



Using Cisco Fabric Services

Cisco Nexus 5000 Series switches provide Cisco Fabric Services (CFS) capability, which simplifies provisioning by automatically distributing configuration information to all switches in the network.

Switch features can use the CFS infrastructure to distribute feature data or configuration data required by the feature.

This chapter contains the following sections:

- [Information About CFS, page 21-1](#)
- [CFS Distribution, page 21-2](#)
- [CFS Support for Applications, page 21-6](#)
- [CFS Regions, page 21-10](#)
- [Configuring CFS Over IP, page 21-12](#)
- [Displaying CFS Distribution Information, page 21-14](#)
- [Default Settings, page 21-16](#)

Information About CFS

Some features in the Cisco Nexus 5000 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 discovering feature capabilities in all CFS capable switches.

Cisco Nexus 5000 Series switches support CFS message distribution over Fibre Channel, IPv4 or IPv6 networks. If the switch is provisioned with Fibre Channel ports, CFS over Fibre Channel is enabled by default. CFS over IP must be explicitly enabled.

CFS provides the following features:

- Peer-to-peer protocol with no client-server relationship at the CFS layer.
- CFS message distribution over Fibre Channel, IPv4 or IPv6 networks.
- Three modes of distribution.
 - Coordinated distributions: Only one distribution is allowed in the network at any given time.

Send feedback to nx5000-docfeedback@cisco.com

- 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.

The following features are supported for CFS distribution over Fibre Channel SANs:

- Three scopes of distribution over SAN fabrics.
 - 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 features require configuration distribution over some specific VSANs. These features can specify to CFS the set of VSANs over which to restrict the distribution.
- Supports a merge protocol that facilitates the merge of feature configuration during a fabric merge event (when two independent SAN fabrics merge).

CFS Distribution

The CFS distribution functionality is independent of the lower layer transport. Cisco Nexus 5000 Series switches support CFS distribution over IP and CFS distribution over Fibre Channel. Features that use CFS are unaware of the lower layer transport.

Additional details are provided in the following sections:

- [CFS Distribution Modes, page 21-2](#)
- [Enabling/Disabling CFS Distribution on a Switch, page 21-3](#)
- [Verifying CFS Distribution Status, page 21-3](#)
- [CFS Distribution over IP, page 21-4](#)
- [CFS Distribution over Fibre Channel, page 21-5](#)
- [CFS Distribution Scopes, page 21-5](#)
- [CFS Merge Support, page 21-6](#)

CFS Distribution Modes

CFS supports three distribution modes to accommodate different feature requirements. Only one mode is allowed at any given time. CFS distribution modes are described in the following sections:

- [Uncoordinated Distribution, page 21-3](#)
- [Coordinated Distribution, page 21-3](#)
- [Unrestricted Uncoordinated Distributions, page 21-3](#)

Send feedback to nx5000-docfeedback@cisco.com

Uncoordinated Distribution

Uncoordinated distributions are used to distribute information that is not expected to conflict with that 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. 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:

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

Enabling/Disabling CFS Distribution on a Switch

If the switch is provisioned with Fibre Channel ports, CFS over Fibre Channel is enabled by default. CFS over IP must be explicitly enabled.

You can globally disable CFS on a switch to isolate the features using CFS from network-wide distributions while maintaining physical connectivity. When CFS is globally disabled on a switch, CFS operations are restricted to the switch.

To globally disable or enable CFS distribution on a switch, perform this task:

	Command	Purpose
Step 1	switch# configure switch(config)#	Enters configuration mode.
Step 2	switch(config)# no cfs distribute	Globally disables CFS distribution (CFS over Fibre Channel or IP) for all applications on the switch.
	switch(config)# cfs distribute	Enables (default) CFS distribution on the switch.

Verifying CFS Distribution Status

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

Send feedback to nx5000-docfeedback@cisco.com

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

CFS Distribution over IP

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 9000 Family switches running release 2.x or later.

Figure 21-1 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 21-1 Network Example 1 with Fibre Channel and IP Connections



Send feedback to nx5000-docfeedback@cisco.com

Figure 21-2 is the same as Figure 21-1 except that node C and node D 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 21-2 Network Example 2 with Fibre Channel and IP Connections





Figure 21-3 is the same as Figure 21-2 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 21-3 Network Example 3 with Fibre Channel and IP Connections



CFS Distribution over Fibre Channel

For FCS distribution over Fibre Channel, the CFS protocol layer resides on top of the FC2 layer. 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 Distribution Scopes

Different applications on the Cisco Nexus 5000 Series switches need to distribute the configuration at various levels. The following levels are available when using CFS distribution over Fibre Channel:

- 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.



Note Logical scope is not supported for FCS distribution over IP.

