



Creating and Managing Clusters

This section explains the steps required to create clusters. To configure other SVC features or to access the switch after initial configuration you can use one of the following CLI or Graphical User Interface (GUI) options:

- Cisco MDS 9000 Family CLI—to use the Cisco MDS CLI, follow the procedure specified in this guide.



Note The rest of this procedure uses the Cisco CLI to configure the Cisco MDS switch.

- IBM™ SAN Volume Controller CLI—to use the IBM CLI, refer to the *IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000 Command-Line Interface User's Guide*.
- IBM SAN Volume Controller GUI—to use the IBM web-based GUI, refer to the *IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000 Configuration Guide*.

This chapter includes the following sections:

- [About CSM Nodes, page 3-2](#)
- [About Clusters, page 3-2](#)
- [Selecting Nodes for a Cluster, page 3-3](#)
- [Isolating Management Traffic, page 3-4](#)
- [Creating a Cluster, page 3-5](#)
- [Adding Nodes to Clusters, page 3-6](#)
- [Verifying Nodes in a Cluster, page 3-7](#)
- [Deleting a node from a Cluster, page 3-8](#)

About CSM Nodes

A node provides virtualization, caching, migration and copy services to the SAN. Nodes are deployed in pairs with each pair of nodes forming an I/O group. The nodes belonging to the same I/O group **MUST** have different power domains. This entails that nodes of the same I/O group should come from two different CSMs. When a node fails within an I/O group then the other node in the I/O group will take over the responsibilities of the failed node. Data loss during a node failure is prevented by mirroring the IO read/write cache info across both the nodes in the I/O group.

About Clusters

Nodes are grouped into clusters of up to 2 pairs of nodes. These nodes are managed as a set (cluster), and present a single point of control for the user for configuration and service activity. For I/O purposes, so as to avoid a single point of loss of availability nodes will be grouped into pairs (I/O groups), with a single pair being responsible for serving I/O on a given VDisk. I/O traffic for a particular VDisk is, at any one time, handled exclusively by the nodes in a single I/O group. Thus, although a cluster may have many nodes within it, the nodes handle i/o in independent pairs. This means that the i/o capability of a cluster scales well, since additional throughput can simply be obtained by adding additional I/O Groups.

There are some circumstances when all the nodes in the cluster do act together rather than in pairs. At any one time, a single node in the cluster is used to manage configuration activity. The node at which the cluster got created will start off as the configuration node. This configuration node manages a cache of the configuration information that describes the cluster configuration and provides a focal point for configuration commands. Similarly, at any one time, a single node acts as the managing node for overall management of the cluster. If the configuration node or managing node fails, another node in the cluster will take over their responsibilities. The nodes also act together to implement the data migration function described in [Chapter 7, “Configuring Copy Services.”](#)

There are several advantages to managing a set of nodes as a cluster.

- All the cluster related configuration operations happen on the config node.
- Individual node operations like node addition, deletion, shutdown, can be done at the config node.
- All the nodes in the cluster run the same software version. Software upgrade can be initiated for the whole cluster instead of having to do this on a per-node basis.

