



Getting Started

This chapter provides an overview of the basic configurations that need to be completed before getting started with IOA-specific configurations:

- [Enabling SSH, page 1](#)
- [Enabling CFS, page 2](#)
- [IP Access Lists, page 2](#)
- [Zone Default Policy, page 2](#)
- [FC-Redirect, page 2](#)
- [Configuring FC-Redirect v2 Mode, page 4](#)
- [Using FC-Redirect with CFS Regions, page 5](#)
- [Using IOA Cluster with IPFC Interface, page 6](#)

Enabling SSH

SSH needs to be enabled on all the IOA switches for Cisco DCNM-SAN to provision IOA. By default, the SSH service is enabled with the RSA key.

Step 1 `switch# configure terminal`
Enters configuration mode.

Step 2 `switch(config)# feature ssh updated`
Enables the use of the SSH service.

For more information about the SSH service, refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.

Enabling CFS

CFS must be enabled on the IOA switches as well as those switches of which the hosts and targets are directly connected to. FC-Redirect internally uses CFS to configure the rules for any given flow in the fabric.

Step 1 switch# **configure terminal**
Enters configuration mode.

Step 2 switch(config)# **cfs distribute**
Enables (default) CFS distribution on the switch.

For more information about CFS, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

IP Access Lists

Cluster communication requires the use of the Management interface. IP ACL configurations must allow UDP and TCP traffic on ports 9333, 9334, 9335, and 9336.

Zone Default Policy

For FC-Redirect to work correctly, the default zone policy on all the switches in the IOA environment must be configured to deny and the initiator-target pairs must be configured in user-defined zones.

FC-Redirect

FC-Redirect Supported Switches

FC-Redirect (FCR) is supported on the following switches and is enabled by default:

- Cisco MDS 9148S Switch
- Cisco MDS 9222i Switch
- Cisco MDS 9250i Switch
- Cisco MDS 9396S Switch
- Cisco MDS 9506 Switch
- Cisco MDS 9509 Switch
- Cisco MDS 9513 Switch
- Cisco MDS 9706 Switch

- Cisco MDS 9710 Switch
- Cisco MDS 9718 Switch

**Note**

Starting from Cisco MDS NX-OS Release 7.3(0)DY(1), FCR is supported on Cisco MDS 9148S and Cisco MDS 9396S switches.

FC-Redirect Unsupported Switches

IOA does not support any FCoE connected devices including devices connected through any of the MDS FCoE linecards.

Starting from Cisco MDS NX-OS Release 5.2(x), you cannot install an FCoE module in a switch that is running DMM, SME, or IOA.

FC-Redirect is not supported on the following switches, which also means that IOA is not supported:

- Cisco MDS 9148 Switch
- Cisco MDS 9140 Switch
- Cisco MDS 9134 Switch
- Cisco MDS 9124 Switch
- Cisco MDS 9120 Switch
- Cisco MDS 9020 Switch

FC-Redirect Requirements

FC-Redirect requirements for IOA include the following:

- The MDS switch with the MSM-18/4 Module installed or the 9222i Switch needs to be running Cisco MDS NX-OS Release 4.2(1) or later or the 9250i Switch needs to be running Cisco MDS NX-OS Release 6.2(5).
- The targets must be connected to a FC-Redirect-capable switch running Cisco MDS NX-OS Release 4.2(1) or later.
- 32 targets per MSM-18/4 Module can be FC-Redirected.
- In FC-Redirect v2 mode, up to 128 hosts per target are supported. If you do not enable FC-Redirect v2, this is limited to 16 hosts per target.
- CFS is enabled by default. Ensure that the CFS is enabled on the switches that have the host and the target connected. Also ensure that the CFS is not disabled on switches that are part of the IOA cluster.
- Advanced zoning capabilities such as quality of service (QoS), logical unit number (LUN) zoning, and read-only LUNs must not be used for FC-Redirect hosts and targets.

**Note**

FCR support for IVR is disabled by default. If FCR support for IVR is enabled on one switch in a fabric, it gets propagated through the entire fabric and is enabled on all the switches. We recommend that you disable FCR support for IVR using the **no fc-redirect ivr-support enable** command in global configuration mode.

Configuring FC-Redirect v2 Mode

To enable the v2 mode in FC-Redirect, use the **fc-redirect version2 enable** command in configuration mode. To disable the v2 mode in FC-Redirect, use the **no** form of the command.

