42 Viewing Cluster States

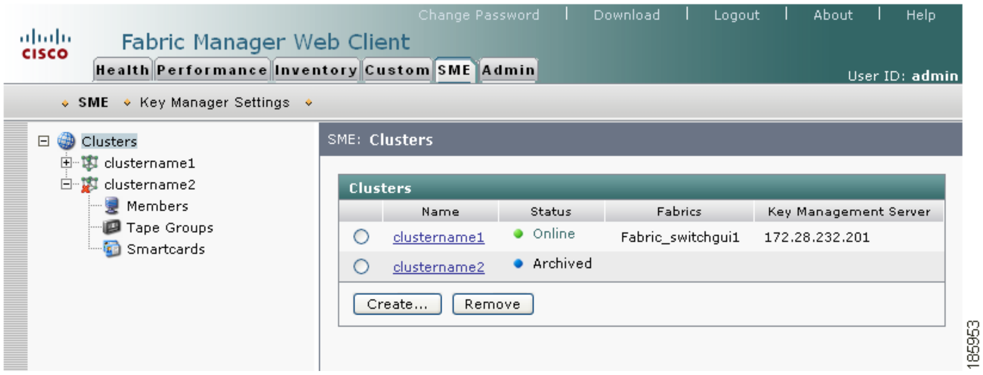

Clusters

	Name	Status	Fabrics	Key Management Server
○	clustername1	● Online	Fabric_switchgui1	172.28.232.201
○	clustername2	● Archived		

Viewing Members in a Cluster

When you view members of a cluster, you see the switches and the interfaces that have been added to a cluster.

To view the Cisco SME interfaces and switches in a cluster, follow these steps:

Step 1 Click **Members** in the navigation pane.

Figure 3-43 Viewing Members (Switches and Interfaces) in a Cluster

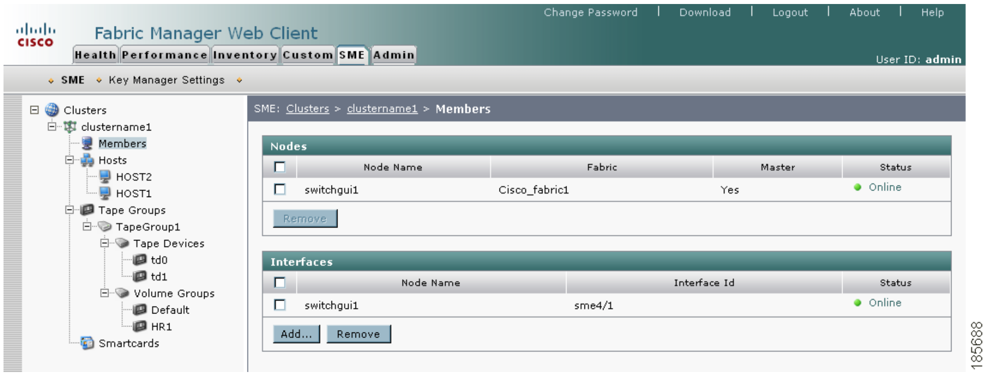

Nodes

□	Node Name	Fabric	Master	Status
□	switchgui1	Cisco_fabric1	Yes	● Online

Interfaces

□	Node Name	Interface Id	Status
□	switchgui1	sme4/1	● Online

Viewing Cluster Information Using Fabric Manager Client

To view Cisco SME cluster information using Fabric Manager Client, follow these steps:

Step 1 In the Physical Attributes pane, select **End Devices > SME Clusters**.

Figure 3-44 Physical Attributes Pane

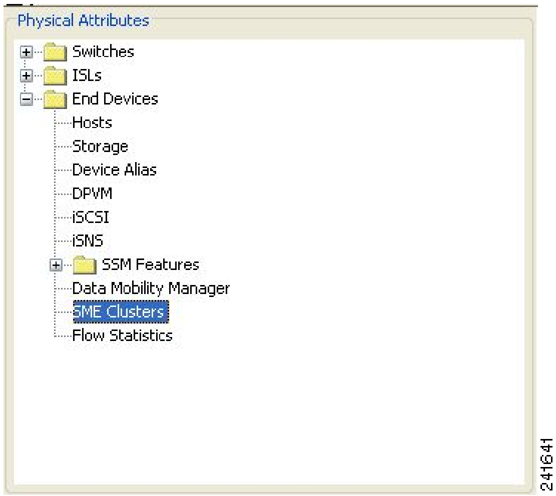

Step 2 Click the **Members** tab to view members in a cluster.