Send feedback to nx5000-docfeedback@cisco.com

- Physical topology level (physical scope)
Some applications (such as NTP) need to distribute the configuration to the entire physical topology.
- Between two selected switches
Some applications operate only between selected switches in the network.

CFS Merge Support

CFS Merge is supported for CFS distribution over Fibre Channel.

An application keeps the configuration synchronized in a SAN 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 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 have a 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.

CFS Support for Applications

The following topics describe the CFS capabilities that support applications:

- [CFS Application Requirements, page 21-6](#)
- [Enabling CFS for an Application, page 21-7](#)
- [Locking the Network, page 21-8](#)
- [Committing Changes, page 21-9](#)
- [Discarding Changes, page 21-9](#)
- [Saving the Configuration, page 21-9](#)
- [Clearing a Locked Session, page 21-9](#)

CFS Application Requirements

All switches in the network must be CFS capable. Switches that are not CFS capable do not receive distributions and result in part of the network not receiving the intended distribution.

Send feedback to nx5000-docfeedback@cisco.com

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 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 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 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
```

Send feedback to nx5000-docfeedback@cisco.com

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 lastly 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 (or 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 holding 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 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.

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 for the specified application:

```
switch# show cfs lock name ntp
Scope      : Physical
-----
```

Send feedback to nx5000-docfeedback@cisco.com

```

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

```

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 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 only supported 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.

The CISCO-CFS-MIB contains SNMP configuration information for any CFS-related functions. Refer to the *Cisco Nexus 5000 Series 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 network to recover from situations where locks are acquired and not released. This function requires Admin permissions.

Send feedback to nx5000-docfeedback@cisco.com

**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

This section contains the following topics:

- [About CFS Regions, page 21-10](#)
- [Example Scenario, page 21-10](#)
- [Managing CFS Regions, page 21-11](#)

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 may 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 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 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 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.

[Send feedback to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

Managing CFS Regions

This section describes how to manage a CFS region. A set of commands are used to complete the following tasks:

- [Creating CFS Regions, page 21-11](#)
- [Assigning Applications to CFS Regions, page 21-11](#)
- [Moving an Application to a Different CFS Region, page 21-11](#)
- [Removing an Application from a Region, page 21-12](#)
- [Deleting CFS Regions, page 21-12](#)

Creating CFS Regions

To create a CFS region, perform this task:

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

Assigning Applications to CFS Regions

To assign an application on a switch to a region, perform this task:

	Command	Purpose
Step 1	switch# configure switch(config)#	Enters configuration mode.
Step 2	switch(config)# cfs region <i>region-id</i>	Creates a region.
Step 3	switch(config-cfs-region)# ntp switch(config-cfs-region)# callhome	Adds application(s).

Moving an Application to a Different CFS Region

To move an application for example, from Region 1 (originating region) with NTP and Call Home applications assigned to it, to Region 2 (target region), perform this task:

	Command	Purpose
Step 1	switch# configure switch(config)#	Enters configuration mode.
Step 2	switch(config)# cfs region <i>region-id</i>	Enters CFS region configuration submenu.
Step 3	switch(config-cfs-region)# ntp switch(config-cfs-region)# callhome	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.”

[Send feedback to nx5000-docfeedback@cisco.com](mailto:nx5000-docfeedback@cisco.com)

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 network into the scope of distribution for the application.

To remove applications from a region, perform this task:

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

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# configure switch(config)#	Enters configuration mode.
Step 2	switch(config)# no cfs region <i>region-id</i>	Deletes the region.



Note

After Step 2, you see the warning, "All the applications in the region will be moved to the default region."

Configuring CFS Over IP

The following sections provide information about configuring CFS over IP:

- [Enabling CFS Over IP, page 21-12](#)
- [Verifying the CFS Over IP Configuration, page 21-13](#)
- [Configuring IP Multicast Address for CFS over IP, page 21-13](#)
- [Verifying IP Multicast Address Configuration for CFS over IP, page 21-14](#)

Enabling CFS Over IP



Note

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

Send feedback to nx5000-docfeedback@cisco.com

To enable or disable CFS over IPv4, perform this task:

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

To enable or disable CFS over IPv6, perform this task:

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

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 : Enabled - mode IPv4
IPv4 multicast address : 239.255.70.83
IPv6 multicast address : ff15::efff:4653
```

Configuring IP Multicast Address 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.

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 ff13:7743:4653.

