



Upgrading CSM Software

When CSMs are present in any switch in the Cisco MDS 9000 Family, several kinds of upgrade may be performed as required—a cluster software upgrade, an automatic upgrade when nodes are added, a service mode upgrade, or a switch software upgrade.

This section also explains the process to manage pWWNs and nWWNs when CSM modules are replaced in any switch in the Cisco MDS 9000 Family switch.

This chapter includes the following sections:

- [Upgrading Clusters, page 8-2](#)
- [Automatic Upgrade during Node Addition, page 8-5](#)
- [Upgrading in Service Mode, page 8-5](#)
- [Upgrading the Switch Software, page 8-7](#)
- [Replacing CSMs, page 8-9](#)

Upgrading Clusters

Cluster software upgrades are done concurrently with user I/O operations and selected management activities.



Note

When software upgrades are in progress, not all nodes in the cluster are operational depending on the upgrade phase. Consequently, the cache operates in write through mode.



Tip

Configuration changes are disallowed from the time the upgrade is started until the upgrade operation has terminated successfully. In case of failure, the upgraded and/or failed nodes must be reverted to the original software version before configuration changes are permitted. If any configuration change is attempted during this time, an error message is issued to indicate that an upgrade is in progress.

Cluster Upgrade Prerequisites

To prepare a switch for cluster software upgrade, follow these steps.

Step 1

Wait for all data migration to complete. The time taken for the data migration depends on the size of the VDisks being migrated.



Caution

Data migration, once started, cannot be stopped.

Step 2

Wait for all FlashCopy mappings to complete or stop them. If you choose to stop the FlashCopy mapping, this will result in the FlashCopy targets going offline. The FlashCopy must be prepared and started again in order to restart FlashCopy. This procedure results in the FlashCopy point-in-time being lost.

Step 3

Stop all remote copy relationships.

Recognizing Failure Situations

Cluster software upgrade will fail if one of the following situations apply:

- Any node configured to be a member of the cluster is not present—in order to upgrade the software the node must either be deleted from the cluster or must be brought online.
- All the nodes configured into the cluster do not currently have the same software level—may happen as result of a failed back out or a service mode upgrade. This behavior cannot be overridden using the **force** option.
- The new SVC software image is not compatible with the current switch or the SVC software level. This behavior cannot be overridden using the **force** option.
- A node has been deleted from the cluster such that any I/O group has only one member—the upgrade process would result in loss of access to data if this were allowed. The **force** option can be used to override this restriction if you do not mind losing access to data during the upgrade.

Cluster Upgrade Guidelines

Be aware of the following guidelines when upgrading clusters:

- Nodes are updated one at a time. One node from each I/O group is updated before the second node in any I/O group is updated.
- During the update of a node it does not participate in I/O activity in the I/O group. Thus all I/O activity for the VDisks in the I/O group is directed to the other node in the I/O group by the host multipathing software. During a node upgrade, the other node in the I/O group will attempt to flush the cache and change operations to write-through mode. This flush is not guaranteed to be successful or complete. Should the remaining node in an I/O group experience a failure when its' partner is being upgraded, then the only valid copy of the modified data in the cache could be lost. If this risk is unacceptable, then the customer must take steps to ensure that no I/O activity takes place during the software upgrade.
- A 30 minute delay is inserted into the procedure between updating the first and second nodes in each I/O group. This allows time for the host multipathing software to rediscover paths to the nodes that were updated first so that when the second set of nodes are updated loss of access does not result.
- The update is not committed until all nodes in the cluster have been successfully updated to the new code level.
- When code upgrade starts an entry is made in the event log and another entry is made when the upgrade completes or fails.
- The config node goes down to upgrade to the new software level. This will cause a different node to become the new config node. At this point, you need to reconnect to the cluster IP address.

Performing the Cluster Upgrade

Step 1 Issue the **cluster name** *cluster-name* **upgrade svc-system** command at the configuration node of the cluster.



