



# CHAPTER 13

## Using the CFS Infrastructure

---

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

This chapter contains the following sections:

- [Information About CFS section, page 13-1](#)
- [Guidelines and Limitations section, page 13-9](#)
- [Default Settings section, page 13-9](#)
- [Configuring CFS section, page 13-9](#)
- [Configuring CFS Regions section, page 13-18](#)
- [Verifying CFS Configurations section, page 13-22](#)
- [Additional References section, page 51-21](#)
- [Field Descriptions for CFS section, page 13-28](#)
- [Monitoring Logs section, page 51-20](#)
- [Feature History for CFS section, page 13-31](#)

## Information About CFS

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

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

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

This section includes the following topics:

- [Cisco MDS NX-OS Features Using CFS section, page 13-2](#)
- [CFS Features section, page 13-2](#)

- [Enabling CFS for an Application section, page 13-3](#)
- [CFS Protocol section, page 13-3](#)
- [CFS Distribution Scopes section, page 13-3](#)
- [CFS Distribution Modes section, page 13-4](#)
- [CFS Connectivity in a Mixed Fabric section, page 13-5](#)
- [Locking the Fabric section, page 13-5](#)
- [Committing Changes section, page 13-5](#)
- [CFS Merge Support section, page 13-6](#)
- [CFS Distribution over IP section, page 13-6](#)
- [Static IP Peers for CFS over IP section, page 13-7](#)
- [About CFS Regions section, page 13-8](#)

## Cisco MDS NX-OS Features Using CFS

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

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

## CFS Features

CFS has the following features:

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

## Enabling CFS for an Application

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

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

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

## CFS Protocol

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

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

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

## CFS Distribution Scopes

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

- VSAN level (logical scope)

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

- Physical topology level (physical scope)

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

- Between selected 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.

## CFS Connectivity in a Mixed Fabric

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

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

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

**Note**

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

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

## Locking the Fabric

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

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

For information on verifying CFS lock status, refer to [“Verifying CFS Lock Status” section on page 13-24](#).

## Committing Changes

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

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

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

- One or more external switches report a successful status—The application applies the changes locally and releases the fabric lock.

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

## CFS Merge Support

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

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

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

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

## CFS Distribution over IP

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

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




---

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

---

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




---

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

---

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

Figure 13-1 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 13-1 Network Example 1 with Fibre Channel and IP Connections**




Figure 13-2 is the same as Figure 13-1 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 13-2 Network Example 2 with Fibre Channel and IP Connections**





Figure 13-3 is the same as Figure 13-2 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 13-3 Network Example 3 with Fibre Channel and IP Connections**



## 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.

## About CFS Regions

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



### Note

---

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

---

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

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

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

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



## Guidelines and Limitations

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

CFS has the following guidelines and limitations:

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

## Default Settings

Table 13-1 lists the default settings for CFS configurations.

**Table 13-1** 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:eff:4653

## Configuring CFS

This section describes the configuration process and includes the following topics:

- [Disabling CFS Distribution on a Switch section, page 13-10](#)
- [Enabling CFS for an Application section, page 13-11](#)

- [Committing Changes](#) section, page 13-12
- [Clearing a Locked Session](#) section, page 13-13
- [Enabling CFS over IP](#) section, page 13-14
- [Configuring IP Multicast Address for CFS over IP](#) section, page 13-14
- [Configuring Static IP Peers for CFS over IP](#) section, page 13-15
- [Adding Peers to List](#) section, page 13-17

## Disabling CFS Distribution on a Switch

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

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

### Restrictions

- 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.

### Detailed Steps

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

	Command	Purpose
Step 1	switch# <b>config t</b> switch(config)#	Enters configuration mode.
Step 2	switch(config)# <b>no cfs distribute</b>	Globally disables CFS distribution for all applications on the switch, including CFS over IP.
	switch(config)# <b>cfs distribute</b>	Enables (default) CFS distribution on the switch.

To globally disable or enable CFS distribution on a switch, 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.

---

## Enabling CFS for an Application

### Restrictions

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

To enable CFS for a feature, 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.

---

**Step 3** Right-click the row you changed to see the pop-up menu. Select **Apply Changes** to apply the CFS configuration change. The CFS tab updates as the CFS changes take effect.

