CHAPTER 2

Using the CFS Infrastructure

This chapter contains the following sections:

About CFS

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

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

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

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

This section includes the following topics:

Cisco MDS NX-OS Features Using CFS

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

- N Port Virtualization
- FlexAttach Virtual pWWN
- NTP
- Dynamic Port VSAN Membership
- Distributed Device Alias Services
- IVR topology
- SAN device virtualization
- TACACS+ and RADIUS
- User and administrator roles
- Port security
- iSNS
- Call Home
- Syslog
- fctimer
- SCSI flow services
- Saved startup configurations using the Fabric Startup Configuration Manager (FSCM)
- Allowed domain ID lists
- RSCN timer
- iSLB

CFS Features

CFS has the following features:

- Peer-to-peer protocol with no client-server relationship at the CFS layer.
- Three scopes of distribution.

  - Logical scope—The distribution occurs within the scope of a VSAN.
  - Physical scope—The distribution spans the entire physical topology.
  - Over a selected set of VSANs—Some applications, such as Inter-VSAN Routing (IVR), require configuration distribution over some specific VSANs. These applications can specify to CFS the set of VSANs over which to restrict the distribution.
- Three modes of distribution.
  - Coordinated distributions—Only one distribution is allowed in the fabric at any given time.
  - Uncoordinated distributions—Multiple parallel distributions are allowed in the fabric except when a coordinated distribution is in progress.
  - Unrestricted uncoordinated distributions—Multiple parallel distributions are allowed in the fabric in the presence of an existing coordinated distribution. Unrestricted uncoordinated distributions are allowed to run in parallel with all other types of distributions.
- Supports a merge protocol that facilitates the merge of application configuration during a fabric merge event (when two independent fabrics merge).

CFS Protocol

The CFS functionality is independent of the lower layer transport. Currently, in Cisco MDS switches, the CFS protocol layer resides on top of the Fiber Channel 22 (FC2) layer and is peer-to-peer with no client-server relationship. CFS uses the FC2 transport services to send information to other switches. CFS uses a proprietary SW_ILS (0x77434653) protocol for all CFS packets. CFS packets are sent to or from the switch domain controller addresses.

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

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

CFS Distribution Scopes

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

- VSAN level (logical scope)

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

- Physical topology level (physical scope)

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

- Between two switches

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

CFS Distribution Modes

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

Uncoordinated Distribution

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

Coordinated Distribution

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

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

Coordinated distribution has two variants:

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

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

Unrestricted Uncoordinated Distributions

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

Disabling CFS Distribution on a Switch

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

You can globally disable CFS on a switch,to isolate the applications using CFS from fabric-wide distributions while maintaining physical connectivity. When CFS is globally disabled on a switch, CFS operations are restricted to the switch and all CFS commands continue to function as if the switch were physically isolated.

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

Step 1 In the Physical Attributes pane, expand **Switches > CFS**.

Step 2 In the information pane, from the drop-down menu, choose **disable** or **enable** for a switch.

Step 3 Click the **Apply Changes** icon to commit the configuration changes.

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

Step 1 Choose **Admin > CFS (Cisco Fabric Services)**.

You see the CFS dialog box with the CFS status for all features on that switch.

Step 2 Uncheck or check the **Globally Enabled** check box to disable or enable CFS distribution on this switch.

Step 3 Click **Apply** to disable CFS on this switch.

Note If a switch's **global state** is disabled, it is isolated from the rest of CFS fabric. Therefore, Fabric Manager will grey out the **Config Action** column to prevent any further CFS operations.

CFS Application Requirements

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

CFS has the following requirements:

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

Enabling CFS for an Application

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

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

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

To enable CFS for a feature using Fabric Manager, follow these steps:

Step 1 Choose a feature on which to enable CFS. For example, expand **Switches > Events**, and then select **CallHome** in the Physical Attributes pane. The Information pane shows that feature with a CFS tab. Click the **CFS** tab to display the CFS state for each switch in the fabric for that feature.

Step 2 Decide on which switch(es) to enable CFS. Set the Admin column to either **enable** to enable CFS or **disable** to disable CFS.

Note Enable CFS for all switches in the fabric or VSAN for the feature that uses CFS.

Fabric Manager retrieves the status of the CFS change and updates the Last Command and Result columns.

To enable CFS for a feature using Device Manager, follow these steps:

Step 1 Choose **Admin > CFS (Cisco Fabric Services)**.

You see the CFS dialog box with the CFS status for all features on that switch.

Step 2 Decide which features need CFS. Set the Command column to either **enable** to enable CFS or **disable** to disable CFS.

Note Enable or disable CFS for all switches in the fabric or VSAN for the feature that uses CFS.

Locking the Fabric

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

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

Committing Changes

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

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

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

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

You can commit changes for a specified feature by setting CFS > Config Action to **commit** for that feature.

To commit changes using Fabric Manager for CFS-enabled features, follow these steps:

Step 1 Choose the feature you want to enable CFS for. For example, choose **Switch > Clock > NTP**

The Information pane shows that feature, with a CFS tab.

Step 2 Click the **CFS** tab to display the CFS state for each switch in the fabric for that feature.

Step 3 In the **Feature** tab, click on the NTP **General** tab and change any configuration for example. Click the **Apply Changes** icon to apply the configuration to the local switch. The change will remain at the pending database at the local switch till further CFS commit is applied.

Step 4 Click **Pending Differences** to check the configuration of this feature on this switch as compared to other switches in the fabric or VSAN that have CFS enabled for this feature (see Figure 2-1).