Tip

If the cluster contains an I/O group with a single node, use the optional **force** keyword at the end of this command.

```
switch(svc)# cluster name SampleCluster upgrade svc-system
scp://userc@171.69.16.22/auto/andusr/userc/hkrel/VegasSW/build/gdb.sb-avanti/isan/targetfs
/m9000-ek9-csm-svc-mzg.1.3.x.bin
userc@171.69.16.22's password:
m9000-ek9-csm-svc-mz 100% |*****| 108 MB 01:23
```

This process may take few minutes as the config node performs the following checks:

- If the new SVC software image is compatible with the config node's running switch software and running SVC software. If any compatibility checks fail, cluster upgrade fails and an appropriate error message is issued.
- If the SVC software image is compatible with the software running on each of the switches where the other nodes in the cluster reside.
- The config node collects the compatibility check result from each node in the cluster and makes sure that all of them have succeeded. If any of nodes report a check failure, then an error message is issued with a table displaying all the check failures across all the nodes in the cluster.

- If all the nodes in the cluster are running the same level of code and all nodes are present.
- If the new SVC software image has been successfully transmitted to all the nodes in the cluster.

If all the checks succeed in Step 1 the upgrade process begins.

Step 2 Verify that cluster upgrade has succeeded by issuing one of two commands in SVC configuration mode:

- the **show node local** command—Ensure that the node software version corresponds to the new node software version.
- the **show cluster SampleCluster nodevpd** command—Ensure that the software upgrade complete event message is in the supervisor syslog.

Sample Cluster Upgrade

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

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

```
switch(svc)# cluster name SampleCluster upgrade svc-system
```

```
scp://userc@171.69.16.22/aut/andusr/userc/hkrel/VegasSW/build/gdb.sb-avanti/isan/targetfs
/m9000-ek9-csm-svc-mzg.1.3.x.bin
```

```
userc@171.69.16.22's password:
```

```
m9000-ek9-csm-svc-mz 100% |*****| 108 MB 01:23
```



Note

The cluster upgrade process has initiated successfully when the prompt returns.

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

```
-----
Node      Cluster      Config Cluster      Node      Sw
          Cluster      node  status      status      version
-----
svc3/1    SampleCluster  No    active      active      1.3 (x)
svc3/2    SampleCluster  No    active      active      1.3 (1)
svc8/1    SampleCluster  Yes   active      active      1.3 (1)
svc8/2    SampleCluster  No    active      active      1.3 (1)
-----
```

In this example, node svc4/1 has completed upgrade to the new version of the svc software.

Automatic Upgrade during Node Addition

If a node is running a SVC software version which is different from the cluster version, and that node is added to the cluster, the config node automatically downloads the cluster software to the new node. If the cluster software is not compatible with the software running on the switch that contains the new node, the add node operation will fail.

Upgrading in Service Mode

Service mode software upgrade provides a recovery alternative if the cluster cannot accept a normal software upgrade. A single node is placed in service access mode to perform a service mode software upgrade. This is in contrast to the normal method for upgrading all the nodes in the cluster.



Caution

This upgrade process is only to be used by experienced switch administrators under the care of customer support representatives.

To upgrade a single node in service mode, follow these steps.

- Step 1** Issue the **node svc slot/node servicemode** command on the required node to place the node in service mode.

```
switch(svc) # node svc 7/2 servicemode
switch(svc) #
```

- Step 2** Issue the **show node local** command in SVC configuration mode to verify that the required node is in servicemode.

```
switch1(svc) # show nodes local
-----
Node      cluster      config  cluster  node      sw
          cluster      node   status   status    version
-----
svc3/1    SampleCluster No      active   active    1.3(1)
svc3/2    SampleCluster No      active   active    1.3(1)
svc7/1    SampleCluster Yes     active   active    1.3(1)
svc7/2    SampleCluster No      unconfigured servicemode 1.3(1)
```

- Step 3** Issue the **node svc** command to begin the upgrade.'

```
switch(svc) # node svc 7/2 upgrade svc-system
scp://userc@171.69.16.22/auto/andusr/userc/hkrel/VegasSW/build/gdb.sb-avanti/isan/targetfs
/m9000-ek9-csm-svc-mzg.1.3.x.bin
```

This command checks if the new SVC software image is compatible with the running switch software and the running SVC software.

If any of these compatibility checks fail, then node upgrade fails with the appropriate error message to the user.