Figure 3-45 Cluster Member Information in Fabric Manager

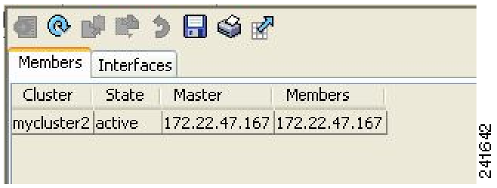

Cluster	State	Master	Members
mycluster2	active	172.22.47.167	172.22.47.167

Step 3 Click the **Interfaces** tab to view information about SME interfaces.

Figure 3-46 Cluster Interface Information in Fabric Manager

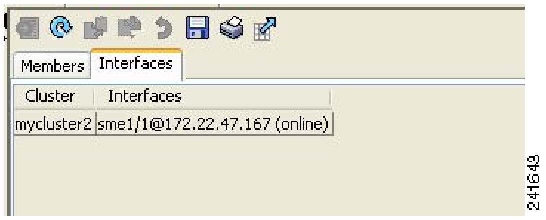

Cluster	Interfaces
mycluster2	sme1/1@172.22.47.167 (online)

Viewing Cluster Information Using Device Manager

To view Cisco SME cluster information using Device Manager, follow these steps:

Step 1 In the Interface menu, select **SME Clusters.**

Step 2 Select the **Cluster** tab to view the cluster name, state, and Master IP address.

Step 3 Select the **Members** tab to view the cluster name, switch, fabric name, and whether or not the cluster/fabric is local.

Step 4 Select **Interfaces** to view cluster interface information.

Step 5 Select **Hosts** to information about the hosts in the cluster.

Figure 3-47 Viewing Cisco SME Cluster Information in Device View

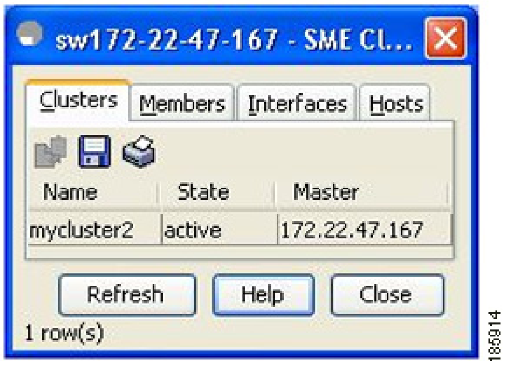

Name	State	Master
mycluster2	active	172.22.47.167

1 row(s)

Figure 3-48 Viewing Cisco SME Member Information in Device View

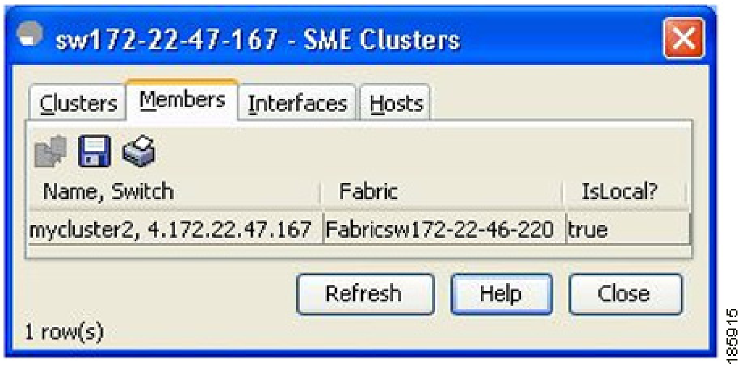

Name, Switch	Fabric	IsLocal?
mycluster2, 4.172.22.47.167	Fabricsw172-22-46-220	true

1 row(s)

Figure 3-49 Viewing Cisco SME Interface Information in Device View

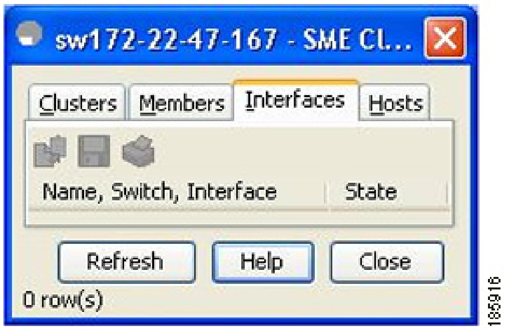

Cluster Quorum and Master Switch Election Overview

This section describes the Cisco SME cluster quorum and the process for electing the master switch in a cluster. The section includes the following:

- Cluster Quorum, page 3-30
- Master Switch Election, page 3-31

In this section, the term *switch* is used to describe a Cisco MDS 9000 Family switch that is part of a Cisco SME cluster. In addition, the following terms are used in this section.

