



CHAPTER 6

Using the CFS Infrastructure

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

This chapter contains the following sections:

- [About CFS, page 6-1](#)
- [Disabling CFS Distribution on a Switch, page 6-4](#)
- [CFS Application Requirements, page 6-5](#)
- [Enabling CFS for an Application, page 6-5](#)
- [Locking the Fabric, page 6-6](#)
- [Committing Changes, page 6-7](#)
- [Discarding Changes, page 6-8](#)
- [Saving the Configuration, page 6-8](#)
- [Clearing a Locked Session, page 6-8](#)
- [CFS Merge Support, page 6-8](#)
- [CFS Distribution over IP, page 6-11](#)
- [CFS Regions, page 6-15](#)
- [Default Settings, page 6-17](#)

About CFS

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

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

Send documentation comments to mdsfeedback-doc@cisco.com

This section includes the following topics:

- [Cisco SAN-OS Features Using CFS, page 6-2](#)
- [CFS Features, page 6-2](#)
- [CFS Protocol, page 6-3](#)
- [CFS Distribution Scopes, page 6-3](#)
- [CFS Distribution Modes, page 6-3](#)

Cisco SAN-OS Features Using CFS

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

- NTP (see the [“NTP CFS Distribution”](#) section on page 5-22).
- Dynamic Port VSAN Membership (see the [“DPVM Database Distribution”](#) section on page 23-5).
- Distributed Device Alias Services (see the [“Device Alias Databases”](#) section on page 24-2).
- IVR topology (see the [“Database Merge Guidelines”](#) section on page 22-36).
- SAN device virtualization (see the [“Configuring SDV”](#) section on page 21-4).
- TACACS+ and RADIUS (see the [“AAA Server Distribution”](#) section on page 32-30).
- User and administrator roles (see the [“Role-Based Authorization”](#) section on page 39-1).
- Port security (see the [“Port Security Configuration Distribution”](#) section on page 39-11).
- iSNS (see the [“iSNS”](#) section on page 42-79).
- Call Home (see the [“Call Home Configuration Distribution”](#) section on page 54-13).
- Syslog (see the [“System Message Logging Configuration Distribution”](#) section on page 55-8).
- fctimer (see the [“About fctimer Distribution”](#) section on page 29-6).
- SCSI flow services (see the [“Configuring SCSI Flow Services”](#) section on page 48-3).
- Saving startup configurations in the fabric using the Fabric Startup Configuration Manager (FSCM) (see the [“Saving Startup Configurations in the Fabric”](#) section on page 8-4).
- Allowed domain ID lists (see the [“About Allowed Domain ID Lists”](#) section on page 18-10).
- RSCN timer (see the [“Configuring the RSCN Timer”](#) section on page 27-10).
- iSLB (see the [“About iSLB Configuration Distribution Using CFS”](#) section on page 42-55).

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.

Send documentation comments to mdsfeedback-doc@cisco.com

- Three modes of distribution.
 - Coordinated distributions: Only one distribution is allowed in the fabric at any given time.
 - Uncoordinated distributions: Multiple parallel distributions are allowed in the fabric except when a coordinated distribution is in progress.
 - Unrestricted uncoordinated distributions: Multiple parallel distributions are allowed in the fabric in the presence of an existing coordinated distribution. Unrestricted uncoordinated distributions are allowed to run in parallel with all other types of distributions.
- Supports a merge protocol that facilitates the merge of application configuration during a fabric merge event (when two independent fabrics merge).

CFS Protocol

The CFS functionality is independent of the lower layer transport. Currently, in Cisco MDS switches, the CFS protocol layer resides on top of the FC2 layer and is peer-to-peer with not 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 (see the [“CFS Distribution over IP” section on page 6-11](#)).

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.
- Betweenselected 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Uncoordinated Distribution

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

Coordinated Distribution

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

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

Coordinated distribution has two variants:

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

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

Unrestricted Uncoordinated Distributions

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

Disabling CFS Distribution on a Switch

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

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

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

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

Send documentation comments to mdsfeedback-doc@cisco.com

Verifying CFS Distribution Status

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

```
switch# show cfs status
Fabric distribution Enabled
```

CFS Application Requirements

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

CFS has the following requirements:

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

Enabling CFS for an Application

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

Applications introduced in Cisco SAN-OS Release 2.0(1b) or later 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*).