Figure 2-1 CFS Pending Difference

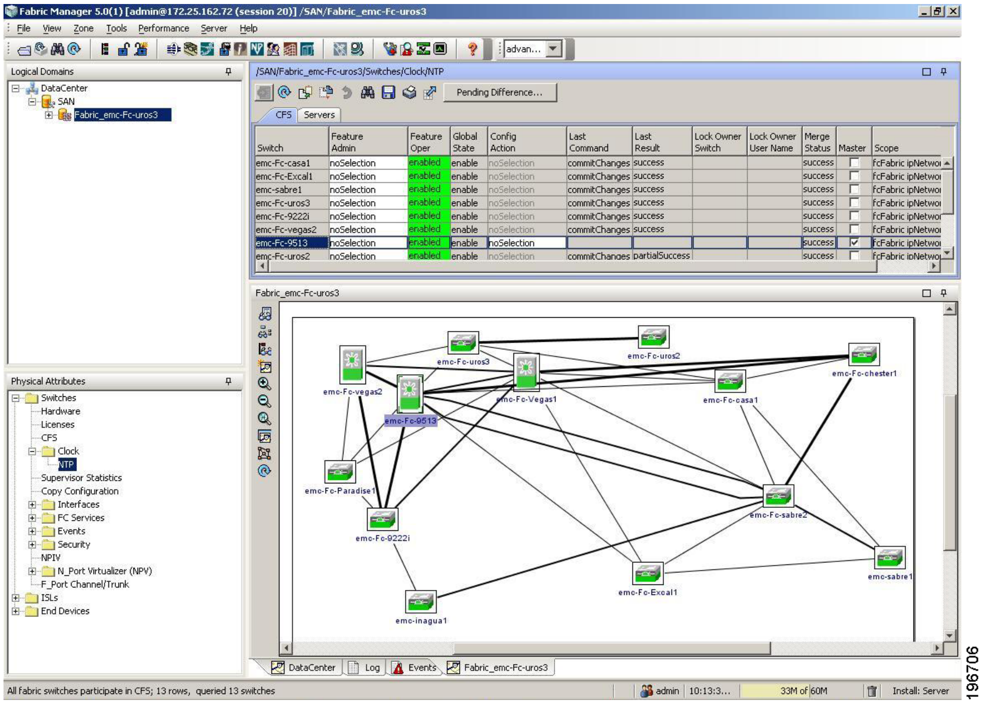

Fabric Manager 5.0(1) [admin@172.25.162.72 (session 20)] /SAN/Fabric_emc-Fc-uros3

/SAN/Fabric_emc-Fc-uros3/Switches/Clock/NTP

Switch	Feature Admin	Feature Oper	Global State	Config Action	Last Command	Last Result	Lock Owner Switch	Lock Owner User Name	Merge Status	Master	Scope
emc-Fc-casa1	noSelection	enabled	enable	noSelection	commitChanges	success			success	☐	fcFabric ipNetwo
emc-Fc-Excal1	noSelection	enabled	enable	noSelection	commitChanges	success			success	☐	fcFabric ipNetwo
emc-sabre1	noSelection	enabled	enable	noSelection	commitChanges	success			success	☐	fcFabric ipNetwo
emc-Fc-uros3	noSelection	enabled	enable	noSelection	commitChanges	success			success	☐	fcFabric ipNetwo
emc-Fc-9222i	noSelection	enabled	enable	noSelection	commitChanges	success			success	☐	fcFabric ipNetwo
emc-Fc-vegas2	noSelection	enabled	enable	noSelection	commitChanges	success			success	☐	fcFabric ipNetwo
emc-Fc-9513	noSelection	enabled	enable	noSelection					success	☑	fcFabric ipNetwo
emc-Fc-uros2	noSelection	enabled	enable	noSelection	commitChanges	partialSuccess			success	☐	fcFabric ipNetwo

All fabric switches participate in CFS; 13 rows, queried 13 switches

Figure 2-2 CFS Pending Difference Screen

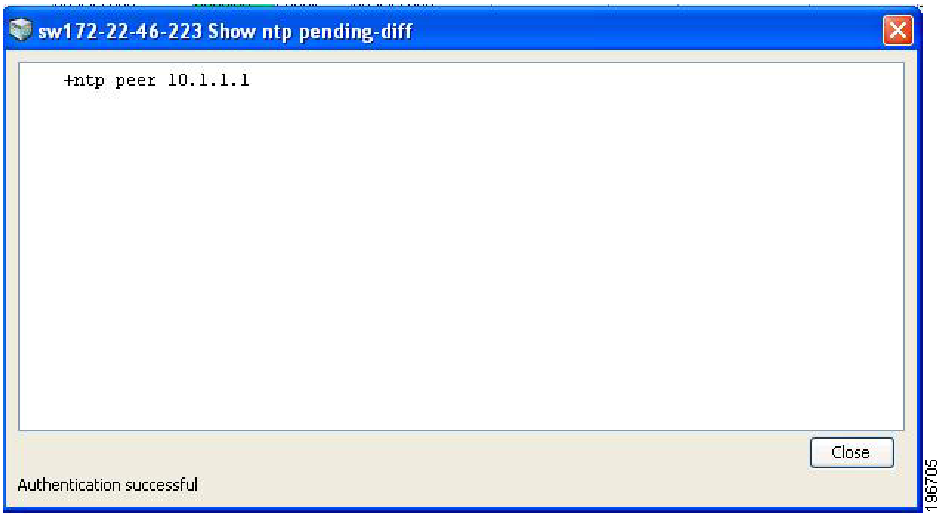

sw172-22-46-223 Show ntp pending-diff

```
+ntp peer 10.1.1.1
```

Authentication successful

Step 5 Go back to the **CFS** tab, right-click the value in the Config Action column in the master switch selected (see Figure 2-1) and select an option from the drop-down menu. (commit, clear lock, abort). For example, right-click the value in the Config Action column and select **commit** to commit the CFS pending changes for that feature and distribute the changes through CFS.

Fabric Manager retrieves the status of the CFS change and updates the Last Command and Last Result columns for the feature or VSAN.