This command is used to increase scalability of FC-Redirect. Disabling v2 mode after it is enabled in the fabric is not recommended. However, if you want to disable v2 mode, you cannot disable it until all FC-Redirect configurations are deleted. FC-Redirect configurations can be deleted only by deleting all corresponding application configurations.

The MDS switches not running Cisco NX-OS 3.3(1c) and later cannot be added to the fabric after the v2 mode is enabled. If the switches are added, all further FC-Redirect configuration changes will fail across the fabric. This could lead to traffic disruption for applications such as IOA, SME, and DMM.

Use the **show fc-redirect configs** command to see the list of applications that create FC-Redirect configurations.

If v2 mode is enabled in the fabric and you want to move a switch to a different fabric, use the **clear fc-redirect decommission-switch** command before moving the switch to a different fabric. If the mode is not enabled, all switches in the new fabric will be converted to v2 mode automatically.

**Note**

Ensure that there are no fabric changes or upgrades in progress. For more information see [Software Requirements](#). Use the **show fc-redirect peer-switches** command (UP state) to see all the switches in the fabric.

Step 1 Enter the following commands:
switch# **configure terminal**

```
switch(config)# fc-redirect version2 enable
```

Step 2 Enter **yes**.

Please make sure to read and understand the following implications before proceeding further:

- 1) This is a Fabric wide configuration. All the switches in the fabric will be configured in Version2 mode. Any new switches added to the fabric will automatically be configured in version2 mode.
- 2) SanOS 3.2.x switches CANNOT be added to the Fabric after Version2 mode is enabled. If any 3.2.x switch is added when Version2 mode is enabled, all further FC-Redirect Configuration changes will Fail across the fabric. This could lead to traffic disruption for applications like SME.
- 3) If enabled, Version2 mode CANNOT be disabled till all FC-Redirect

configurations are deleted. FC-Redirect configurations can be deleted ONLY after all the relevant application configurations are deleted. Please use the command 'show fc-redirect configs' to see the list of applications that created FC-Redirect configurations.

4) 'write erase' will NOT disable this command. After 'write erase' on ANY switch in the fabric, the user needs to do:

```
'clear fc-redirect decommission-switch'
```

on that that switch. Without that, if the user moves the switch to a different fabric it will try to convert all the switches in the fabric to Version2 mode automatically. This might lead to Error conditions and hence Traffic disruption.

Do you want to continue? (Yes/No) [No] **Yes**

Step 3 Enter yes.

Before proceeding further, please check the following:

- 1) All the switches in the fabric are seen in the output of 'show fc-redirect peer-switches' command and are in 'UP' state.
 - 2) All switches in the fabric are running SanOS version 3.3.x or higher.
 - 3) Please make sure the Fabric is stable ie., No fabric changes/upgrades in progress
- Do you want to continue? (Yes/No) [No] **Yes**

Using FC-Redirect with CFS Regions

The FC-Redirect feature uses Cisco Fabric Services (CFS) regions to distribute the FC-Redirect configuration. By default, the configuration is propagated to all FC-Redirect-capable switches in the fabric. CFS regions can be used to restrict the distribution of the FC-Redirect configuration.



Note

Using FC Redirect with CFS regions is an optional configuration only if the number of switches in the SAN exceeds the scalability limit supported by IOA. As of MDS NX-OS Release 4.2(1), the number of switches supported in a fabric is 34.

To learn more about CFS regions, refer to the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

Guidelines for Designing CFS Regions For FC-Redirect

To design CFS regions for FC-Redirect, follow these guidelines:

- Ensure that the CFS region configuration for FC-Redirect can be applied to all FC-Redirect-based applications. The applications include Cisco SME, Cisco DMM, Cisco IOA, and any future applications.
- Ensure that all FC-Redirect-capable switches, that are connected to the hosts, targets, and the application switches (switches with MSM-18/4 Modules in a cluster), are configured in the same region.

- All switches in the region must have a common VSAN.
- For existing IOA installations, refer to [Configuring CFS Regions For FC-Redirect](#), on page 6 for steps on migrating to CFS regions.
- Remove all instances of the previous configurations when a switch is moved to a region or moved out of a region.