- Step 4** Issue the **show node local** command in SVC configuration mode to verify that the required node is running with new node software.

```
switch1(svc) # show nodes local
-----
Node      cluster      config  cluster  node      sw
          cluster      node   status   status    version
```

```

-----
svc3/1    SampleCluster    No    active    active    1.3(1)
svc3/2    SampleCluster    No    active    active    1.3(1)
svc7/1    SampleCluster    Yes   active    active    1.3(1)
svc7/2    SampleCluster    No    active    active    1.3(x)

```

**Note**

The node automatically exits the service mode when the upgrade is completed.

Sample Service Mode Upgrade

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

```

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

```

```
switch(svc)# node svc 7/2 servicemode
```

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

```

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

```

```
switch(svc)# node svc 7/2 upgrade svc-system
```

```
scp://userc@171.69.16.22/auto/andusr/userc/hkrel/VegasSW/build/gdb.sb-avanti/isan/targetfs/m9000-ek9-csm-svc-mzg.1.3.x.bin
```

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

```

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

```

Upgrading the Switch Software

This upgrade requires a dual-supervisor MDS 9500 director. This procedure only upgrades the switch software—not the SVC software that is running on the nodes in the cluster.

Use the **install all** command to do switch software upgrade irrespective of the modules present in the switch.

Performing the Switch Upgrade

To perform an automated software upgrade on any switch, follow these steps:

Step 1 Log into the switch through the console, Telnet, or SSH port of the active supervisor.

Step 2 Perform the upgrade by issuing the **install all** command.

```
switch# install all system bootflash:system-image kickstart bootflash:kickstart-image
```

At this point, the following events happen:

- The installer checks if the new switch system software is compatible with the SVC software for both nodes in a CSM and for all CSMs in the switch.
- You will receive the following descriptive information on the intended changes to your system:
 - a compatibility assessment for each module in the switch
 - a upgrade assessment
 - a module version comparison

See the “[Sample Switch Upgrade](#)” section on page 8-7 for an example of this information.

- Once the effects of the command are presented, you can choose to continue or cancel when you see this question (the default is **no**):

```
Do you want to continue y/n? [n] :y
```

Step 3 Select **y** if you choose to continue or **n** if you choose to cancel the upgrade.

- **y** = each CSM module is upgraded—one at a time, with a gap of 30 minutes between each to ensure that only one node in the I/O group is down at any time. This allows time for the host multipathing software to rediscover paths to the modules containing nodes that were upgraded first.
 - If incompatibility warnings exist, the incompatible CSM nodes are shut down after the upgrade.
 - If incompatibilities do not exist, the CSM nodes are upgraded.
- **n** = the installation process is aborted.

Step 4 Exit the switch console and open a new terminal session to view the upgraded supervisor module using the **show module** command.

Sample Switch Upgrade

```
switch# install all system
scp://usery@171.69.16.22/auto/vwsvkd/usery/m9500-sf1ek9-mzg.1.3.x.bin
For scp://usery@171.69.16.22, please enter password:
```

```

Copying image from scp://usery@171.69.16.22/auto/vwsvkd/usery/m9500-sflek9-mzg.1.3.x.bin
to bootflash:///m9500-sflek9-mz
.1.3.x.bin.
[#####] 100% -- SUCCESS

Verifying image bootflash:///b96d
[#####] 100% -- SUCCESS

Verifying image bootflash:///m9500-sflek9-mzg.1.3.x.bin
[#####] 100% -- SUCCESS

Extracting "svcl" version from image bootflash:///m9500-sflek9-mzg.1.3.x.
in.
[#####] 100% -- SUCCESS

Extracting "slc" version from image bootflash:///m9500-sflek9-mzg.1.3.x.bi
.
[#####] 100% -- SUCCESS

Extracting "system" version from image bootflash:///m9500-sflek9-mzg.1.3.x
bin.
[#####] 100% -- SUCCESS

Extracting "kickstart" version from image bootflash:///b96d.
[#####] 100% -- SUCCESS

Extracting "loader" version from image bootflash:///b96d.
[#####] 100% -- SUCCESS

Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
-----  -
      2  yes      disruptive      reset  To be installed sys image is incompatible with
node 1 running image
      2  yes      disruptive      reset  To be installed sys image is incompatible with
node 2 running image

Images will be upgraded according to following table:
Module      Image      Running-Version      New-Version  Upg-Required
-----  -
      1      slc          1.3(1)          1.3(x)      yes
      1      bios      v1.0.8(08/07/03)  v1.1.0(10/24/03)  yes

      2      svclc          1.3(1)          1.3(x)      yes
      2      svcsb          1.3(3)          1.3(3)      no
      2      svcsb          1.3(3)          1.3(3)      no
      2      bios      v1.0.8(08/07/03)  v1.1.0(10/24/03)  yes

      3      svclc          1.3(1)          1.3(x)      yes
      3      svcsb          1.3(1)          1.3(1)      no
      3      svcsb          1.3(1)          1.3(1)      no
      3      bios      v1.0.8(08/07/03)  v1.1.0(10/24/03)  yes

      5      system          1.3(1)          1.3(x)      yes
      5  kickstart          1.3(1)          1.3(1)      no
      5      bios      v1.0.8(08/07/03)  v1.1.0(10/24/03)  yes
      5      loader          1.2(2)          1.2(2)      no

      6      system          1.3(1)          1.3(x)      yes
      6  kickstart          1.3(1)          1.3(1)      no
      6      bios      v1.0.8(08/07/03)  v1.1.0(10/24/03)  yes
      6      loader          1.2(2)          1.2(2)      no