To commit changes using Device Manager for CFS-enabled features, follow these steps:

Step 1 Choose the feature you want to enable CFS for in Device Manager. For example, choose **Admin > NTP**

Step 2 In the **Feature** tab, click on the NTP **General** tab and change any configuration for example. Click the **Apply Changes** icon to apply the configuration to the local switch. The change will remain at the pending database at the local switch till further CFS commit is applied.

Step 3 Choose **Admin** > **CFS** (Cisco Fabric Services)

Step 4 In the CFS table, click **Pending Differences** to check the configuration of this feature on this switch as compared to other switches in the fabric or VSAN that have CFS enabled for this feature.

Step 5 For each applicable feature, set the Command column to **commit** to commit the configuration changes for that feature and distribute the c hanges through CFS, or set it to **abort** to discard the changes for that feature and release the fabric lock for CFS for that feature.

Device Manager retrieves the status of the CFS change and updates the Last Command and Result columns.

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

Discarding Changes

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

You can discard changes for a specified feature by setting theCommand column value to **disable** for that feature then clicking **Apply**.

Saving the Configuration

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

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

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

Clearing a Locked Session

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

To clear locks using Fabric Manager, follow these steps:

Step 1 Click the **CFS** tab.

Step 2 Select **clearLock** from the Config Action drop-down list for each switch that you want to clear the lock (see Figure 2-3).

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

Figure 2-3 Clearing Locks

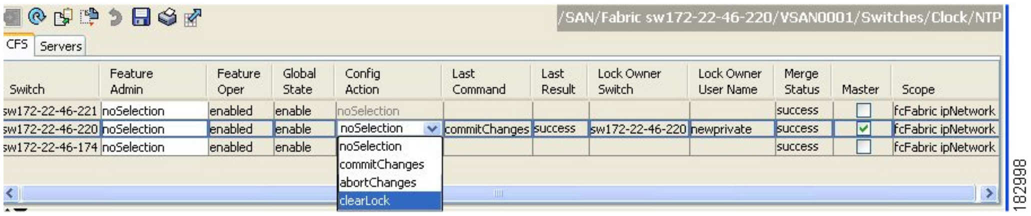

/SAN/Fabric sw172-22-46-220/VSAN0001/Switches/Clock/NTP

Switch	Feature Admin	Feature Oper	Global State	Config Action	Last Command	Last Result	Lock Owner Switch	Lock Owner User Name	Merge Status	Master	Scope
sw172-22-46-221	noSelection	enabled	enable	noSelection					success	☐	fcFabric ipNetwork
sw172-22-46-220	noSelection	enabled	enable	noSelection	commitChanges	success	sw172-22-46-220	newprivate	success	☑	fcFabric ipNetwork
sw172-22-46-174	noSelection	enabled	enable						success	☐	fcFabric ipNetwork

Caution Exercise caution when using this function to clear locks in the fabric. Any pending configurations in any switch in the fabric is flushed and lost.

CFS Merge Support

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

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

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

In case of a successful merge, the merged database is distributed to all switches in the combined fabric and the entire new fabric remains in a consistent state. You can recover from a merge failure by starting a distribution from any of the switches in the new fabric. This distribution restores all peers in the fabric to the same configuration database.

Displaying CFS Configuration Information

To display the status of CFS distribution on the switch using Device Manager, follow these steps:

Step 1 Choose **Admin > CFS (Cisco Fabric Services)**.

You see the CFS dialog box. This dialog box displays the distribution status of each feature using CFS, which currently registered applications are using CFS, and the result of the last successful merge attempt.

Step 2 Select a row and click **Details** to view more information about the feature.

CFS Distribution over IP

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

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

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

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

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

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

Figure 2-4 shows a network with both Fibre Channel and IP connections. Node A forwards an event to node B over Fibre Channel. Node B forwards the event node C and node D using unicast IP. Node C forwards the event to node E using Fibre Channel.

Figure 2-4 Network Example 1 with Fibre Channel and IP Connections

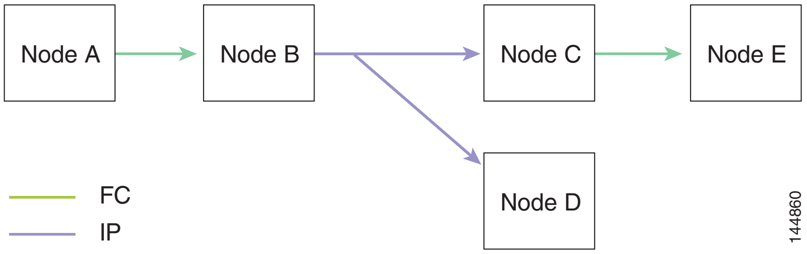

Figure 2-5 is the same as Figure 2-4 except that node D and node E are connected using Fibre Channel. All processes is the same in this example because node B has node C and node D the distribution list for IP. Node C does not forward to node D because node D is already in the distribution list from node B.

Figure 2-5 Network Example 2 with Fibre Channel and IP Connections

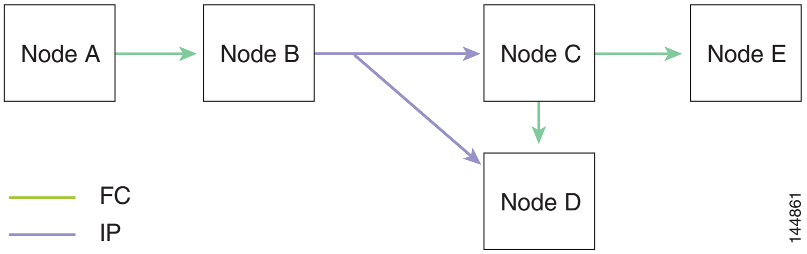

Figure 2-6 is the same as Figure 2-5 except that node D and node E are connected using IP. Both node C and node D forward the event to E because the node E is not in the distribution list from node B.