Configuring CFS Regions For FC-Redirect

To configure the CFS regions for FC-Redirect, do the following tasks:

Step 1 Configure a switch in the CFS region as shown in the following example:

```
switch# config t  
switch(config)# cfs region 2  
switch(config-cfs-region)# fc-redirect  
switch(config)# end
```

Repeat this step for all the switches that are included in the specified region.

Step 2 Confirm that all the required switches are available in the CFS region by entering the **show fc-redirect peer-switches** command.

Step 3 To migrate existing Cisco IOA installations to CFS regions for FC-Redirect, delete all the existing FC-Redirect configurations created by the switches in other regions from each switch. To remove the configurations, perform the following steps:

- a) Obtain a list of all FC-Redirect configurations by entering the **show fc-redirect configs** command.
 - b) Remove all configurations created by the switches in other regions by using the **clear fc-redirect configs** command. The configurations are removed from the switches but the switches remain active in the region in which they are created.
-

Using IOA Cluster with IPFC Interface

Internet protocol over Fibre Channel (IPFC) provides IP forwarding or in-band switch management over a Fibre Channel interface (instead of management using the Gigabit Ethernet mgmt 0 interface). You can use IPFC to specify that IP frames be transported over Fibre Channel using encapsulation techniques. IP frames are encapsulated into Fibre Channel frames so that cluster management information can transmit across the Fibre Channel network without using an overlay Ethernet network.

When an IOA cluster communicates via the IPFC interface, the cluster management messages can be sent and received on Fibre Channel ISLs by encapsulating cluster management messages in Fibre Channel frames instead of using the management interface.

**Note**

- Configuring IOA cluster with the IPFC interface is optional and is supported in Cisco MDS NX-OS Release 5.0(4c) or later. Support for GUI for configuring IOA cluster with the IPFC interface might be added in the future releases.
- You must configure the nodes in an IOA cluster either to use an IPFC interface or a management interface. We do not recommend using the combination of two interface configurations.

Task Flow for Configuring IOA Cluster To Use the IPFC Interface

To configure IOA cluster using the IPFC interface, follow these steps:

-
- Step 1** Create an IPFC interface.
- a) Create a VSAN to use for in-band management.
 - b) Configure an IPv4 address and subnet mask for the VSAN interface.
 - c) Enable IPv4 routing.
 - d) Verify connectivity.
- Step 2** Create an IOA cluster.
- Step 3** Change the local node to use IPFC interface's IPv4 address.
- Step 4** Add the IOA interfaces to the cluster.
- Step 5** Add the remote node with IPFC interface IPv4 address.
- Step 6** Add the IOA interface of the remote cluster.
-

Configuring an IOA Cluster To Use the IPFC Interface

The process of configuring an IOA cluster to use the IPFC interface involves a number of configuration tasks that should be completed in the following order:

Creating a VSAN Interface and Configuring IPv4 Addresses

The first step in the process of configuring an IOA cluster to use the IPFC interface is to create a VSAN interface and configure IPv4 addresses.

To create an interface VSAN, perform this task:

-
- Step 1** Switch# **configure terminal**
Enters configuration mode
- Step 2** Switch(config)# **interface vsan 1**

Configures the interface for the specified VSAN (1).

Step 3 Switch (config-if)# **ip address 10.1.1.1 255.255.255.0**
Configures the IPV4 address and netmask for the selected interface.

Step 4 Switch (config-if)# **no shutdown**
Enables the interface.

After creating the VSAN and configuring the IPv4 address, use the **show interface vsan** command to verify the configuration:

```
sw-231-14# show interface vsan 1
vsan1 is up, line protocol is up
  WWPN is 10:00:00:0d:ec:18:a1:05, FCID is 0xec03c0
  Internet address is 10.1.1.1/24
  MTU 1500 bytes, BW 1000000 Kbit
  0 packets input, 0 bytes, 0 errors, 0 multicast
  6 packets output, 384 bytes, 0 errors, 0 dropped
sw-231-14#
```

Enabling IPv4 Routing

To enable IPv4 routing, perform this task:

Step 1 Switch# **configure terminal**
Enters configuration mode.

Step 2 Switch(config)# **ip routing**
Enables IPV4 routing.

Step 3 Switch(config)# **no ip routing**
Disables IPV4 routing.

