



Using Cisco Fabric Services

This chapter contains the following sections:

- [Information About CFS, on page 1](#)
- [CFS Distribution, on page 2](#)
- [CFS Support for Applications, on page 3](#)
- [CFS Regions, on page 6](#)
- [Configuring CFS over IP, on page 9](#)
- [Default Settings for CFS, on page 10](#)

Information About CFS

Some features in the Cisco Nexus Series switch require configuration synchronization with other switches in the network to function correctly. Synchronization through manual configuration at each switch in the network can be a tedious and error-prone process.

Cisco Fabric Services (CFS) provides a common infrastructure for automatic configuration synchronization in the network. It provides the transport function and a set of common services to the features. CFS has the ability to discover CFS-capable switches in the network and to discover feature capabilities in all CFS-capable switches.

Cisco Nexus Series switches support CFS message distribution over IPv4 or IPv6 networks.

CFS provides the following features:

- Peer-to-peer protocol with no client-server relationship at the CFS layer.
- CFS message distribution over IPv4 networks.
- Three modes of distribution.
 - Coordinated distributions—Only one distribution is allowed in the network at any given time.
 - Uncoordinated distributions—Multiple parallel distributions are allowed in the network except when a coordinated distribution is in progress.
 - Unrestricted uncoordinated distributions—Multiple parallel distributions are allowed in the network in the presence of an existing coordinated distribution. Unrestricted uncoordinated distributions are allowed to run in parallel with all other types of distributions.

The following features are supported for CFS distribution over IP:

- One scope of distribution over an IP network:
 - Physical scope—The distribution spans the entire IP network.

CFS Distribution

The CFS distribution functionality is independent of the lower layer transport. Cisco Nexus Series switches support CFS distribution over IP. Features that use CFS are unaware of the lower layer transport.

CFS Distribution Modes

CFS supports three distribution modes to accommodate different feature requirements:

- Uncoordinated Distribution
- Coordinated Distribution
- Unrestricted Uncoordinated Distributions

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 information from a peer. Parallel uncoordinated distributions are allowed for a feature.

Coordinated Distribution

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

- A network lock is acquired.
- The configuration is distributed and committed.
- The network lock is released.

Coordinated distribution has two variants:

- CFS driven—The stages are executed by CFS in response to a feature request without intervention from the feature.
- Feature driven—The stages are under the complete control of the feature.

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

Verifying the 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 : Enabled - mode IPv4
IPv4 multicast address : 239.255.70.83

Distribution over Ethernet : Enabled
```

CFS Support for Applications

CFS Application Requirements

All switches in the network must be CFS capable. Switches that are not CFS capable do not receive distributions, which results in part of the network not receiving the intended distribution. CFS has the following requirements:

- Implicit CFS usage—The first time that you issue a CFS task for a CFS-enabled application, the configuration modification process begins and the application locks the network.
- 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 network. 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 the CFS distribution state differs between applications. If CFS distribution is disabled for an application, that application does not distribute any configuration and does not accept a distribution from other switches in the network.
- 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 network, and to release the network 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.

Applications have the distribution enabled by default.

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

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
rscn             No       Logical
fctimer         No       Physical-fc
syslogd         No       Physical-all
callhome        No       Physical-all
fcdomain        Yes      Logical
device-alias    Yes      Physical-fc

Total number of entries = 8
```