Figure 2-6 Network Example 3 with Fibre Channel and IP Connections

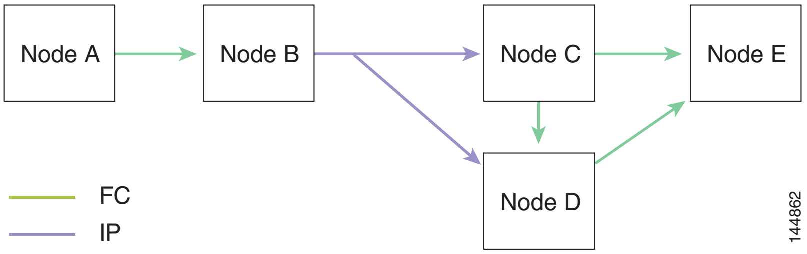

Configuring Static IP Peers for CFS over IP

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

To enable CFS over IP on the switches that do not support multicast forwarding, multicast forwarding has to be enabled on the Ethernet IP switches all along the network that physically connects the switch. In such cases, you can configure static IP peers for CFS distribution over IP.

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

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

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

Cisco MDS Fabric Manager discovers NPV devices by reading the name server database on the NPV core switch, which is also used to manage the static peer list at an NPV switch for CFS distribution over IP using static peers.

Fabric Manager 4.1(1) and later provides a one-time configuration wizard to manage the peer list of the discovered NPV peers on a switch. When the peer list is configured on a switch, CFS enables distribution using the IP static peers on all members of the list and propagates the peer list to all members on the list.

Note If a new NPV switch is added to the fabric, you must launch the NPV CFS Setup wizard to update the list, because Fabric Manager does not update the list automatically.

Adding Peers to List

To configure the static IP peers list using Fabric Manager, follow these steps:

Step 1 From the Fabric Manager menu, select **Tools** > **Other** > **NPV CFS Setup**.

Figure 2-7 NPV CFS Setup Menu

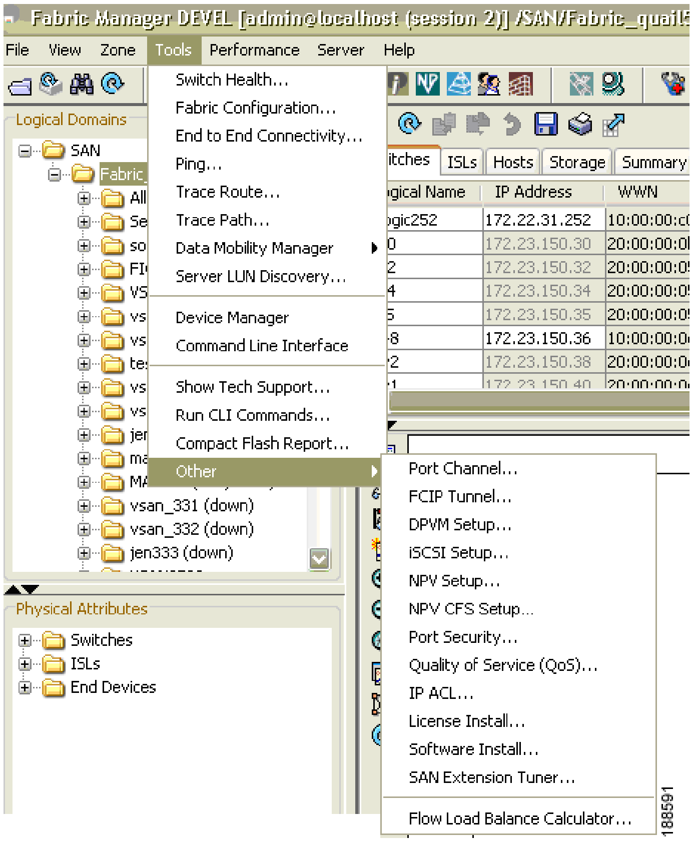

The NPV Device Selection dialog box is displayed with the list of NPV device peers retrieved from the switch including the device name, device IP address, and the status of the peer.

Figure 2-8 NPV Device Selection

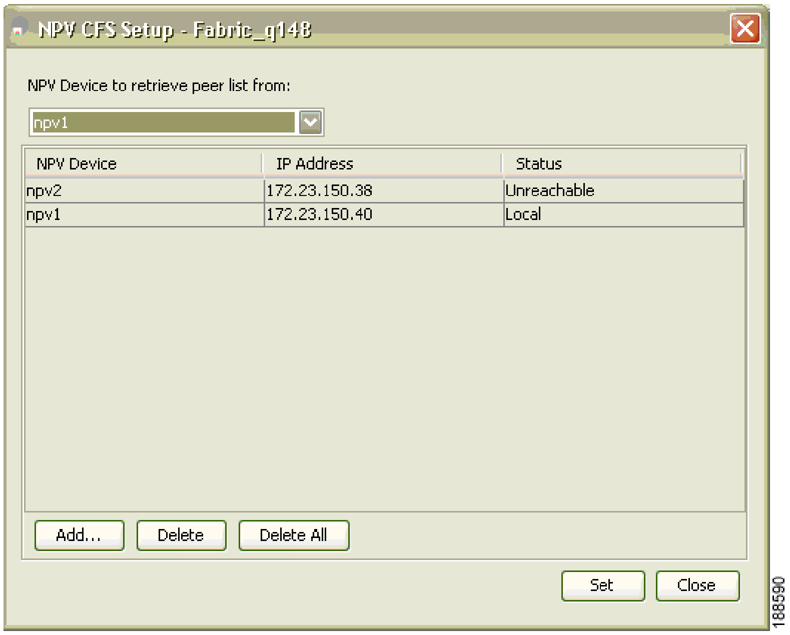

NPV Device to retrieve peer list from:

npv1

NPV Device	IP Address	Status
npv2	172.23.150.38	Unreachable
npv1	172.23.150.40	Local

Step 2 From the **NPV Device to retrieve peer list from** drop-down list box, select the device to retrieve the peer list from.