Send feedback to nx5000-docfeedback@cisco.com

To configure an IP multicast address for CFS over IPv4, perform this task:

	Command	Purpose
Step 1	switch# configure switch(config)#	Enters configuration mode.
Step 2	switch(config)# cfs ipv4 mcast-address <i>ipv4-address</i> Distribution over this IP type will be affected Change multicast address for CFS-IP ? Are you sure? (y/n) [n] y	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)# no cfs ipv4 mcast-address <i>ipv4-address</i> Distribution over this IP type will be affected Change multicast address for CFS-IP ? Are you sure? (y/n) [n] y	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, perform this task:

	Command	Purpose
Step 1	switch# configure switch(config)#	Enters configuration mode.
Step 2	switch(config)# cfs ipv6 mcast-address <i>ipv6-address</i> Distribution over this IP type will be affected Change multicast address for CFS-IP ? Are you sure? (y/n) [n] y	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)# no cfs ipv6 mcast-address <i>ipv6-address</i> Distribution over this IP type will be affected Change multicast address for CFS-IP ? Are you sure? (y/n) [n] y	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.

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
```

Displaying CFS Distribution Information

The **show cfs merge status name** command displays the merge status for a given application. The following example displays the output for an application distributing in logical scope. It shows the merge status in all valid VSANs on the switch. The command output shows the merge status as one of the following: Success, Waiting, or Failure or In Progress. In case of a successful merge, all the switches in

Send feedback to nx5000-docfeedback@cisco.com

the network are shown under the local network. In case of a merge failure or a merge in progress, the local network and the remote network involved in the merge are indicated separately. The application server in each network that is mainly responsible for the merge is indicated by the term Merge Master.

```
switch# show cfs merge status name port-security

Logical [VSAN 1] Merge Status: Failed
Local Fabric
-----
Domain Switch WWN          IP Address
-----
238    20:00:00:05:30:00:6b:9e  10.76.100.167  [Merge Master]

Remote Fabric
-----
Domain Switch WWN          IP Address
-----
236    20:00:00:0e:d7:00:3c:9e  10.76.100.169  [Merge Master]

Logical [VSAN 2] Merge Status: Success
Local Fabric
-----
Domain Switch WWN          IP Address
-----
211    20:00:00:05:30:00:6b:9e  10.76.100.167  [Merge Master]
1      20:00:00:0e:d7:00:3c:9e  10.76.100.169

Logical [VSAN 3] Merge Status: Success
Local Fabric
-----
Domain Switch WWN          IP Address
-----
221    20:00:00:05:30:00:6b:9e  10.76.100.167  [Merge Master]
103    20:00:00:0e:d7:00:3c:9e  10.76.100.169
```

The following example of the **show cfs merge status name** command output displays an application using the physical scope with a merge failure. The command uses the specified application name to display the merge status based on the application scope.

```
switch# show cfs merge status name ntp

Physical Merge Status: Failed
Local Fabric
-----
Switch WWN          IP Address
-----
20:00:00:05:30:00:6b:9e  10.76.100.167  [Merge Master]

Remote Fabric
-----
Switch WWN          IP Address
-----
20:00:00:0e:d7:00:3c:9e  10.76.100.169  [Merge Master]
```

The **show cfs peers** command output displays all the switches in the physical network in terms of the switch WWN and the IP address. The local switch is indicated as Local.

```
switch# show cfs peers

Physical Fabric
-----
Switch WWN          IP Address
-----
20:00:00:05:30:00:6b:9e  10.76.100.167  [Local]
```

Send feedback to nx5000-docfeedback@cisco.com

```
20:00:00:0e:d7:00:3c:9e 10.76.100.169
```

```
Total number of entries = 2
```

The **show cfs peers name** command displays all the peers for which a particular application is registered with CFS. The command output shows all the peers for the physical scope or for each of the valid VSANs on the switch, depending on the application scope. For physical scope, the switch WWNs for all the peers are indicated. The local switch is indicated as Local.

```
switch# show cfs peers name ntp
```

```
Scope      : Physical
-----
Switch WWN          IP Address
-----
20:00:00:44:22:00:4a:9e 172.22.92.27   [Local]
20:00:00:05:30:01:1b:c2 172.22.92.215
```

The following example **show cfs peers name** command output displays all the application peers (all switches in which that application is registered). The local switch is indicated as Local.

```
switch# show cfs peers name port-security
```

```
Scope      : Logical [VSAN 1]
-----
Domain    Switch WWN          IP Address
-----
124      20:00:00:44:22:00:4a:9e 172.22.92.27   [Local]
98       20:00:00:05:30:01:1b:c2 172.22.92.215
```

```
Total number of entries = 2
```

```
Scope      : Logical [VSAN 3]
-----
Domain    Switch WWN          IP Address
-----
224      20:00:00:44:22:00:4a:9e 172.22.92.27   [Local]
151      20:00:00:05:30:01:1b:c2 172.22.92.215
```

```
Total number of entries = 2
```

Default Settings

Table 21-1 lists the default settings for CFS configurations.

Table 21-1 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.

Send feedback to nx5000-docfeedback@cisco.com

Table 21-1 *Default CFS Parameters (continued)*

Parameters	Default
IPv4 multicast address	239.255.70.83.
IPv6 multicast address	ff15::eff:4653.

Send feedback to nx5000-docfeedback@cisco.com