Node ID

Every switch in a cluster has a node ID. Cisco SME assigns a node ID to every new switch as it is added to the cluster. The switch where the cluster is created is assigned the node ID of 1. This is the master switch. When a new switch is added to the cluster, it is assigned the next available higher node ID. For example, when a second switch is added to the cluster it gets the node ID of 2 and the third switch gets the node ID of 3, and so on.

Cluster View

The cluster view is the set of switches that are part of the operational cluster.

Cluster Quorum

For a cluster to be operational, it must include more than half the number of configured switches in the cluster view. Thus, in an N-node cluster, N/2 + 1 nodes form a cluster quorum.

If N is even, the cluster quorum requires N/2 nodes and additionally, the presence of the switch with the lowest node ID.

The quorum logic ensures that in the event of cluster partitions, at most one partition can be operational. All other switches are nonoperational. This guarantees the consistency of the cluster.

Master Switch Election

When a cluster is created, the switch on which the cluster is created becomes the cluster master switch. When the master switch fails or is rebooted, another switch takes over as the master switch. The master election logic uses the node ID and the latest cluster configuration to determine which switch in the cluster will become the master switch. The master election logic is describe as follows:

- If the master switch fails in an operational cluster, the switch with the next lowest node ID takes over as the master switch. Note that in an operational cluster, all the switches run the same cluster configuration.
  - When the previous master switch comes back online and joins the cluster, it does not immediately become the master.
- When all the switches of a cluster are coming up, the switch that has the latest cluster configuration becomes the master switch. If there are multiple switches with the same configuration, the switch with the lowest node ID is chosen to be the master switch.
  - Once a master switch is chosen and the cluster is operational (there is a quorum), even if a switch with a lower node ID joins the cluster at a later time, the master switch does not change.

    For example, there are 3 switches S1, S2 and S3 with node IDs 1, 2 and 3 respectively. If switches S2 and S3 form a quorum then switch S2 becomes the master switch. Even if switch S1 with the node ID of 1 comes up and joins the cluster at a later time, switch S2 continues to be the master. However, if switch S2 goes down for any reason, switch S1 will become the master switch.

Two-Switch Cluster Scenarios

According to the cluster quorum logic (see Cluster Quorum, page 3-30), a cluster with 2 configured switches can be operational if both switches are operational or the switch with the lowest node ID is operational.

In the latter case, the switch with the lowest node ID is the master of the 1-switch cluster. The other switch could have failed or simply lost connectivity to the operational switch. In either case, the switch with the higher node ID would become nonoperational. If the node with the lower node ID failed, the other switch cannot form an operational cluster.

The examples that follow describe these scenarios. The first three examples consider single switch failures.

1. Assume that in a 2-switch cluster with switches S1 (node ID 1) and S2 (node ID 2), S1 is the master (the master has the lower node ID).

   When the switches lose connectivity between them, the master switch S1 continues to be operational since it has the lower node ID and can form an (N/2) switch cluster. Switch S2 becomes non-operational.
2. Assume that in a 2-switch cluster with switches S1 (node ID 1) and S2 (node ID 2), S2 is the master (note that the master has the higher node ID because it has the latest configuration when both the switches came online).

   When the switches lose connectivity between them, switch S2 becomes non-operational and S1 takes over as the master to form a 1-switch cluster. This is consistent with the quorum logic in a 2-switch cluster (N/2 with lowest node ID).
3. Assume that in a 2-switch cluster with switches S1 (node ID 1) and S2 (node ID 2). If S1 fails (regardless of which switch was the master), S2 will also become non-operational as long as S1 is down.

When S1 comes up, S1 and S2 will form a 2-switch cluster.

The next set of examples consider reboots of both switches (S1 with node ID 1 and S2 with node ID 2).

Caution If you perform any configuration change on a cluster, you must save the running configuration to the startup configuration by entering the copy running-config startup-config CLI command on all switches before rebooting them. Otherwise, the cluster may not form correctly after the reboot (see example 6.).

