



Replacing a Cisco Nexus 5010 cluster switch

Replacing a defective Cisco Nexus 5010 cluster switch in a cluster network is a nondisruptive procedure (NDU).

Before you begin

The following conditions must exist before performing the switch replacement in the current environment and on the replacement switch:

- Existing cluster and network configuration:
 - The Nexus 5010 cluster infrastructure must be redundant. Refer to the *Cisco Ethernet Switch* web page to ensure that you have the latest reference configuration file (RCF) and NX-OS versions on your switches.
 - All cluster ports must be in the up state.
 - Management connectivity must be present on both switches.
 - Console access to both cluster switches must be in place.
 - All cluster logical interfaces (LIFs) must be up and must not have been migrated.
- Nexus 5010 replacement switch:
 - Management network connectivity on the replacement switch must be functional.
 - Console access to the replacement switch must be in place.
 - All relevant switch ports for node connection must be disabled on all relevant ports.
 - All Inter-Switch Link (ISL) ports must be enabled.
 - The desired reference configuration file (RCF) and NX-OS operating system image switch must be loaded onto the switch.
 - Initial customization of the switch must be complete.

About this task

This procedure replaces an existing Nexus 5010 cluster switch (cs1 in this procedure) with a new Nexus 5010 switch (cs-new).

The examples in this procedure use the following switch and node nomenclature:

- The names of the existing Nexus 5010 cluster switches are cs1 and cs2.
- The name of the new Nexus 5010 cluster switch is cs-new.
- The node names of the nodes are node x for every node in a cluster.
- The cluster:: $*>$ prompt indicates the name of the cluster.
- The node-facing ports are:
 - Ports e1/1 through e1/12 on Nexus 5010 switches without expansion modules
 - Ports e1/1 through e1/12 and ports e2/1 through e2/6 on Nexus 5010 switches with an expansion module
- The ISL ports are e1/13 through e1/20 on Nexus 5010 switches whether an expansion module is present or not
- The names of the cluster LIFs connected to cs1 and cs2 are clus1 and clus2.
- The cluster ports used in this procedure are e1a and e2a. Refer to the *Hardware Universe* for the actual cluster ports supported on your platforms.

Steps

1. Install the appropriate RCF and image on the Nexus 5010 cluster switch cs-new and make any necessary site preparations.

This optional step is to verify, download, and install the appropriate versions of the RCF and NX-OS software for the new switch. If you have verified that the new switch is correctly set up and does not need updates to the RCF and NX-OS software, continue to step 2.

- a. Go to the [Cisco Ethernet Switch](#) page on The NetApp Support Site.
 - b. Note your switch and the required software versions in the table on that page.
 - c. To download the appropriate version of the RCF file, click on *Data ONTAP 8.X or later Cluster and Management Network Switch Reference Configuration Files*.
 - d. On the *Description* page, click on **CONTINUE**, accept the license agreement, and follow the instructions on the *Download* page to download the RCF file.
 - e. To download the appropriate version of the NX-OS software, click on *Data ONTAP 8.X or later Cluster Network Switch*.
2. On the new switch, shut down all of the ports that will be connected to the node (cluster ports 1/1 through 1/12 and expansion module ports 2/1 through 2/6).

If the switch that you are replacing is not functional and powered down, go to step 3. The LIFs on the cluster nodes should have already failed over to the other cluster port for each node.

Example

This example shows ports 1/1 through 1/12 and 2/1 through 2/6 with an expansion module being shut down.

```
cs-new# configure
cs-new(config)# interface ethernet 1/1-12, ethernet 2/1-6
cs-new(config-if-range)# shutdown
cs-new(config-if-range)# exit
cs-new(config)# exit
```

3. Set the privilege level of the command session to advanced and enter `y` at the prompt.

Example

The following example shows the privilege level being set to advanced:

```
cluster::> set -privilege advanced
Warning: These advanced commands are potentially dangerous; use them only when
        directed to do so by NetApp personnel.
Do you wish to continue? (y or n): y
cluster::*>
```

Note: You must be in advanced mode to use the clustered Data ONTAP commands in this procedure.

4. Verify that all cluster ports are transmitting and receiving traffic through the console or SSH connection.

Example

This example shows that both ports e1a and e2a in the cluster are transmitting and receiving traffic.

```
cluster::*> system node run -node * ifstat e1a
cluster::*> system node run -node * ifstat e2a
```

5. On the console of all nodes, use the `network interface migrate` command to migrate `clus1` to port `e2a`.

```
cluster::*> network interface migrate -vserver nodex -lif clus1 -source-node nodex -dest-
node nodex -dest-port e2a
```

6. Use the `network interface show -role cluster` command to verify that the migration took place.

The LIFs are migrated if `clus1`'s Current Port column shows `e2a` (the same port as `clus2`) and the Is Home column shows false.

Example

The following example shows the output for a cluster with two nodes:

```
cluster::*> network interface show -role cluster
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
node1						
	clus1	up/up	10.10.0.1/24	node1	e2a	false
	clus2	up/up	10.10.0.2/24	node1	e2a	true
node2						
	clus1	up/up	10.10.0.1/24	node2	e2a	false
	clus2	up/up	10.10.0.2/24	node2	e2a	true

7. Use the `cluster show` command to show the status of the node members.