```


7	svclc	1.3(1)	1.3(x)	yes
7	svcsb	1.3(1)	1.3(1)	no
7	svcsb	1.3(1)	1.3(1)	no
7	bios	v1.0.8(08/07/03)	v1.1.0(10/24/03)	yes

Do you want to continue with the installation (y/n)? [n] **y**



Caution

If you type **yes** at this point, the switch upgrade will proceed. Since the software running on the nodes 2/1 and 2/2 are not compatible, after the switchover the nodes 2/1, 2/2 are shutdown. Nodes 3/1, 3/2, 7/1, and 7/2 on the other hand will be up and running since there were no incompatibility warning messages.



Tip

If the cluster spans multiple switches, we recommend that all switches run the same version of the switch software. When upgrading switch software in a multi-switch environment, be sure to update one switch at a time.

Replacing CSMs

When you replace the CSM, you must ensure that the new CSM is using valid nWWNs and pWWNs. You may choose to install the new CSM in a different slot or in the same slot. The process to replace the CSM differs based on this decision.



Tip

To avoid the need to reconfigure servers and controllers, we recommend that you configure the replacement nodes with the same nWWNs and pWWNs as the replaced nodes. The procedure provide in this section follows this recommendation.



Caution

If nodes being replaced are given the same nWWNs or pWWNs as previous nodes that were participating in a cluster, they must be added to same I/O group and the same cluster as the nodes being replaced. Adding nodes with the same nWWNs or pWWNs (as the replaced nodes) to a different I/O group or cluster, can result in data corruption. Refer to the *IBM TotalStorage Subsystem Device Driver User's Guide* for more information.

If the nodes being replaced are given new nWWNs or pWWNs, then perform the following additional steps after adding the nodes back to the cluster:

- Step 1** At the end of the recovery process, follow the SDD procedure to discover the new paths and to check that each VPath is presenting the correct number of paths. Refer to the *IBM TotalStorage Subsystem Device Driver User's Guide* for more information.
- Step 2** You may also need to modify the configuration of your disk controllers. If your controller uses a mapping technique to present its RAID arrays or partitions to the cluster, you must modify the port groups that belong to the cluster because the nWWN or pWWN's of the node have changed.

Prerequisites to Replacing a CSM

Step 1 Use the **show nodes local** and **show interface svc slot/node** commands to obtain the following information for each of the nodes in the CSM:

- The slot number in which the CSM is located
- The node number within the CSM (1 or 2)
- The node name—If the default name was used, then you cannot keep the same node name.
- The cluster (name) to which this node belongs
- The I/O group to which this node belongs
- The applicable nWWN and pWWNs.
- The current software version on the node

Document this information in an easily accessible location—you will be using this information to upgrade the software after replacing a CSM.

Step 2 Verify that the node is not functional.