DCNM-SAN retrieves the status of the CFS change and updates the Last Result column.

---

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.

---

**Step 3** Click **Pending Differences** to compare the configuration of this feature on this switch to other switches in the fabric or VSAN that have CFS enabled for this feature. Close the Show Pending Diff pop-up window.

**Step 4** Click **Apply** to apply the CFS configuration change.

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

---

## Committing Changes

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

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

To commit changes 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 the NTP **General** tab and change any configuration. Click the **Apply Changes** icon to apply the configuration to the local switch. The change remains at the pending database at the local switch until 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.
  - Step 5** Click the **CFS** tab, right-click the value in the Config Action column of the master switch that is selected 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.

DCNM-SAN 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 the NTP **General** tab and change any configuration. Click the **Apply Changes** icon to apply the configuration to the local switch. The change remains at the pending database at the local switch until 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 changes 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 using the **abort** command for that feature.

You can discard changes for a specified feature by setting the Command column value to **disable** for that feature, and 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.

To clear the CFS locks, follow these steps:

	Command	Purpose
Step 1	switch# <b>config t</b> switch(config)#	Enters configuration mode.
Step 2	Switch# <b>Conf t</b> Switch(conf)# <b>dpvm abort</b>	Aborts the configuration from the switch where the configuration lock was acquired previously. This method clears the CFS locks in the entire fabric.  Clears the CFS locks for the DPVM application in the entire fabric.
	Switch# <b>Conf t</b> Switch# <b>clear dpvm session</b>	Clears the sessions from any switch in the fabric.  Clears the CFS locks for the DPVM application.

To clear locks, 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.
- Step 3** Click **Apply Changes** icon to save the change.

## Troubleshooting Tips

- 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.

## Enabling CFS over IP

### Detailed Steps

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

	Command	Purpose
Step 1	switch# <b>config t</b> switch(config)#	Enters configuration mode.
Step 2	switch(config)# <b>cfs ipv4 distribute</b>	Globally enables CFS over IPv4 for all applications on the switch.
	switch(config)# <b>no cfs ipv4 distribute</b> This will prevent CFS from distributing over IPv4 network. Are you sure? (y/n) [n] <b>y</b>	Disables (default) CFS over IPv4 on the switch.

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

	Command	Purpose
Step 1	switch# <b>config t</b> switch(config)#	Enters configuration mode.
Step 2	switch(config)# <b>cfs ipv6 distribute</b>	Globally enables CFS over IPv6 for all applications on the switch.
	switch(config)# <b>no cfs ipv6 distribute</b>	Disables (default) CFS over IPv6 on the switch.

## Configuring IP Multicast Address for CFS over IP

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



### Note

CFS distributions for application data use directed unicast.

### Detailed Steps

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

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

	Command	Purpose
Step 1	switch# <b>config t</b> switch(config)#	Enters configuration mode.
Step 2	switch(config)# <b>cfs ipv4 mcast-address 239.255.1.1</b> Distribution over this IP type will be affected Change multicast address for CFS-IP ? Are you sure? (y/n) [n] <b>y</b>	Configures the IPv4 multicast address for CFS distribution over IPv4. The ranges of valid IPv4 addresses are 239.255.0.0 through 239.255.255.255 and 239.192/16 through 239.251/16.
	switch(config)# <b>no cfs ipv4 mcast-address 239.255.1.1</b> Distribution over this IP type will be affected Change multicast address for CFS-IP ? Are you sure? (y/n) [n] <b>y</b>	Reverts to the default IPv4 multicast address for CFS distribution over IPv4. The default IPv4 multicast address for CFS is 239.255.70.83.

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

	Command	Purpose
Step 1	switch# <b>config t</b> switch(config)#	Enters configuration mode.
Step 2	switch(config)# <b>cfs ipv6 mcast-address ff15::e244:4754</b> Distribution over this IP type will be affected Change multicast address for CFS-IP ? Are you sure? (y/n) [n] <b>y</b>	Configures the IPv6 multicast address for CFS distribution over IPv6. The range of valid IPv6 addresses is ff15::/16 (ff15::0000:0000 through ff15::ffff:ffff) and ff18::/16 (ff18::0000:0000 through ff18::ffff:ffff).
	switch(config)# <b>no cfs ipv6 mcast-address ff15::e244:4754</b> Distribution over this IP type will be affected Change multicast address for CFS-IP ? Are you sure? (y/n) [n] <b>y</b>	Reverts to the default IPv6 multicast address for CFS distribution over IPv6. The default IPv6 multicast address for CFS over IP is ff15::efff:4653.