```
cluster::*> cluster show
```

Node	Health	Eligibility	Epsilon
node1	true	true	false
node2	true	true	false

8. Use the `network port modify` command to shut down cluster port `e1a` on all of the nodes.

Example

The following example shuts down port `e1a` on `nodex`.

```
cluster::*> network port modify -node nodex -port e1a -up-admin false
```

9. Shut down the ISL ports 13 through 20 on the Nexus 5010 cluster switch `cs2`.

Example

The following example shows the ISL ports 13 through 20 on `cs2` being shut down:

```
cs2# configure
cs2(config)# interface ethernet 1/13-20
cs2(config-if-range)# shutdown
cs2(config-if-range)# exit
cs2(config)# exit
```

10. Disconnect *all of the cables* from the cluster ports and ISLs from the Nexus 5010 `cs1` cluster switch and reconnect them to the appropriate ports on the Nexus 5010 `cs-new` switch.
11. Bring up the ISLs 13 through 20 between the `cs-new` and `cs2` Nexus 5010 switches and verify the port channel operation status.

Example

The following example shows the ISL ports 13 through 20 being opened:

```
cs2# configure
cs2(config)# interface ethernet 1/13-20
cs2(config-if-range)# no shutdown
```

```
cs2(config-if-range)# exit
cs2(config)# exit
```

- Use the `show port-channel summary` command on both the cs-new and cs2 switches to verify that the port-channel members have a status of (P).

Example

The following example shows results for the Nexus 5010 cs2 switch:

```
cs2# show port-channel summary
Flags: D - Down          P - Up in port-channel (members)
       I - Individual    H - Hot-standby (LACP only)
       s - Suspended     r - Module-removed
       S - Switched      R - Routed
       U - Up (port-channel)
       M - Not in use. Min-links not met
-----
Group Port-          Type      Protocol  Member Ports
Channel
-----
1      Pol(SU)        Eth       LACP      Eth1/13(P)  Eth1/14(P)  Eth1/15(P)
                               Eth1/16(P)  Eth1/17(P)  Eth1/18(P)
                               Eth1/19(P)  Eth1/20(P)
```

The following sample results are for the new Nexus 5010 switch cs-new:

```
cs-new# show port-channel summary
Flags: D - Down          P - Up in port-channel (members)
       I - Individual    H - Hot-standby (LACP only)
       s - Suspended     r - Module-removed
       S - Switched      R - Routed
       U - Up (port-channel)
       M - Not in use. Min-links not met
-----
Group Port-          Type      Protocol  Member Ports
Channel
-----
1      Pol(SU)        Eth       LACP      Eth1/13(P)  Eth1/14(P)  Eth1/15(P)
                               Eth1/16(P)  Eth1/17(P)  Eth1/18(P)
                               Eth1/19(P)  Eth1/20(P)
```

- Bring up the ports on the Nexus 5010 cs-new switch that are associated with the cluster nodes.

Example

The following example shows ports 1/1 through 1/12 and 2/1 through 2/6 with an expansion module being brought up on the Nexus 5010 cs-new switch:

```
cs-new #configure
cs-new(config)# interface ethernet 1/1-12, ethernet 2/1-6
cs-new(config-if-range)# no shut
cs-new(config-if-range)# exit
cs-new(config)# exit
```

- Use the `network port modify` command to enable the first cluster port, e1a, on all nodes in the cluster.

Example

The following example shows port e1a being brought up on node.x:

```
cluster::*> network port modify -node nodex -port e1a -up-admin true
```

15. Use the `network port show -role cluster` command to verify that on all nodes that the cluster ports are up. The following example shows the output for a cluster with two nodes:

```
cluster::*> network port show -role cluster
```

Node	Port	Role	Link	MTU	Auto-Negot Admin/Oper	Duplex Admin/Oper	Speed (Mbps) Admin/Oper

node1	e1a	clus1	up	9000	true/true	full/full	auto/10000
	e2a	clus2	up	9000	true/true	full/full	auto/10000
node2	e1a	clus1	up	9000	true/true	full/full	auto/10000
	e2a	clus2	up	9000	true/true	full/full	auto/10000

4 entries were displayed.

16. On all nodes, use the `network interface revert` command to revert `clus1` (which was previously migrated) back to `e1a`.

Example

```
cluster::*> network interface revert -vserver nodex -lif clus1
```

17. Use the `network interface show` command to ensure that all cluster LIFs are up and operational and display `true` in the `Is Home` column.

Bringing up the first node is successful if the `Is Home` column is `true` for both cluster interfaces and they show the correct port assignments, which in this example are `e1a` and `e2a`.

The following example shows the typical output for two nodes in a cluster:

```
cluster::*> network interface show -role cluster
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home

node1	clus1	up/up	10.10.10.1/16	node1	e1a	true
	clus2	up/up	10.10.10.2/16	node1	e2a	true
node2	clus1	up/up	10.10.11.1/16	node2	e1a	true
	clus2	up/up	10.10.11.2/16	node2	e2a	true

4 entries were displayed.

18. Use `cluster show` to display information about all nodes in a cluster.

Example

This example shows a two-node cluster and neither node has `epsilon` in this case. In clusters with more than two nodes one of the nodes will hold `epsilon` as `true`.

```
cluster::*> cluster show
```

Node	Health	Eligibility	Epsilon

node1	true	true	false
node2	true	true	false

Related information

[Cisco Ethernet Switch web page](#)
[Hardware Universe](#)

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Customer Fitness, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.