After enabling IPv4 routing, use the show ip routing to verify the configuration.

```
sw-231-14(config)# show ip routing
ip routing is enabled
```

Verifying Connectivity

To verify the connectivity, use the show ip route and ping commands.

```
sw-231-14# show ip route
Codes: C - connected, S - static
C 10.1.1.0/24 is directly connected, vsan1
sw-231-14# ping 10.1.1.2
PING 10.1.1.2 (10.1.1.2) 56(84) bytes of data.
64 bytes from 10.1.1.2: icmp_seq=1 ttl=64 time=0.875 ms
64 bytes from 10.1.1.2: icmp_seq=2 ttl=64 time=0.866 ms
64 bytes from 10.1.1.2: icmp_seq=3 ttl=64 time=0.884 ms
64 bytes from 10.1.1.2: icmp_seq=4 ttl=64 time=0.875 ms
```



```
--- 10.1.1.2 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3023ms
rtt min/avg/max/mdev = 0.866/0.875/0.884/0.006 ms
```

Creating an IOA Cluster and IOA Interface in the Local Node

To create an IOA cluster and IOA interface in the local node, perform this task:



Note

To configure an IOA cluster, you can use the name of the switch if the network supports DNS service. The IOA cluster requires switch name to IP address resolution.

-
- Step 1** Switch# **configure terminal**
Enters configuration mode
- Step 2** Switch(config)# **ioa cluster cluster name**
Creates IOA cluster with specific name.
- Step 3** Switch(config-ioa-cl)# **node switchname/ip address ip-address 10.1.1.1**
Adds or changes the node address from the mgmt0 address to the IPFC interface address.
- Step 4** Switch(config-ioa-cl-node)# **int ioa 1/1**
Adds IOA interfaces to the cluster.
-

Verifying Cluster Configuration

To verify the cluster configuration, use the **show ioa cluster name node summary** command.

```
sw-231-14# sh ioa cluster cltr1 node sum
-----
Switch           Site           Status           Master           Node ID
-----
sw-231-14(L)     site2          online           yes              1
```

To verify the IP address of the node, use the **show ioa cluster name node** command.

```
sw-231-14# show ioa cluster cltr1 node
Node sw-231-14 is local switch
Node ID is 1
IP address is 10.1.1.1
Status is online
Belongs to Site site2
Node is the master switch
```

Adding a Remote Node and IOA Interface to the Remote Node

To add a remote node, perform this task:

-
- Step 1** Switch# **configure terminal**
Enters configuration mode.
- Step 2** Switch(config)# **ioa cluster** *cluster name*
Enter IOA cluster.
- Step 3** Switch(config-ioa-cl)# **node** <switchname/ip address> **ip-address 10.1.1.2**
Adds remote node to the cluster with the IPFC interface address.
- Step 4** Switch(config-ioa-cl-node)# **int ioa 4/1**
Adds IOA interfaces to the cluster.
-

Verifying the Cluster Configuration

To verify the node configuration, use the **show ioa cluster name node summary** command:

```
sw-231-14# show ioa cluster cltr1 node summary
-----
Switch                Site                Status              Master              Node ID
-----
sw-231-14(L)          site2               online              yes                 1
sw-231-19             site1               online              no                  2
```

To verify the IP address of the node, use the **show ioa cluster name node** command:

```
Node sw-231-14 is local switch
Node ID is 1
IP address is 10.1.1.1
Status is online
Belongs to Site site2
Node is the master switch
Node sw-231-19 is remote switch
Node ID is 2
IP address is 10.1.1.2
Status is online
Belongs to Site site1
Node is not master switch
sw-231-14#
```


To see all of the configured interfaces in the IOA cluster, use the **show ioa cluster name interface summary** command:

```
sw-231-14# show ioa cluster cltr1 interface summary
-----
Switch                Interface           Status              Flows
-----
sw-231-14(L)          ioa1/1              up                  0
sw-231-14(L)          ioa1/2              up                  0
sw-231-19             ioa4/1              up                  0
sw-231-19             ioa4/2              up                  0
sw-231-14#
```

Configuration Example

This section includes an example for creating an IOA cluster using IPFC interface. [Figure 1: Configuration Example, on page 11](#) illustrates the IOA cluster configuration used in this example. The sample topology shows the FC ISL between sw-231-14 and sw-231-19 switches.