## Configuring Static IP Peers for CFS over IP

### Detailed Steps

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

	Command	Purpose
Step 1	switch# <b>config t</b> switch(config)#	Enters configuration mode.

	Command	Purpose
Step 2	<pre>switch(config)# <b>cfs static-peers</b> WARNING: This mode will stop dynamic discovery and rely only on these peers. Do you want to continue? (y/n) [n] <b>y</b> switch(config-cfs-static)#</pre>	Enters CFS static peers configuration mode and disables dynamic discovery of peers using multicast forwarding.
	<pre>switch(config)# <b>no cfs static-peers</b> WARNING: This mode will disable static IP peer configuration and start dynamic discovery of the peers. Do you want to continue? (y/n) [n] <b>y</b> switch(config)#</pre>	Disables CFS static peer discovery and enables dynamic peer discovery using multicast forwarding on all switches.
Step 3	<pre>switch(config-cfs-static)# <b>ip address</b> <b>1.2.3.4</b> switch(config-cfs-static)# <b>ip address</b> <b>1.2.3.5</b> switch(config-cfs-static)# <b>end</b> switch#</pre>	Adds the IP address to the static peers list and marks the switch as CFS-capable. To display the static IP peers list, use the <b>show cfs static peers</b> command.
	<pre>switch(config-cfs-static)# <b>no ip address</b> <b>1.2.3.3</b> switch(config-cfs-static)# <b>end</b> switch#</pre>	Removes the IP address from the static peers list and moves the switch to dynamic peer discovery using multicast forwarding.
Step 4	<pre>switch# <b>show cfs static peers</b></pre>	Displays the IP address, WWN, and the status of CFS static peer request: <ul style="list-style-type: none"> <li>• Discovery Inprogress</li> <li>• Local</li> <li>• Reachable</li> <li>• Unreachable</li> <li>• Local IP not present</li> <li>• Rediscovery and distribution disabled</li> </ul>

**Note**

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

Cisco DCNM-SAN 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.

DCNM for SAN 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 DCNM-SAN does not update the list automatically.



## Adding Peers to List

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

---

**Step 1** From the DCNM-SAN menu, select **Tools > NPV CFS Setup**.

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.

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

A 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.

**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.

**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, follow these steps:

---

**Step 1** From the DCNM-SAN menu, select **Tools > NPV CFS Setup**.

The NPV CFS Setup wizard is launched.

**Step 2** From the **NPV Device to retrieve peer list from** the drop-down list, 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. A message box is displayed:

**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.

---

# Configuring CFS Regions

This section contains the following topics:

- [Managing CFS Regions section, page 13-18](#)
- [Creating CFS Regions section, page 13-18](#)
- [Assigning Applications to CFS Regions section, page 13-20](#)
- [To assign an application on a switch to a region, perform this task: Moving an Application to a Different CFS Region section, page 13-20](#)
- [Removing an Application from a Region section, page 13-21](#)
- [Assigning Features to CFS Regions section, page 13-19](#)
- [Moving a Feature to a Different Region section, page 13-20](#)
- [Removing a Feature from a Region section, page 13-21](#)
- [Deleting CFS Regions section, page 13-21](#)

## Managing CFS Regions

This section describes how to use DCNM-SAN for managing CFS regions. DCNM-SAN 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.



### Note

CFS always works within an individual fabric when no CFS region is applied, or within an individual CFS region when the CFS region exist. Even when a SAN or data center (higher than fabric) node or scope is selected, DCNM-SAN only shows the switches for the first fabric under the selected scope.

## Creating CFS Regions

To create a CFS region, perform this task:

	Command	Purpose
Step 1	switch# <b>config t</b> switch(config)#	Enters configuration mode.
Step 2	switch(config)# <b>cfs region 4</b>	Creates a region, for example, number 4.