Physical Topology

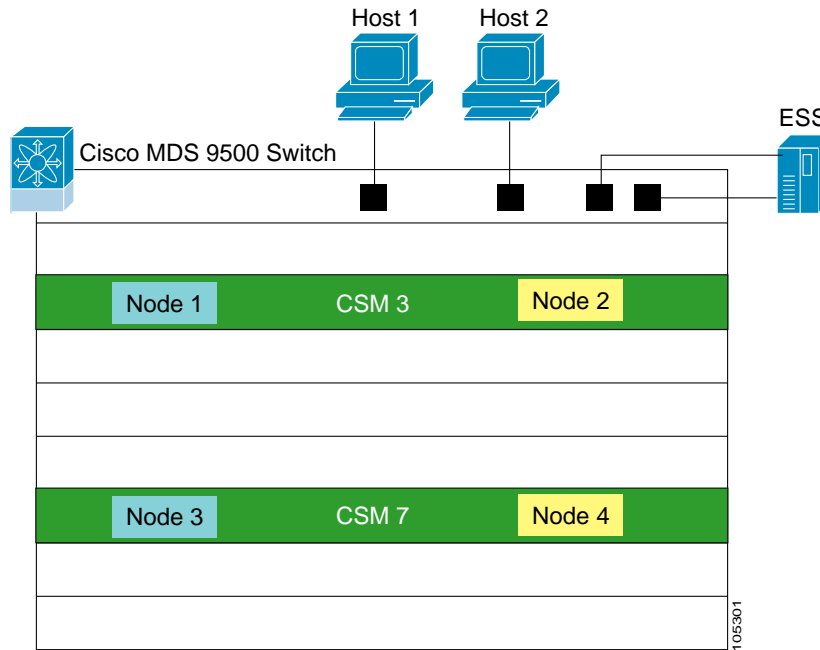
In [Figure 3-1](#), CSMs reside in slots 3 and 7 in a Cisco MDS 9500 Series switch. CSM 3 has two nodes identified as **interface svc 3/1 and interface svc 3/2**. CSM 7 has two nodes identified as **interface svc 7/1 and interface svc 7/2**. These four interfaces are configured to form a 4-node cluster.

- I/O group 1 includes interface svc 3/1 (Node1) and interface svc 7/1 (Node3).
- I/O group 2 is made up of interface svc 3/2 (Node 2) and interface svc 7/2 (Node 4).

These two I/O groups form a SVC cluster. So SVC interfaces 3/1, 3/2, 7/1, and 7/2 belong to one cluster

[Figure 3-1](#) also shows two hosts and a back-end storage device. This physical topology serves as an example in the following sections to understand SVC configurations.

Figure 3-1 Sample SVC Configuration Using Cisco MDS Switches



Selecting Nodes for a Cluster

To configure a 4-node sample configuration, follow these steps.

Step 1 Change to SVC configuration mode.

```
switch1# svc-config
```

Step 2 Display the available nodes in the local switch (switch 1).

```
switch1(svc)# show nodes local
```

```
-----
Node          cluster      config  cluster  node      sw
              cluster     node    status   status    version
-----
svc3/1        CSM 3       No      unconfigured free      1.3 (1)
svc3/2        CSM 3       No      unconfigured free      1.3 (1)
svc7/1        CSM 7       No      unconfigured free      1.3 (1)
svc7/2        CSM 7       No      unconfigured free      1.3 (1)
-----
```

Isolating Management Traffic

Management traffic is isolated for CSM nodes that are part of the same cluster using a traffic domain separator (for example VSANs or zones).

To create a separate traffic domain using VSANs, follow these steps.

Step 1 Exit the SVC configuration mode and enter the configuration mode.

```
switch1(svc)# exit
switch1# config t
switch1(config)#
```

Step 2 Configure VSAN 2 for CSM management traffic

```
switch1(config)# vsan database
switch1(config-vsan-db)# vsan 2
switch1(config-vsan-db)# exit
switch1(config)#
```

Step 3 Configure the management N-ports on VSAN 2 for all 4 SVC nodes.

```
switch1(config)# interface svc3/1
switch1(config-if)# mgmt vsan 2
switch1(config-if)# no mgmt vsan 1 <--- You have to explicitly remove from vsan 1
switch1(config-if)# no shut
switch1(config-if)# exit
switch1(config)# interface svc3/2
switch1(config-if)# mgmt vsan 2
switch1(config-if)# no mgmt vsan 1 <--- You have to explicitly remove from vsan 1
switch1(config-if)# no shut
switch1(config-if)# exit
switch1(config)# interface svc7/1
switch1(config-if)# mgmt vsan 2
switch1(config-if)# no mgmt vsan 1 <--- You have to explicitly remove from vsan 1
switch1(config-if)# no shut
switch1(config-if)# exit
switch1(config)# interface svc7/2
switch1(config-if)# mgmt vsan 2
switch1(config-if)# no mgmt vsan 1 <--- You have to explicitly remove from vsan 1
switch1(config-if)# no shut
switch1(config-if)# exit
switch1(config)#
```