Figure 1: Configuration Example



Creating an Interface VSAN

The following example creates an interface VSAN and configures IP address on sw-231-14 and enables IP routing:

```
sw-231-14(config)# int vsan 1
sw-231-14(config-if)# ip address 10.1.1.1 255.255.255.0
sw-231-14(config-if)# no shut
sw-231-14(config-if)# exit
sw-231-14(config)# ip routing
sw-231-14(config)#
```

The following example creates an interface VSAN and configure IP address on sw-231-19 and enables IP routing.

```
sw-231-19(config)# int vsan 1
sw-231-19(config-if)# ip address 10.1.1.12 255.255.255.0
sw-231-19(config-if)# no shut
sw-231-19(config-if)# exit
sw-231-19(config)# ip routing
```

Verifying the Configuration

The following example verifies the configuration of sw-231-14 using the **show interface** command.

```
sw-231-14# show interface vsan 1
vsan1 is up, line protocol is up
WWPN is 10:00:00:0d:ec:18:a1:05, FCID is 0xec03c0
Internet address is 10.1.1.1/24
MTU 1500 bytes, BW 1000000 Kbit
758 packets input, 110841 bytes, 0 errors, 42 multicast
651 packets output, 122577 bytes, 0 errors, 0 dropped
sw-231-14#
```

The following example verifies the configuration of sw-231-19 using the **show interface** command:

```
sw-231-19# show interface vsan 1
vsan1 is up, line protocol is up
WWPN is 10:00:00:05:30:01:9f:09, FCID is 0xc60000
Internet address is 10.1.1.12/24
MTU 1500 bytes, BW 1000000 Kbit
675 packets input, 124613 bytes, 0 errors, 36 multicast
```

```
755 packets output, 111785 bytes, 0 errors, 0 dropped
sw-231-19#
```

Verifying the Connectivity

The following example verifies the connectivity using the **ping** command:

```
sw-231-14# ping 10.1.1.2
PING 10.1.1.2 (10.1.1.2) 56(84) bytes of data.
64 bytes from 10.1.1.2: icmp_seq=1 ttl=64 time=0.868 ms
64 bytes from 10.1.1.2: icmp_seq=2 ttl=64 time=0.898 ms
64 bytes from 10.1.1.2: icmp_seq=3 ttl=64 time=0.906 ms
--- 10.1.1.2 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2017ms
rtt min/avg/max/mdev = 0.868/0.890/0.906/0.038 ms
sw-231-14#
```

Configuring IOA Site on Switch sw-231-14

The following example configures IOA site on switch sw-231-14:

```
sw-231-14(config)# ioa site-local site2
sw-231-14(config)#
```

Configuring IOA Site on Switch sw-231-19

The following example configures IOA site on switch sw-231-19:

```
sw-231-19(config)# ioa site-local site1
sw-231-19(config)#
```

Configuring IOA Cluster cltr1 on Switch sw-231-14

The following example configures IOA cluster cltr1 on switch sw-231-14:

```
sw-231-14(config)# ioa cluster cltr1
2011 Apr 8 05:00:46 sw-231-14 %CLUSTER-2-CLUSTER_LEADER_ANNOUNCE: Node 0x1 is the new
Master of cluster 0x2e05000dec18a133 of 1 nodes
2011 Apr 8 05:00:46 sw-231-14 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x2e05000dec18a133
now has quorum with 1 nodes
```

Changing the Node to Use IPFC Interface Address

The following example forces the node to use IPFC interface address:

```
sw-231-14(config-ioa-cl)# node sw-231-14 ip-address 10.1.1.1
sw-231-14(config-ioa-cl-node)# ex
```

Adding a Remote Node to the IOA Cluster

The following example adds a remote node to an IOA cluster:

```
sw-231-14(config-ioa-cl)# node sw-231-19 ip-address 10.1.1.2
2011 Apr 8 05:02:47 sw-231-14 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x2e05000dec18a133
```

```

now has quorum with 1 nodes
2011 Apr  8 05:02:52 sw-231-14 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x2e05000dec18a133
now has quorum with 2 nodes
sw-231-14(config-ioa-cl-node)# ex

```

Adding an IOA Interface to the Switch sw-231-14

The following example adds an IOA interfaces on the switch sw-231-14:

```

sw-231-14(config-ioa-cl)# node sw-231-14
sw-231-14(config-ioa-cl-node)# int ioa 1/1
sw-231-14(config-ioa-cl-node)# ex