To create a CFS region, follow these steps:

- Step 1** Expand **Switches** and select **CFS** from the Physical Attributes pane.  
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.
- 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.

---

## Assigning Features to CFS Regions

### Restrictions

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 [“Moving a Feature to a Different Region” section on page 13-20](#).

To assign a feature to a region, follow these steps:

---

**Step 1** Expand **Switches** and select **CFS** from the Physical Attributes pane.

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.

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.

**Step 3** Click the **Create Row** button on the toolbar.**Step 4** From the drop-down list, 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.

---

## Assigning Applications to CFS Regions

### Detailed Steps

To assign an application on a switch to a region, perform this task: **Moving an Application to a Different**

	Command	Purpose
Step 1	switch# <b>config t</b> switch(config)#	Enters configuration mode.
Step 2	switch(config)# <b>cfs region 4</b>	Creates a region, for example, number 4.
Step 3	switch(config-cfs-region)# <b>ntp</b> switch(config-cfs-region)# <b>callhome</b>	Adds application(s).

## CFS Region

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

### Detailed Steps

To move an application, perform this task:

	Command	Purpose
Step 1	switch# <b>config t</b> switch(config)#	Enters configuration mode.
Step 2	switch(config)# <b>cfs region 2</b>	Enters Region 2.
Step 3	switch(config-cfs-region)# <b>ntp</b> switch(config-cfs-region)# <b>callhome</b>	Indicates application(s) to be moved into Region 2 that originally belong to Region 1. For example, here, the NTP and Call Home applications are moved to Region 2.



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

## Moving a Feature to a Different Region

### Prerequisites

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, follow these steps:

- 
- Step 1** Expand **Switches** and select **CFS** from the Physical Attributes pane.  
The information pane displays the Global, IP Multicast, Feature by Region, and All Regions tabs.

- Step 2** Click the **Feature by Region** tab.
  - 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 an Application from a Region

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

### Detailed Steps

To remove applications from Region 1, perform this task:

	Command	Purpose
<b>Step 1</b>	switch# <b>config t</b> switch(config)#	Enters configuration mode.
<b>Step 2</b>	switch(config)# <b>cfs region 1</b>	Enters Region 1.
<b>Step 3</b>	switch(config-cfs-region)# <b>no ntp</b> switch(config-cfs-region)# <b>no callhome</b>	Removes application(s) that belong to Region 1, which you want to move.

## Removing a Feature from a Region

To remove a feature from a region, 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.
  - Step 3** Click **Yes** to confirm row deletion from the table in view.
- 

## Deleting CFS Regions

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

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

	Command	Purpose
Step 1	switch# <b>config t</b> switch(config)#	Enters configuration mode.
Step 2	switch(config)# <b>no cfs region 4</b> WARNING: All applications in the region will be moved to default region. Are you sure? (y/n) [n]	Deletes the Region 4.



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

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

## Verifying CFS Configurations

To display the CFS configuration information, perform the following task:

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

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

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

This section includes the following topics:

- [Verifying CFS Distribution Status section, page 13-23](#)

- [Verifying Application Registration Status section, page 13-23](#)
- [Verifying CFS Lock Status section, page 13-24](#)
- [Verifying the CFS over IP Configuration section, page 13-24](#)
- [Verifying IP Multicast Address Configuration for CFS over IP section, page 13-25](#)
- [Verifying Static IP Peer Configuration section, page 13-25](#)
- [Verifying CFS Regions section, page 13-25](#)
- [Displaying CFS Configuration Information section, page 13-25](#)

## Verifying CFS Distribution Status

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

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

## Verifying Application Registration Status

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



### Note

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

```
switch# show cfs application
-----
Application      Enabled   Scope
-----
ntp              No       Physical-all
fscm             Yes      Physical-fc
islb            No       Physical-fc
role            No       Physical-all
rscn            No       Logical
radius          No       Physical-all
fctimer         No       Physical-fc
syslogd         No       Physical-all
callhome        No       Physical-all
fcdomain        No       Logical
device-alias    Yes      Physical-fc

Total number of entries = 11
```

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

```
switch# show cfs application name ntp
```

```

Enabled      : Yes
Timeout      : 5s
Merge Capable : Yes
Scope        : Physical
Region       : Default