Step 4 Show the FC name server database

```
switch1# show fcns database vsan 2
VSAN 2:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0x770000      N     22:32:00:05:30:00:11:69 (Cisco)           scsi-fcp:both svc
0x770001      N     22:33:00:05:30:00:11:69 (Cisco)           scsi-fcp:both svc
0x770002      N     22:34:00:05:30:00:11:69 (Cisco)           scsi-fcp:both svc
0x770003      N     22:35:00:05:30:00:11:69 (Cisco)           scsi-fcp:both svc
Total number of entries = 4
```

Step 5 Default zone configuration is permit for VSAN 2.

```
switch1(config)# zone default-zone permit vsan 2
switch1(config)#
```

Creating a Cluster

Create a cluster called SampleCluster using one node for the cluster creating. This example uses **interface svc3/1** to begin the cluster creation process. It also uses the **10.1.1.100** IP address that is in the same subnet as switch 1's management's IP network.



Note

If the cluster spans multiple switches, the switch management IP address should be in the same subnet as the cluster IP address, because the cluster IP address can move to any switch (based on the SVC config node).

To create a cluster, follow these steps.

Step 1 Create a cluster using the **cluster add** command in SVC configuration mode.

```
switch1# svc-config
switch1(svc)# cluster add SampleCluster ip 10.1.1.100 node svc3/1
Cluster creation going on. Please wait....--> This process takes a few seconds.
```

Step 2 Verify the newly-created cluster for the configured node.

```
switch1(svc)# show nodes local
```

Node	cluster	config	cluster node	node status	sw status	version
svc3/1	SampleCluster	Yes	active	active		1.3 (1)
svc3/2		No	unconfigured	free		1.3 (1)
svc7/1		No	unconfigured	free		1.3 (1)
svc7/2		No	unconfigured	free		1.3 (1)

Step 3 Display the configured cluster.

```
switch1(svc)# show cluster SampleCluster nodes
Node node1 is online(3)
Node WWN is 22:26:00:05:30:00:11:69
Serial number is JAB072006AQ
Unique id is 01:00:07:20:30:36:41:51
Node is in config mode
Node is part of iogroup id 1 name io_grp0
```

The configured node is the only node in this cluster.

Adding Nodes to Clusters

Once the initial node is used to create a cluster, you can add other required nodes to the same cluster. You can determine which nodes are available by issuing the **show cluster *cluster-name* nodes candidate** command.

To add other nodes to a cluster, follow these steps.

Step 1 Add the other nodes to the cluster.

- a. Verify which nodes can be added to the cluster. This step is important if you have several SVC nodes in the SAN,.

```
switch1(svc)# show cluster SampleCluster nodes candidate
```

```
-----
NODE                               NWWN
-----
switch1.7.2                        21:28:00:05:30:00:11:69
switch1.7.1                        21:26:00:05:30:00:11:69
switch1.3.2                        21:2a:00:05:30:00:11:69
```

This example has 3 other SVC nodes in the SAN that are candidate nodes for this cluster. The node name is an encoding of the <switch-name>.<slot-number>.<node-ID>. For example: switch1.7.2 is in the switch named switch1 at slot 7 node 2.



Caution

Do not add two nodes from the same CSM to the same I/O group of a cluster. Cisco MDS SVC does not allow this configuration as both nodes will be contained in one power domain. If both nodes are configured in the same I/O group of one cluster and a power failure occurs, both nodes will fail.