If the NPV device in the list retrieved from the switch is present in the fabric, then one of the following statuses is displayed: Local, Reachable, Unreachable, or Discovery in Progress. If the NPV device is not present in the fabric, then the status is displayed as Not in Fabric.

Note If the status is displayed as Not in Fabric, you must remove the device from the list.

Step 3 Click **Add**.

The following dialog box is displayed with the list of all the NPV devices in the fabric that are not included in the current peer list. By default, all the switches in the list are selected.

Figure 2-9 Peer Selection

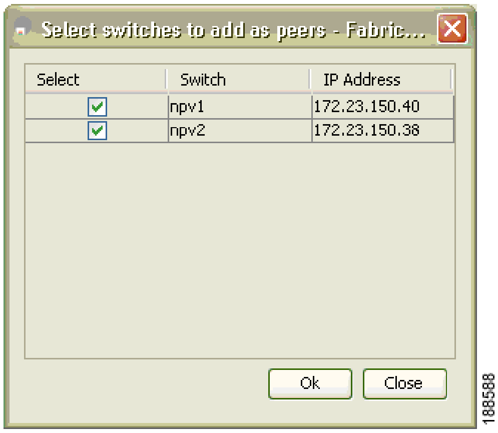

Select	Switch	IP Address
☑	npv1	172.23.150.40
☑	npv2	172.23.150.38

Step 4 Select the peers, and then click **Ok** to add the peers to the list.

The peers are added to the list with To Be Added status.

Figure 2-10 Confirm Peer Selection

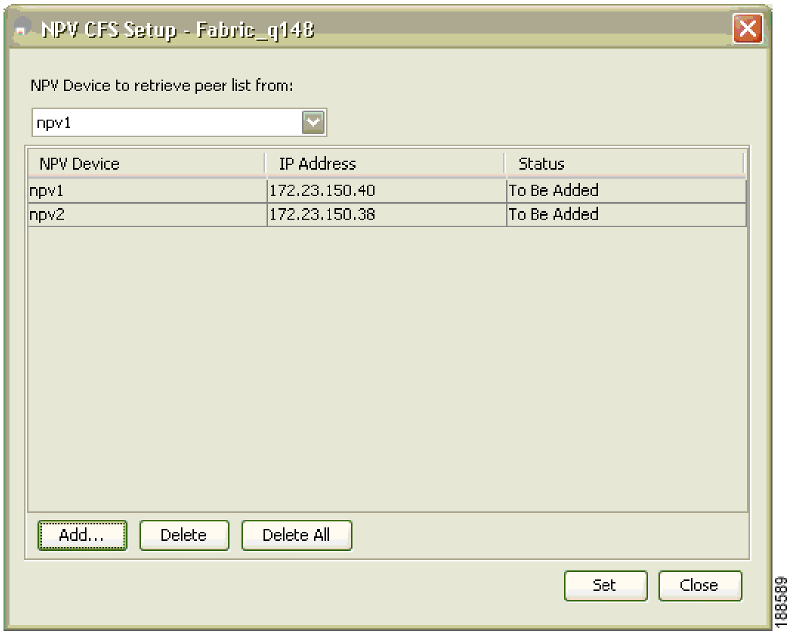

NPV Device to retrieve peer list from: npv1

NPV Device	IP Address	Status
npv1	172.23.150.40	To Be Added
npv2	172.23.150.38	To Be Added

Step 5 Click **Set** to confirm adding the peers to the list and start the peers list propagation by CFS.

Removing an NPV Device from the Peer List

To delete a peer from the IP peer list using Fabric Manager, follow these steps:

Step 1 From Fabric Manager menu, select **Tools** > **Other** > **NPV CFS Setup**.

The NPV CFS Setup wizard is launched.

Step 2 From the **NPV Device** to retrieve peer list from drop-down list box, select the device to retrieve the peer list from which you want to delete a peer.

Step 3 Do one of the following tasks to mark the peer or local host as deleted:

- To delete a peer from the peer list, select the peer from the list, and then click **Delete**.
- To delete the local host from the peer list, select the local NPV device and click **Delete**, or select all the peers in the list, and then click **Delete All**.

Step 4 Click **Yes** to delete the peer from the list.

Step 5 Click **Set** in the NPV CFS wizard. The following message box is displayed:

Figure 2-11 Start Dynamic Peer Discovery

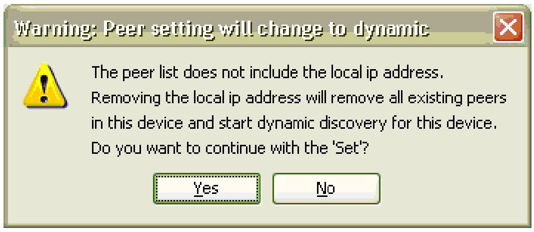

Warning: Peer setting will change to dynamic

The peer list does not include the local ip address. Removing the local ip address will remove all existing peers in this device and start dynamic discovery for this device. Do you want to continue with the 'Set'?

Step 6 Click **Yes** to remove the deleted peer or local host from all the other NPV device peer lists, and start dynamic peer discovery using multicast in the deleted peer.

CFS Regions

This section contains the following topics:

- About CFS Regions, page 2-17
- Managing CFS Regions, page 2-19
- Creating CFS Regions, page 2-19
- Assigning Features to CFS Regions, page 2-20
- Moving a Feature to a Different Region, page 2-21
- Removing a Feature from a Region, page 2-21
- Deleting CFS Regions, page 2-22

About CFS Regions

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

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

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

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

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

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

Managing CFS Regions

This section describes how to use Fabric Manager for managing CFS regions. Fabric Manager provides a comprehensive view of all the switches, regions, and the features associated with each region in the topology. To complete the following tasks, use the tables under the All Regions and Feature by Region tabs:

- Creating CFS Regions, page 2-19
- Assigning Features to CFS Regions, page 2-20
- Moving a Feature to a Different Region, page 2-21
- Removing a Feature from a Region, page 2-21
- Deleting CFS Regions, page 2-22

Creating CFS Regions

To create a CFSregion using Fabric Manager, follow these steps:

Step 1 Expand the **Switches** folder in the **Physical Attributes** pane and click **CFS**.