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 fscm

Enabled          : Yes
Timeout         : 100s
Merge Capable   : No
Scope           : Physical-fc
```

Locking the Network

When you configure (first-time configuration) a feature (application) that uses the CFS infrastructure, that feature starts a CFS session and locks the network. When a network is locked, the switch software allows configuration changes to this feature only from the switch that holds the lock. If you make configuration changes to the feature from another switch, the switch 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 network lock but forget to end the session, an administrator can clear the session. If you lock a network at any time, your username is remembered across restarts and switchovers. If another user (on the same machine) tries to perform configuration tasks, that user's attempts are rejected.

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.

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

Committing Changes

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

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 network 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 network. The network lock is not released.

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

Discarding Changes

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

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

Saving the Configuration

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



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

Clearing a Locked Session

You can clear locks held by an application from any switch in the network to recover from situations where locks are acquired and not released. This function requires Admin permissions.

**Caution**

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

CFS Regions

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 network spans a vast geography, you might need to localize or restrict the distribution of certain profiles among a set of switches based on their physical proximity. CFS regions allow you to create multiple islands of distribution within the network 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 network.

**Note**

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

Example Scenario

The Smart Call Home application triggers alerts to network administrators when a situation arises or something abnormal occurs. When the network covers many geographies, and there are multiple network administrators who are each responsible for a subset of switches in the network, the Smart Call Home application sends alerts to all network administrators regardless of their location. For the Smart 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. You can achieve this scenario 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 network. You can configure regions from 1 through 200. The default region maintains backward compatibility.

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

Creating CFS Regions

You can create a CFS region.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# cfs region <i>region-id</i>	Creates a region.

Assigning Applications to CFS Regions

You can assign an application on a switch to a region.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# cfs region <i>region-id</i>	Creates a region.
Step 3	switch(config-cfs-region)# <i>application</i>	Adds application(s) to the region. Note You can add any number of applications on the switch to a region. If you try adding an application to the same region more than once, you see the "Application already present in the same region" error message.

Example

The following example shows how to assign applications to a region:

```
switch# configure terminal
switch(config)# cfs region 1
switch(config-cfs-region)# ntp
switch(config-cfs-region)# callhome
```

Moving an Application to a Different CFS Region

You can move an application from one region to another region.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# cfs region <i>region-id</i>	Enters CFS region configuration submenu.

	Command or Action	Purpose
Step 3	switch(config-cfs-region)# <i>application</i>	Indicates application(s) to be moved from one region into another. Note If you try moving an application to the same region more than once, you see the "Application already present in the same region" error message.

Example

The following example shows how to move an application into Region 2 that was originally assigned to Region 1:

```
switch# configure terminal
switch(config)# cfs region 2
switch(config-cfs-region)# ntp
```

Removing an Application from a Region

Removing an application from a region is the same as moving the application back to the default region (Region 0), which brings the entire network into the scope of distribution for the application.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# cfs region <i>region-id</i>	Enters CFS region configuration submenu.
Step 3	switch(config-cfs-region)# no <i>application</i>	Removes application(s) that belong to the region.

Deleting CFS Regions

Deleting a region nullifies the region definition. All the applications bound by the region are released back to the default region.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# no cfs region <i>region-id</i>	Deletes the region. Note You see the, "All the applications in the region will be moved to the default region" warning.

Configuring CFS over IP

Enabling CFS over IPv4

You can enable or disable CFS over IPv4.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# cfs ipv4 distribute	Globally enables CFS over IPv4 for all applications on the switch.
Step 3	(Optional) switch(config)# no cfs ipv4 distribute	Disables (default) CFS over IPv4 on the switch.

Verifying the CFS Over IP Configuration

The following example show how to verify the CFS over IP configuration:

```
switch# show cfs status
Distribution : Enabled
Distribution over IP : Enabled - mode IPv4
IPv4 multicast address : 239.255.70.83
```

Configuring IP Multicast Addresses for CFS over IP

All CFS over IP enabled switches with similar multicast addresses form one CFS over IP network. 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.

Configuring IPv4 Multicast Address for CFS

You can configure a CFS over IP multicast address value for IPv4. The default IPv4 multicast address is 239.255.70.83.

Procedure

	Command or Action	Purpose
Step 1	switch# configure terminal	Enters global configuration mode.
Step 2	switch(config)# cfs ipv4 mcast-address <i>ipv4-address</i>	Configures the IPv4 multicast address for CFS distribution over IPv4. The ranges of valid IPv4

	Command or Action	Purpose
		addresses are 239.255.0.0 through 239.255.255.255 and 239.192/16 through 239.251/16.
Step 3	(Optional) switch(config)# no cfs ipv4 mcast-address <i>ipv4-address</i>	Reverts to the default IPv4 multicast address for CFS distribution over IPv4. The default IPv4 multicast address for CFS is 239.255.70.83.

Verifying the IP Multicast Address Configuration for CFS over IP

The following example shows how to verify the IP multicast address configuration for CFS over IP:

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

Default Settings for CFS

The following table lists the default settings for CFS configurations.

Table 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

The CISCO-CFS-MIB contains SNMP configuration information for any CFS-related functions. See the MIB reference for your platform.