```

## Verifying CFS Lock Status

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

```

switch# show cfs lock

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

Application: port-security
Scope      : Logical
-----
VSAN   Domain   IP Address      User Name      User Type
-----
1      238      10.76.100.167  admin          CLI/SNMP v3
2      211      10.76.100.167  admin          CLI/SNMP v3
Total number of entries = 2

```

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

### Example 13-1 Displays the Lock Information for the Specified Application

```

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

```

## Verifying the CFS over IP Configuration

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

```

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

```



## Verifying IP Multicast Address Configuration for CFS over IP

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

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

## Verifying Static IP Peer Configuration

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

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

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

```
switch# show cfs static peers
-----
IP address      WWN name                Status
-----
1.2.3.4         00:00:00:00:00:00:00:00 Discovery Inprogress
1.2.3.5         20:00:00:0d:ec:06:55:b9 Reachable
1.2.3.6         20:00:00:0d:ec:06:55:c0 Local
```

## Verifying CFS Regions

To display the CFS regions, perform this task:

	Command	Purpose
Step 1	switch# <b>confi g t</b> switch(config)#	Enters configuration mode.
Step 2	switch(config)# <b>show cfs region brief</b>	Displays brief information about the CFS regions.
Step 3	switch(config)# <b>show cfs region</b>	Displays detailed information about the CFS regions.

## Displaying CFS Configuration Information

To display the status of CFS distribution on the switch, 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, the currently registered applications that 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.

---

## Configuration Examples for CFS

The example in this section show how to configure CFS.

### CFS Example Using DCNM for SAN

This procedure is an example of what you see when you use DCNM-SAN 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.

**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. DCNM-SAN 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.

**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.

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

If you make any changes, the status automatically changes to **Pending**.

**Step 10** Click the **Commit CFS Pending Changes** icon to save the changes.

**Step 11** The status changes to **Running**.

**Step 12** From the Config Action drop-down list, select **abortChanges** for each switch that you enabled for CFS .



**Note** DCNM-SAN 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, follow these steps:

- Step 1** Choose the feature that needs a merge master for CFS. For example, 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

### Restrictions

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**.

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, 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**.

## Field Descriptions for CFS

This section describes the field descriptions for CFS.

### Cisco Fabric Services (CFS) Features

Field	Description
Globally Enabled	Check this box to allow CFS on this switch to distribute feature configurations to other switches. Uncheck the box to prevent CFS from distributing the configuration to other switches.
Feature	The name of the CFS-capable feature.
Status	Status of the CFS-capable feature.
Command	The action to be triggered for the feature. Actions include: <ul style="list-style-type: none"> <li>• noop—No operation.</li> <li>• enable—Enable CFS distribution on the switch.</li> <li>• disable—Disable CFS distribution on the switch.</li> <li>• commit—Commit changes made since the session began.</li> <li>• abort—Discard changes made, and close the session.</li> <li>• clear—Discard changes made without closing the session.</li> </ul>
Type	The last CFS feature scope type used.
VSAN Id	The ID of the VSAN on which this feature is running.
RegionId	The distribution region ID that this CFS capable feature maps to. This region is required to be defined prior to its usage.

Field	Description
View Config Changes As	Determines whether to view the changes as running or pending. A pending configuration exists until a Commit or Abort action is triggered for that feature. If the value is running then all subsequent configuration retrieval for this feature will be from the running configuration on the local device. If the value is pending then all subsequent configuration retrieval for this feature will be from the pending configuration on the local device.
LastCommand	The last action performed on this feature.
Result	Result of the action performed on the CFS-capable feature.
Scope	The value of this object represents the attributes of a CFS-capable feature as registered with the CFS infrastructure. <ul style="list-style-type: none"> <li>• fcFabric—Indicates that the CFS based distribution for a feature spans the entire FC (Fibre Channel) fabric.</li> <li>• ipNetwork—Indicates that the CFS based distribution for a feature spans the entire IP network.</li> <li>• vsanScope—Indicates that the CFS based distribution for a feature is done on per VSAN basis and restricted to a specific VSAN in a FC (Fibre Channel) fabric.</li> </ul>
PendingConfOwnerAddr	The address of the device in the fabric where the pending configuration exists for the feature.