The information pane displays the Global, IP Multicast, Feature by Region, and All Regions tabs.

Step 2 Click the **All Regions** tab.

The tab displays a list of Switches and RegionIds.

Step 3 Click the **Create Row** button on the toolbar.

Figure 2-12 shows the Create a Region dialog box.

Figure 2-12 Create a Region Dialog Box

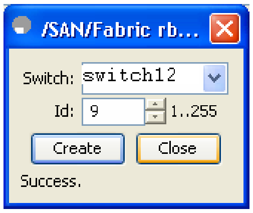

/SAN/Fabric rb...

Switch: switch12

Id: 9 1..255

Success.

Step 4 From the drop-down list, select the switch and choose a RegionId from the range.

Step 5 Click **Create**.

Upon successful creation of the region, Success is displayed at the bottom of the dialog box.

Note CFS always works within individual fabric when no CFS region is applied, or within individual CFS region when CFS region exists. Therefore, even when a SAN or DataCenter (higher than fabric) node or scope is selected, Fabric Manager will only show the switches for the first fabric under the selected scope.

Assigning Features to CFS Regions

To assign a feature to a region using Fabric Manager, follow these steps:

Step 1 Expand the Switches folder in the Physical Attributes pane and click **CFS**.

The information pane displays the Global, IP Multicast, Feature by Region, and All Regions tabs.

Step 2 Click the **Feature by Region** tab.

This tab lists all the switches along with their corresponding Feature and RegionId.

Step 3 Click the **Create Row** button on the toolbar.

Figure 2-13 shows the Assign a Feature dialog box.

Figure 2-13 Assign a Feature Dialog Box

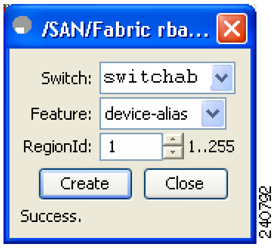

/SAN/Fabric rba...

Switch: switchab

Feature: device-alias

RegionId: 1 1..255

Success.

Step 4 From the drop-down box, select a switch.

The features running on the selected switch are listed in the Feature drop-down list.

Step 5 Select a feature on that switch to associate a region.

Step 6 From the RegionID list, select the region number to associate a region with the selected feature.

Step 7 Click **Create** to complete assignment of a switch feature to the region.

Upon successful assignment of feature, “Success” is displayed at the bottom of the dialog box.

When a feature is assigned to a new region using the **Feature by Region** tab, a new row with the new region is created automatically in the table under the **All Regions** tab. Alternatively, you can create a region using the **All Regions** tab.

Note In the **Feature by Region** tab, when you try to reassign a feature on a switch to another region by clicking **Create Row**, an operation failed message is shown. The error message states that an entry already exists. However, moving a feature to a different region is a different task and it is described in the next section.

Moving a Feature to a Different Region

Before moving a feature to a new region, create the new region in the All Regions tab. That is, a new row has to be added in the All Regions tab with the new Region ID.

To move a feature to a different region using Fabric Manager, follow these steps:

Step 1 Expand the Switches folder in the Physical Attributes pane and select CFS.

The information pane displays the Global, IP Multicast, Feature by Region, and All Regions tabs.

Step 2 Click the **Feature by Region** tab.

Figure 2-14 shows the Feature by Region tab, which lists all the switches along with their feature and region details.

Figure 2-14 Feature by Region Tab

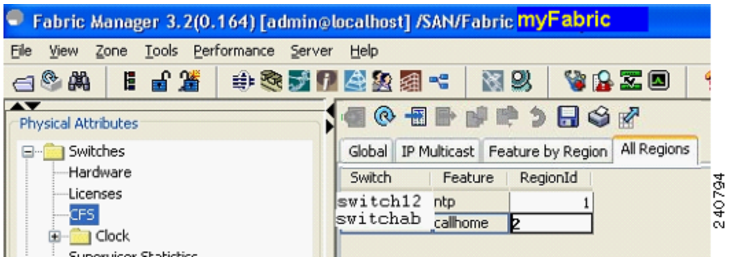

Switch	Feature	RegionId
switch12	ntp	1
switchab	callhome	2

Step 3 Double-click the RegionId cell in the required row.

The cursor blinks in the cell prompting a change in the value.

Step 4 Change the RegionId value to the required region.

Step 5 Click the **Apply Changes** button on the tool bar to commit the change.

Removing a Feature from a Region

To remove a feature from a region using Fabric Manager, follow these steps:

Step 1 Click the **Feature by Region** tab and select the required row.

Step 2 Click the **Delete Row** button on the toolbar.

Figure 2-15 shows a confirmation dialog box.

Figure 2-15 Removing from a Region

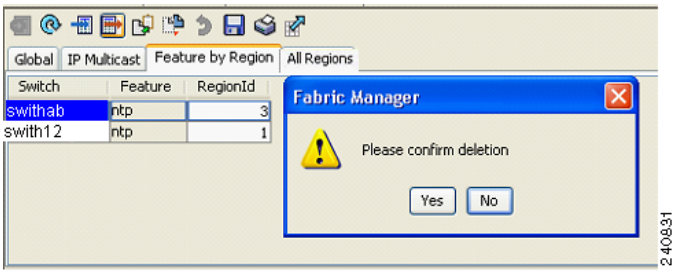

Switch	Feature	RegionId
swithab	ntp	3
swith12	ntp	1

Please confirm deletion

Step 3 Click **Yes** to confirm row deletion from the table in view.

Deleting CFS Regions

To delete an entire region, follow these steps

Step 1 Click the **All Regions** tab and select the required row.

Step 2 Click **Delete Row**.