- b. Add more nodes (switch 1-slot 3-node 2, switch 1-slot 7-node 1, and switch 1-slot 7-node 2) to the newly-created cluster by entering the configuration submode for the selected cluster (SampleCluster)

```
switch1(svc)# cluster config SampleCluster
switch1(svc-cluster)# node nwwn 21:28:00:05:30:00:11:69 iogroup 1
switch1(svc-cluster)# node nwwn 21:26:00:05:30:00:11:69 iogroup 2
switch1(svc-cluster)# node nwwn 21:2a:00:05:30:00:11:69 iogroup 2
switch(svc-cluster)# exit
```

Verifying Nodes in a Cluster

After the cluster is created with the required nodes, you can verify the status of each node in the cluster, status of each node, and the associated I/O group to ensure the configuration is functioning as desired.

The node state transitions from adding to pending to online during the cluster creation process.

To verify the nodes in a cluster, follow these steps.

Step 1 Verify the nodes in the cluster.

```
switch1(svc)# show cluster SampleCluster nodes
Node node1 is online(3)
  Node WWN is 22:26:00:05:30:00:11:69
  Serial number is JAB072006AQ
  Unique id is 01:00:07:20:30:36:41:51
  Node is in config mode
  Node is part of iogroup id 1 name io_grp0

Node node2 is online(3)
  Node WWN is 21:28:00:05:30:00:11:69
  Serial number is JAB076607H8
  Unique id is 01:00:07:66:30:37:48:38
  Node is in non config mode
  Node is part of iogroup id 1 name io_grp0

Node node3 is pending(2)
  Node WWN is 21:26:00:05:30:00:11:69
  Serial number is JAB071007H8
  Unique id is 01:00:07:10:30:37:48:38
  Node is in non config mode
  Node is part of iogroup id 2 name io_grp1

Node node4 is adding(6)
  Node WWN is 00:00:00:00:00:00:00:00
  Serial number is JAB076606AQ
  Unique id is 01:00:07:66:30:36:41:51
  Node is in non config mode
  Node is part of iogroup id 2 name io_grp1
```

Step 2 Verify the I/O group status

```
switch1(svc)# show cluster SampleCluster iogroup
ID      Name                Node count  VLUN count
-----
1      io_grp0                2           0
2      io_grp1                2           0
3      io_grp2                0           0
4      io_grp3                0           0
5      recovery_io_grp        0           0
```



Note The recovery_io_group is an internal iogroup used during cluster recovery.

Step 3 Display the status of all nodes in the switch.

```
switch(svc)# show nodes local
```

Node	Cluster	Config node	Cluster status	Node status	Sw version
svc3/1	SampleCluster	Yes	Active	Active	1.3(1)
svc3/2	SampleCluster	No	Active	Active	1.3(1)
svc7/1	SampleCluster	No	Active	Active	1.3(1)
svc7/2	SampleCluster	No	active	Active	1.3(1)

The 4-node cluster is now created with the nodes communicating with each other in VSAN 2. All nodes in the switch are active and are part of cluster named SampleCluster. The SVC config node is svc3/1 (see Figure 3-2).

Figure 3-2 Creating a 4-Node Cluster



Deleting a node from a Cluster

A node has to be removed from a cluster using the **no node** command in the cluster config mode.

To delete a node that is online, follow these steps.

Step 1 Enter the cluster configuration mode for the required cluster.

```
switch1# svc-config
switch1(svc)# cluster config SampleCluster
```

Step 2 Issue the **no node nwwn** command.

```
switch1(svc-cluster)# no node nwwn 21:28:00:05:30:00:11:69 iogroup 1
```

To delete a node that is in an offline state, follow these steps.

Step 1 Enter the cluster configuration mode for the required cluster.

```
switch1# svc-config
switch1(svc)# cluster config SampleCluster
```

Step 2 Issue the **no node name** command.

```
switch1(svc-cluster)# no node name node3
```


When you delete a node in a cluster, the node is removed from the cluster state. In addition, the local state of the deleted node is also updated to indicate that it is no longer a part of any cluster.

If the node is offline, the local state of the deleted node should be explicitly updated using the **node svc x/y delete** command.

Deleting a Cluster

The MDS CLI does not use an explicit command to delete a cluster. The cluster is automatically deleted when the last node in the cluster is deleted.