```

Adding an IOA Interface to the Switch sw-231-19

The following example adds an IOA interface on the switch sw-231-19:

```

sw-231-14(config-ioa-cl)# node sw-231-19
sw-231-14(config-ioa-cl-node)# int ioa 4/1
sw-231-14(config-ioa-cl-node)# exit

```

Verifying the Cluster Configuration

The following example verifies the cluster configuration using the **show cluster name node summary** command:

```

sw-231-14# show ioa cluster cltrl node summary
-----
Switch           Site           Status         Master         Node ID
-----
sw-231-14(L)    site2         online         yes            1
sw-231-19      site1         online         no             2

```

Verifying the IP Address

The following example verifies the IP address that is configured on the switch using the **show ioa cluster cluster name node** command:

```

sw-231-14# show ioa cluster cltrl node
Node sw-231-14 is local switch
  Node ID is 1
  IP address is 10.1.1.1
  Status is online
  Belongs to Site site2
  Node is the master switch
Node sw-231-19 is remote switch
  Node ID is 2
  IP address is 10.1.1.2
  Status is online
  Belongs to Site site1
  Node is not master switch

```

Verifying the IOA Interface

The following example verifies the IOA interface that is configured on the switch using the **show ioa cluster cluster name interface summary** command:

```
sw-231-14# show ioa cluster cltr1 int summary
-----
Switch                Interface        Status        Flows
-----
sw-231-14 (L)         ioa1/1          up            0
sw-231-14 (L)         ioa1/2          up            0
sw-231-19             ioa4/1          up            0
sw-231-19             ioa4/2          up            0
sw-231-14#
```

Task Flow for Converting an Existing IOA Cluster to Use the IPFC Interface

To convert an existing IOA cluster to use the IPFC Interface, follow these steps:

- 1 Shut down IOA cluster on both the nodes.
- 2 Remove the IOA cluster that is configured on remote node.
- 3 Remove the remote node from the cluster on local switch and convert it as a single node cluster.
- 4 Change the local node to use IPFC by entering the commands **node id id nodename** and **ip-address IPFC address**.
- 5 Bring the single node cluster by **no-shut** on local switch.
- 6 Add the remote node and its interface on local switch.
- 7 Verify using **show** commands.

Configuration Example for Converting IOA Cluster to Use the IPFC Interface

Verifying the IOA Cluster Configuration

The following example verifies the IOA cluster configuration that is configured on the switch using **show ioa cluster cluster name node summary** command:

```
sw-231-14(config)# show ioa cluster cltnew node summary
-----
Switch                Site            Status        Master        Node ID
-----
sw-231-14 (L)         site2           online        yes           1
sw-231-19             site1           online        no            2
```

Verifying the IP Address

The following example verifies the IP address that is configured on the switch using the **show ioa cluster cluster name node** command:

```
sw-231-14(config)# show ioa cluster cltnew node
Node sw-231-14 is local switch
Node ID is 1
IP address is 172.25.231.14
Status is online
```

```

    Belongs to Site site2
    Node is the master switch
Node sw-231-19 is remote switch
    Node ID is 2
    IP address is 172.25.231.19
    Status is online
    Belongs to Site site1
    Node is not master switch

```

Verifying the Flow Status

The following example verifies the status of the flows using the **show ioa cluster *cluster name* flows** command. The nodes in this example are using mgmt0 interface address.

```

sw-231-14(config)# show ioa cluster cltnew flows
-----
Host WWN,          VSAN    WA  TA  Comp  Status  Switch,Interface
Target WWN
-----
21:01:00:1b:32:22:55:df, 1      Y   Y   N   online  sw-231-14, ioa1/1
21:01:00:0d:77:dd:f8:9d, 1
sw-231-19, ioa4/1

```

Shutting Down IOA Cluster on a Local Node

The following example shuts down the IOA cluster on a local node using **shut down** command:

```

sw-231-14(config)# ioa cluster cltnew
sw-231-14(config-ioa-cl)# shut down
This change can be disruptive. Please ensure you have read the IOA Cluster Recovery Procedure"
in the configuration guide. -- Are you sure you want to continue? (y/n) [n] y
2011 Apr  8 05:36:41 sw-231-14 %CLUSTER-2-CLUSTER_LOCAL_NODE_EXIT: Local Node 0x1 has left
the Cluster 0x2e06000dec18a133