This action removes all entries pertaining to that switch and region in the table under Feature by Region tab.

Figure 2-16 shows a confirmation dialog box.

Figure 2-16 Deleting CFS Regions

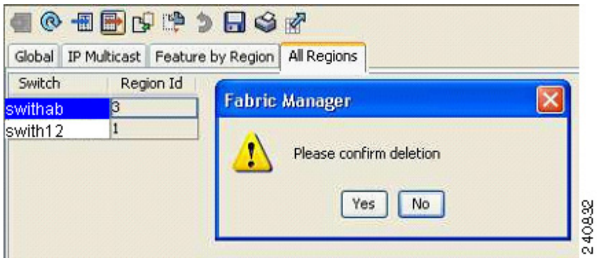

Switch	Region Id
swithab	3
swith12	1

Please confirm deletion

Step 3 Click **Yes** to confirm deletion of the region.

CFS Example Using Fabric Manager

This procedure is an example of what you see when you use Fabric Manager to configure a feature that uses CFS.

Step 1 Select the CFS-capable feature you want to configure. For example, expand a **VSAN,** and then 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 **CFS** tab.

You see the CFS configuration and status for each switch (see Figure 2-17).

Figure 2-17 CFS Configuration

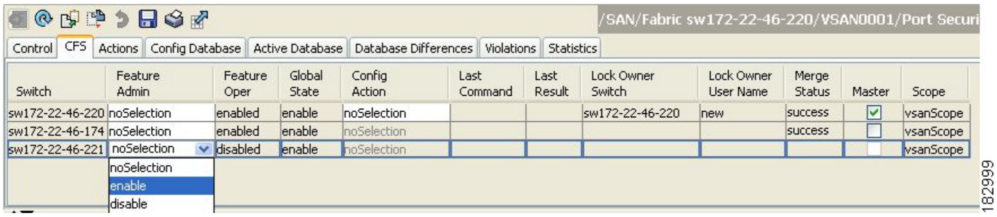

Switch	Feature Admin	Feature Oper	Global State	Config Action	Last Command	Last Result	Lock Owner Switch	Lock Owner User Name	Merge Status	Master	Scope
sw172-22-46-220	noSelection	enabled	enable	noSelection			sw172-22-46-220	new	success	☑	vsanScope
sw172-22-46-174	noSelection	enabled	enable	noSelection					success	□	vsanScope
sw172-22-46-221	noSelection noSelection enable disable	disabled	enable	noSelection						□	vsanScope

Step 3 From the Feature Admin drop-down list, select **enable** for each switch.

Step 4 Repeat step 3 for all switches in the fabric.

Note A warning is displayed if you do not enable CFS for all switches in the fabric for this feature.

Step 5 Check the **Master** check box for the switch to act as the merge master for this feature.

Note If you click any other tab in the information pane and then click the CFS tab, the Master check box will no longer be checked. Fabric Manager does not cache the CFS Master information.

Step 6 From the Config Action drop-down list, select **commit Changes** for each switch that you enabled for CFS.

Step 7 Click the **Servers tab** in the Information pane.

You see the configuration for this feature based on the master switch (see Figure 2-18).

Step 8 Modify the feature configuration. For example, right-click the name in the Master column and select **Create Row** to create a server for NTP.

a. Set the ID and the Name or IP Address for the NTP server.

b. Set the **Mode** radio button and optionally check the **Preferred** check box.

c. Click **Create** to add the server.

Figure 2-18 Servers Tab

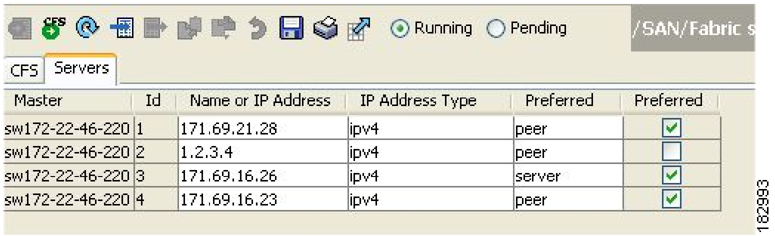

Master	Id	Name or IP Address	IP Address Type	Preferred	Preferred
sw172-22-46-220	1	171.69.21.28	ipv4	peer	☑
sw172-22-46-220	2	1.2.3.4	ipv4	peer	☐
sw172-22-46-220	3	171.69.16.26	ipv4	server	☑
sw172-22-46-220	4	171.69.16.23	ipv4	peer	☑

Step 9 Click the **Delete Row** icon to delete a row.

If you make any changes, the status automatically changes to **Pending** (see Figure 2-19).

Figure 2-19 Status Change to Pending

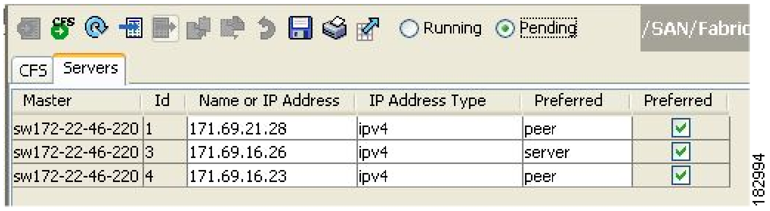

Master	Id	Name or IP Address	IP Address Type	Preferred	Preferred
sw172-22-46-220	1	171.69.21.28	ipv4	peer	☑
sw172-22-46-220	3	171.69.16.26	ipv4	server	☑
sw172-22-46-220	4	171.69.16.23	ipv4	peer	☑

Step 10 Click the **Commit CFS Pending Changes** icon to save the changes (see Figure 2-20).