4. After a reboot, if both switches S1 and S2 come up about the same time, a 2-switch cluster will be formed.
   a. If the cluster configurations are the same, S1 (with the lower node ID) will become the master.
   b. If the cluster configurations are different, the switch with the latest cluster configuration will become the master.
5. After a reboot, if switch S2 comes up first, it will not be able to form a cluster until S1 also comes up. After that, the algorithm explained in the previous case will be used.
6. After a reboot, if switch S1 comes up first, it will form a 1-switch cluster (N/2 with lowest node ID). When S2 comes up, it will join the cluster to form a 2-switch cluster.

   When S2 comes up and if it happens to have the latest cluster configuration in the startup configuration (this can happen if you did not save the running configuration to the startup configuration on S1 but did so on S2), it will not be able to join the cluster formed by S1. You may be required to follow the recovery procedures described in Chapter 9, "Cisco SME Troubleshooting" to bring S2 back into the cluster.

Caution It is critical that you save the running configuration on all switches before a reboot.

Three-Switch Cluster Scenarios

In a 3-switch cluster, the quorum requires 2 switches to be in the cluster view (N/2 + 1). The examples below explain three scenarios in a 3-switch cluster with switches S1 (node ID 1), S2 (node ID 2) and S3 (node ID 3). S1 is the master switch.

1. In a 3-switch operational cluster, if switch S3 fails or loses connectivity with the other two switches, then S3 becomes nonoperational. Switches S1 and S2 will form an operational cluster. When S3 comes up again, it will rejoin the cluster.
2. In a 3-switch operational cluster, if the master switch S1 fails or loses connectivity with the other two switches, then S1 becomes nonoperational. Switches S2 and S3 will form an operational cluster and S2 will be the master. When S1 comes up again, it will rejoin the cluster. Note that S2 will continue to be the master.
3. If two switches fail, the cluster will become nonoperational.

The examples below consider reboots on all switches in the cluster.

Caution If you perform any configuration change on a cluster, you must save the running configuration to the startup configuration by entering the copy running-config startup-config CLI command on all switches before rebooting them. Otherwise, the cluster may not form correctly after the reboot.

4. After a reboot, if all switches come up at about the same time, first a 2-switch cluster will be formed and later the third switch will be added.
   a. If the cluster configurations are the same, S1 (with the lower node ID) will become the master switch and form the 2-switch cluster first; and then add the third switch.
   b. If the cluster configurations are different, the switch that is running the latest configuration will become the master switch and then form a 2-switch cluster; and then add the third switch.
5. After a reboot, if the switches come up one at a time, a 2-switch cluster will be formed after the first two switches are up. Later, when the third switch comes online, it will join the cluster.

   If the third switch happens to be running the latest cluster configuration in the startup configuration (this can happen if you save the running configuration only on this switch but not on the other two), the third switch will not be able to join the cluster. You may be required to follow the recovery procedures described in Chapter 9, "Cisco SME Troubleshooting" to bring this switch back into the cluster.

Caution It is critical that you save the running configuration on all switches before a reboot.

Four-Switch Cluster Scenarios

The four-switch cluster scenario is very similar to the examples above. The cluster will be operational if the cluster view has at least 3 switches (N/2 + 1), or if the cluster view has 2 switches including the switch with the lowest node ID (N/2 with lowest node ID).

In-Service Software Upgrade (ISSU) In a Two-Node Cluster

In-Service Software Upgrade (ISSU) is a comprehensive, transparent software upgrade application that allows you to deploy bug fixes and add new features and services without any disruption to the traffic.

In a cluster comprising of the MDS 9222i switches as nodes, if the nodes are not able to communicate, then the node having the lowest node identifier (node ID) remains in the cluster while the other node leaves the cluster. However, when an ISSU is performed on a node having the lowest node identifier, a complete loss of the cluster results because both the nodes leave the cluster.

This undesirable situation is addressed in a two-node cluster as follows:

- The upgrading node sends a message to the other node of the intent to leave the cluster. The upgrading node can either be a master node or a slave node.
- The remaining node remains in the cluster and performs the role of the master node if it was a slave node. This node continues to remain in the cluster with the quorum intact.
- After the ISSU is completed and the switches boots up, the upgraded node rejoins the cluster as a slave node.

Note This feature is tied to the internals of ISSU logic and no additional commands need to be executed.