Step 3 Verify that the other node in its I/O group is operational before deleting this node.

Step 4 Delete the nodes to be replaced from their cluster(s).

Replacing a CSM in the Same Slot

If a CSM is replaced by another CSM in the same slot in any switch in the Cisco MDS 9500 Series or in a Cisco MDS 9216 Switch, the same nWWNs and pWWNs are automatically assigned for both interfaces on the new CSM. No further configuration is required.



Caution

If the replacement nodes are assigned the same nWWNs and pWWNs as the replaced nodes, be sure to assign the nodes to the same I/O group and cluster as before. Otherwise, data corruption may occur. If the information on which I/O group and cluster the previous nodes were part of is not available, contact your reseller (if applicable) or customer service for assistance.

If a CSM is replaced by another CSM in the same slot on the same chassis, and you *do not wish to retain the same nWWNs and pWWNs*, follow this process.

Step 1 Remove the CSM from the slot as directed in the *Cisco MDS 9216 Switch* or the *Cisco MDS 9500 Series Hardware Installation Guide*.

Step 2 Use the **svc purge-wwn module** command to erase the nWWNs and pWWNs from the original slot. Issue the command after the module has been removed or when it is in the powered-down state.

```
switch# svc purge-wwn module <slot-num>
```

The old nWWNs and pWWNs will be lost from the system and never reassigned for any purpose in that chassis.

- Step 3** Replace the new CSM in the same slot as directed in the *Cisco MDS 9216 Switch* or the *Cisco MDS 9500 Series Hardware Installation Guide*. The interfaces on this new CSM will have brand new nWWN and pWWNs.

Replacing a CSM in a Different Slot

If a CSM is replaced by another CSM in a different slot on the same or different chassis, and the same nWWNs and pWWNs are to be retained, follow this process.

- Step 1** Document the nWWN and pWWN values for the CSM before removing it.
- Step 2** Remove the CSM from the slot as directed in the *Cisco MDS 9216 Switch* or the *Cisco MDS 9500 Series Hardware Installation Guide*.
- Step 3** Use the following command to erase the nWWNs and pWWNs from the original slot. Issue the command when the module has already been removed or when it is in the powered-down state.

```
switch# svc purge-wwn module <slot-num>
```

The old nWWNs and pWWNs will be lost from the system and never reassigned for any purpose in that chassis.

- Step 4** Replace the new CSM in the desired slot as directed in the *Cisco MDS 9216 Switch* or the *Cisco MDS 9500 Series Hardware Installation Guide*.
- Step 5** Wait for the module to initialize.
- Step 6** Configure the VSAN for the initiator, target, and management N-ports to match the replaced nodes:

```
switch# config t
switch (config)# interface svc new-slot-num / node-num
switch (config-if)# initiator vsan vsan-id
switch (config-if)# target vsan vsan-id
switch (config-if)# mgmt vsan vsan-id
```

- Step 7** Set the nWWN and pWWN values for the two interfaces within the module with the following commands:

```
switch (config-if)# shutdown
switch (config-if)# nwwn saved-nwwn-value
switch (config-if)# initiator vsan vsan-id pwwn saved-pwwn-value
switch (config-if)# target vsan vsan-id pwwn saved-pwwn-value
switch (config-if)# mgmt vsan vsan-id pwwn saved-pwwn-value
```



Tip

The related SVC interface must be shut down before setting the WWNs.

- Step 8** Reload the CSM using the **reload module slot-number** command for the new nWWNs to take effect.

Post-Replacement Verification

To perform a post-replacement verification, follow these steps.

-
- Step 1** Use the **show nodes local** command to verify that the nodes in the CSM have initialized without errors. If node does not initialize, the software version of the nodes on the replacement CSM may not be compatible with the version running on the switch.
- Step 2** Wait for the nodes to initialize.
- Step 3** Add the nodes back to their clusters (see the [Chapter 8, “Upgrading CSM Software”](#)).

**Tip**

If the previous node used the default name, you cannot reassign the default name to the new node. If you assigned a name to the old node, the new node can be assigned the same name.

**Note**

The Node ID of the replacement node is different from the node ID of the replaced node.

- Step 4** Use the SDD management tool on the host systems to verify that all paths are online.
-