Figure 2-20 Commit CFS Pending Changes

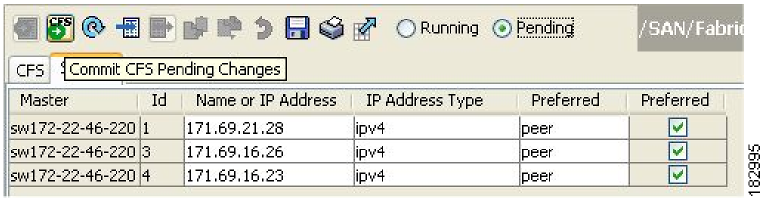

Master	Id	Name or IP Address	IP Address Type	Preferred	Preferred
sw172-22-46-220	1	171.69.21.28	ipv4	peer	☑
sw172-22-46-220	3	171.69.16.26	ipv4	peer	☑
sw172-22-46-220	4	171.69.16.23	ipv4	peer	☑

Step 11 The status changes to **Running** (see Figure 2-21).

Figure 2-21 Status Change to Running

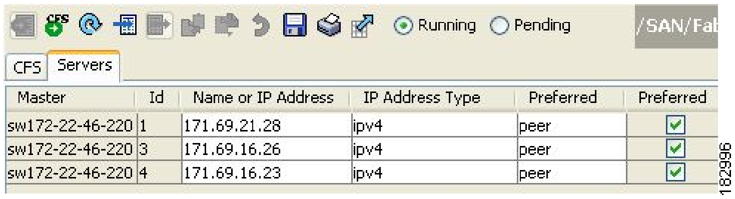

Master	Id	Name or IP Address	IP Address Type	Preferred	Preferred
sw172-22-46-220	1	171.69.21.28	ipv4	peer	☑
sw172-22-46-220	3	171.69.16.26	ipv4	peer	☑
sw172-22-46-220	4	171.69.16.23	ipv4	peer	☑

Step 12 From the Config Action drop-down list, select **abortChanges** for each switch that you enabled for CFS (see Figure 2-22).

Figure 2-22 Commit Configuration Changes

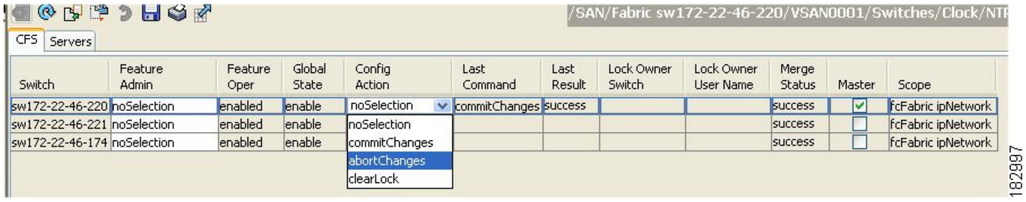

Switch	Feature Admin	Feature Oper	Global State	Config Action	Last Command	Last Result	Lock Owner Switch	Lock Owner User Name	Merge Status	Master	Scope
sw172-22-46-220	noSelection	enabled	enable	noSelection	commitChanges	success			success	☑	fcFabric ipNetwork
sw172-22-46-221	noSelection	enabled	enable	noSelection commitChanges abortChanges clearLock					success	☐	fcFabric ipNetwork
sw172-22-46-174	noSelection	enabled	enable						success	☐	fcFabric ipNetwork

Note Fabric Manager does not change the status to pending if **enable** is selected, because the pending status does not apply until the first actual change is made.

Step 13 Click the **Apply Changes** icon to commit the configuration changes for that feature and distribute the changes through CFS.

Note When using CFS with features such as DPVM and device alias, you must select **commit** at the end of each configuration. If the session is locked, you must exit the feature by selecting **abort**.

To configure the master or seed switch for distribution for each feature using Fabric Manager, follow these steps:

Step 1 Choose the feature that needs a merge master for CFS. For example, expand **Switches,** expand **Events** and select **CallHome** from the Physical Attributes pane.

The Information pane shows that feature including a CFS tab.

Step 2 Click the **CFS** tab to display the CFS state for each switch in the fabric for that feature.

Step 3 Check the Master column check box for the switch to act as the merge master for this feature.

Step 4 Click the **Apply Changes** icon to select this switch as master for future CFS distributions.

.

CFS Example Using Device Manager

This procedure is an example of what you see when you use Device Manager to configure a feature that uses CFS. For specific procedures for features that use CFS, refer to that feature’s documentation.

To configure a feature that uses CFS using Device Manager, follow these steps:

Step 1 Open the dialog box for any CFS-capable feature. Device Manager checks to see whether CFS is enabled. It also checks to see if there is a lock on the feature by checking for at least one entry in the Owner table. If CFS is enabled and there is a lock, Device Manager sets the status to “pending” for that feature. You see a dialog box displaying the lock information.

Step 2 Click **Continue** or **Cancel** when prompted. If you continue, Device Manager remembers the CFS status.

Step 3 Choose **Admin > CFS (Cisco Fabric Services)** to view the user name of the CFS lock holder.

Step 4 Click the locked feature and click **Details**.

Step 5 Click the **Owners** tab and look in the **UserName** column.

Note Device Manager does not monitor the status of the feature across the fabric until you click **Refresh**. If a user on another CFS-enabled switch attempts to configure the same feature, they do not see the "pending" status. However, their configuration changes are rejected by your switch.

Step 6 If CFS is enabled and there is no lock, Device Manager sets the status to running for that feature.

You then see a dialog box for the feature. As soon as you perform a creation, deletion, or modification, Device Manager changes the status to pending and displays the updated information from the pending database.

Step 7 View the CFS table for a feature. Device Manager only changes the status to running when **commit**, **clear,** or **abort** is selected and applied. Device Manager will not change the status to "pending" if **enable** is selected, because the pending status does not apply until the first actual change is made.

The **Last Command** and **Result** fields are blank if the last command is **noOp**.

Note When using CFS with features like DPVM and device alias, you must select **commit** at the end of each configuration. If the session is locked, you must exit the feature by selecting **abort**.

Default Settings

Table 2-1 lists the default settings for CFS configurations.

Table 2-1 Default CFS Parameters

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