```

Shutting Down the IOA Cluster on the Remote Node

The following example shuts down the IOA cluster on the remote node using the **shut down** command:

```

sw-231-19(config)# ioa cluster cltnew
sw-231-19(config-ioa-cl)# shut
This change can be disruptive. Please ensure you have read the IOA Cluster Recovery Procedure"
in the configuration guide. -- Are you sure you want to continue? (y/n) [n] y
2011 Apr  8 05:37:03 sw-231-19 %CLUSTER-2-CLUSTER_LOCAL_NODE_EXIT: Local Node 0x2 has left
the Cluster 0x2e06000dec18a133
sw-231-19(config-ioa-cl)# exit

```

Removing the IOA Cluster from the Remote Node

The following example removes the IOA cluster from the remote node using the **no ioa cluster *cluster name*** command:

```

sw-231-19(config)# no ioa cluster cltnew

```



Note

In case of removal of IOA cluster from the master node, clear all the FC-redirect configurations from the fabric, refer to [“Configuring CFS Regions For FC-Redirect, on page 6”](#).

Verifying the IOA Cluster in the Remote Node

The following example verifies the absence of the IOA cluster on the remote node using the **show ioa cluster cluster name** command:

```
sw-231-19(config)# show ioa cluster
sw-231-19(config)#
```

Removing the Remote Node from the Cluster in the Local Switch

The following example removes the remote node from the cluster in the local switch:

```
sw-231-14(config-ioa-cl)# no node sw-231-19
sw-231-14(config-ioa-cl)# show ioa cluster cltnew node summary
-----
Switch                Site                Status                Master                Node ID
-----
sw-231-14(L)          --                unknown (cluster is offline)                1
```

Changing the Local Node Configuration To Use IPFC Address

The following example changes the local node to use the IPFC address:

```
sw-231-14(config-ioa-cl)# node id 1 sw-231-14 ip-address 10.1.1.1
sw-231-14(config-ioa-cl-node)# exit
```

Activating the Single Node Cluster

The following example activates the single node cluster:

```
sw-231-14(config-ioa-cl)# no shut
This change can be disruptive. Please ensure you have read the IOA Cluster Recovery Procedure"
in the configuration guide. -- Are you sure you want to continue? (y/n) [n] y
sw-231-14(config-ioa-cl)# 2011 Apr 8 05:39:17 sw-231-14 %CLUSTER-2-CLUSTER_LEADER_ANNOUNCE:
Node 0x1 is the new Master of cluster 0x2e06000dec18a133 of 1 nodes
2011 Apr 8 05:39:17 sw-231-14 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x2e06000dec18a133
now has quorum with 1 nodes
```

Adding Remote Node with IPFC Address

The following example adds a remote node with IPFC address:

```
sw-231-14(config-ioa-cl)# node sw-231-19 ip-address 10.1.1.2
2011 Apr 8 05:39:36 sw-231-14 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x2e06000dec18a133
now has quorum with 1 nodes
2011 Apr 8 05:39:41 sw-231-14 %CLUSTER-2-CLUSTER_QUORUM_GAIN: Cluster 0x2e06000dec18a133
now has quorum with 2 nodes
```

Adding IOA Interfaces to the Remote Node

The following example adds the IOA interfaces to the remote node:

```
sw-231-14(config-ioa-cl-node)# int ioa 4/1
```



```
sw-231-14(config-ioa-cl-node)# end
sw-231-14#
```

Verifying the Cluster Nodes

The following example verifies the status of the IOA clusters using the **show ioa cluster *cluster name* node summary** command:

```
sw-231-14# show ioa cluster cltnew node summary
-----
Switch                Site                Status              Master              Node ID
-----
sw-231-14(L)          site2               online              yes                  1
sw-231-19             site1               online              no                    2
```

Verifying the Flow Status

The following example verifies the status of the IOA clusters using the **show ioa cluster *cluster name* flows** command:

```
sw-231-14# show ioa cluster cltnew flows
-----
Host WWN,              VSAN    WA  TA  Comp  Status  Switch,Interface
Target WWN
-----
21:01:00:1b:32:22:55:df, 1          Y    Y    N    online  sw-231-14, ioa1/1
21:01:00:0d:77:dd:f8:9d, 1          Y    Y    N    online  sw-231-19, ioa4/1
sw-231-14#
```