Send documentation comments to mdsfeedback-doc@cisco.com



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 lastly the distribution scope.

```
switch# show cfs application name ntp

Enabled          : Yes
Timeout          : 5s
Merge Capable    : Yes
Scope            : Physical
```

Locking the Fabric

When you configure (first time configuration) a Cisco SAN-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 SAN-OS software does not allow any configuration changes from a switch, other than the switch holding the lock, to this Cisco SAN-OS feature 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.

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.

Send documentation comments to mdsfeedback-doc@cisco.com

```
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 6-1 *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
```

Committing Changes

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

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

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

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

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

Send documentation comments to mdsfeedback-doc@cisco.com

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.

Saving the Configuration

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



Caution

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

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

Clearing a Locked Session

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



Caution

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

CFS Merge Support

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

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

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

Send documentation comments to mdsfeedback-doc@cisco.com

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.

Verifying CFS Merge Status

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 the fabric are shown under the local fabric. In case of a merge failure or a merge in progress, the local fabric and the remote fabric involved in the merge are indicated separately. The application server in each fabric 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
```

Send documentation comments to mdsfeedback-doc@cisco.com

```
-----
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 fabric 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]
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
```

Send documentation comments to mdsfeedback-doc@cisco.com

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.

- Keep-alive 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 6-1 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 6-1 Network Example 1 with Fibre Channel and IP Connections




Figure 6-2 is the same as Figure 6-1 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.

Send documentation comments to mdsfeedback-doc@cisco.com

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





Figure 6-3 is the same as Figure 6-2 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 6-3 Network Example 3 with Fibre Channel and IP Connections



Enabling CFS Over IP

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

	Command	Purpose
Step 1	switch# config t 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 This will prevent CFS from distributing over IPv4 network. Are you sure? (y/n) [n] y	Disables (default) CFS over IPv4 on the switch.

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

	Command	Purpose
Step 1	switch# config t 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.

Send documentation comments to mdsfeedback-doc@cisco.com

Verifying the CFS Over IP Configuration

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

```
switch# show cfs status
Fabric distribution Enabled
IP distribution Enabled mode ipv4
```

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 keep-alive 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.

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

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# cfs ipv4 mcast-address 239.255.1.1 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 239.255.1.1 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, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# cfs ipv6 mcast-address ff15::e244:4754 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 ff15::e244:4754 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.

Send documentation comments to mdsfeedback-doc@cisco.com

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
```

Send documentation comments to mdsfeedback-doc@cisco.com

CFS Regions

This section contains the following topics:

- [About CFS Regions, page 6-15](#)
- [Managing CFS Regions, page 6-16](#)
- [Creating CFS Regions, page 6-16](#)
- [Assigning Applications to CFS Regions, page 6-16](#)
- [Moving an Application to a Different CFS Region, page 6-16](#)
- [Removing an Application from a Region, page 6-17](#)
- [Deleting CFS Regions, page 6-17](#)

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 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 Scenario: The callhome 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 callhome application sends alerts to all Network Administrators regardless of their location. For the callhome 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.

Send documentation comments to mdsfeedback-doc@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 6-16](#)
- [Assigning Applications to CFS Regions, page 6-16](#)
- [Moving an Application to a Different CFS Region, page 6-16](#)
- [Removing an Application from a Region, page 6-17](#)
- [Deleting CFS Regions, page 6-17](#)

Creating CFS Regions

To create a CFS region, follow these steps:

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

Assigning Applications to CFS Regions

To assign an application on a switch to a region, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# cfs region 4	Creates a region, for example, number 4.
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 callhome applications assigned to it, to Region 2 (target region), follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# cfs region 2	Enters the Region 2.
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 callhome 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 documentation comments to mdsfeedback-doc@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 fabric into the scope of distribution for the application.

To remove applications from Region 1, follow these steps:

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

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, follow these steps:

	Command	Purpose
Step 1	switch# config t switch(config)#	Enters configuration mode.
Step 2	switch(config)# no cfs region 4 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.”

Default Settings

Table 6-1 lists the default settings for CFS configurations.

Table 6-1 **Default CFS Parameters**

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

Send documentation comments to mdsfeedback-doc@cisco.com