#### Related Topics

[Using the CFS Infrastructure](#)

[Disabling CFS Distribution on a Switch](#)

[Enabling CFS for an Application](#)

## Cisco Fabric Services (CFS) IP Multicast

Field	Description
IP Address Type	The IP address type (IPv4, IPv6, or DNS).

Field	Description
Multicast Address Domain	The multicast address domain to which the CFS distribution is restricted. There is a default multicast address for both IPv4 and IPv6 through which the keepalive messages are sent and received to discover the CFS capable switches over IP. All switches with similar multicast address form one CFS-over-IP fabric. The default multicast address for IPv4 is 239.255.70.83 and range supported is [239.255.0.0 to 239.255.255.255] The default multicast address for IPv6 is ff13::7743:4653 and the supported range is [ff13::0000:0000 to ff13::ffff:ffff].
Action	Specifies the current operating mode employed in CFS for distribution over the corresponding type of Internet address. By setting the value of this object to enable, CFS will enable its capability to distribute the application data across the fabric over the corresponding type of Internet address. By setting the value of this object to disable, CFS will disable its capability to distribute the data across the fabric over the corresponding type of Internet address.

## Cisco Fabric Services (CFS) Feature by Region

Field	Description
Feature	Identifies the name of a CFS-capable feature within a distribution region.
RegionId	Identifies a CFS distribution region.

## Cisco Fabric Services (CFS) All Region

Field	Description
RegionId	Identifies a CFS distribution region.

## Cisco Fabric Services (CFS) Owner

Field	Description
Feature, VSAN	The name of the CFS-capable feature, and the VSAN in which the feature is enabled or committed.
Name or IP Address	The name or IP address of the switch on which the feature is enabled or committed.
UserName	The name of the user who enabled or committed the feature.
Type	The last CFS feature scope type used.

## Cisco Fabric Services (CFS) Merge

Field	Description
Feature	The name of the CFS-capable feature.
CFS Merge Status Value	The result of the last fabric merge that occurred.

## Additional References

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

- [MIBs section, page 13-31](#)

## MIBs

MIBs	MIBs Link
<ul style="list-style-type: none"> <li>• CISCO-CFS-CAPABILITY-MIB</li> <li>• CISCO-CFS-MIB</li> </ul>	To locate and download MIBs, go to the following URL: <a href="http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html">http://www.cisco.com/en/US/products/ps5989/prod_technical_reference_list.html</a>

## Feature History for CFS

Table 13-2 lists the release history for this feature. Only features that were introduced or modified in Release 3.x or a later release appear in the table.

**Table 13-2** Feature History for CFS

Feature Name	Releases	Feature Information
CFS Pending Difference	5.0(1a)	Added the new CFS pending difference screen shot.
IP Static Peers for CFS over IP	4.1(1a)	New NPV CFS Setup wizard. Added IP static peers configuration steps for CFS distribution over IP.
CFS enhancements	3.2(1)	Support for CFS regions.
CFS over IP	3.0(1)	Allows CFS distributions over IP connections.
CFS support for allowed domain ID lists	3.0(1)	Allows the allowed domain ID lists to be distributed in the fabric using the CFS infrastructure.
CFS support for RCSN	3.0(1)	Allows the RCSN timer value to be distributed in the fabric using the CFS infrastructure.
CFS Regions	3.2(1)	Region tabs added to CFS and dialog boxes for creating a region and assigning a feature to a region  Enables users to configure a CFS region that contains a subset of switches in a switch fabric.

**Table 13-2** Feature History for CFS (continued)

Feature Name	Releases	Feature Information
CFS support for allowed domain ID lists	3.0(1)	Allowed DomainIds tab under VSAN, Domain Manager Enables the allowed domain ID lists to be distributed in the fabric using the CFS infrastructure.
CFS over IP	3.0(1)	Enables CFS distributions over IP connections.
CFS support for RCSN	3.0(1)	CFS tab under VSAN, Domain Manager, Advanced Enables the RCSN timer value to be distributed in the fabric using the CFS infrastructure.
CFS Regions	3.2.(1)	Added All Regions and Feature by Region tabs Allows to view and manage CFS Regions. You can create regions, assign features to region and move features from one region to other using the All Regions and Feature by Region tabs.