



Replacing a Cisco® Nexus 5596 cluster switch

Replacing a defective Cisco Nexus 5596 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 5596 cluster infrastructure must be redundant and fully functional on both switches.
Note: You can refer to the *Cisco Ethernet Switch* web page to verify 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 exist on both switches.
 - All cluster logical interfaces (LIFs) must be **up** and must not have been migrated.
- Nexus 5596 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 on to the switch.
 - Initial customization of the switch must be complete.
Any previous site customizations, such as SMTP, SNMP, and SSH should be copied to the new switch.

About this task

This procedure replaces an existing Nexus 5596 cluster switch (cs1 in this procedure) with a new Nexus 5596 switch (cs-new).

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

- The names of the existing Nexus 5596 cluster switches are cs1 and cs2.
- The name of the new Nexus 5596 cluster switch is cs-new.
- The node names of the nodes are node *x* for every node in the cluster.
- The `cluster::*>` prompt indicates the name of the cluster.
- The node-facing ports are e1/1 through e1/40.
- The ISL ports are e1/41 through e1/48.
- 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.
You can 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 5596 cluster switch cs-new and make any necessary site customizations.

This optional step allows you 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, go to the next step:

- a. See the *Cisco Ethernet Switch* page on the NetApp Support Site.

[Cisco Ethernet Switch](#)

- b. Note your switch and the required software versions in the table on that page.
- c. Download the appropriate version of the RCF file.
- d. Click **CONTINUE** on the **Description** page, accept the license agreement, and follow the instructions on the **Download** page to download the RCF file.
- e. Download the appropriate version of the image software.

See the *Data ONTAP 8.X or later Cluster and Management Network Switch Reference Configuration Files* download page, and then you can click the appropriate version.

To find the correct version, see the *Data ONTAP 8.X or later Cluster Network Switch* download page.

2. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message. If AutoSupport is not enabled on this cluster, then go to the next step: **system node autosupport invoke -node * -type all -message MAINT=xh**

x is the duration of the maintenance window in hours.

Note: The message will notify technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

Example

This command suppresses automatic case creation for two hours:

```
cluster::*> system node autosupport invoke -node * -type all -message MAINT=2h
```

3. On the new switch, shut down all of the ports that will be connected to the node cluster ports 1 through 40.

If the switch that you are replacing is not functional and powered down, go to step 4 on page 2. The LIFs on the cluster nodes should have already failed over to the other cluster port for each node.

Example

The following example shows ports 1/1 through 1/40 are shut down:

```
cs-new# configure
cs-new(config)# interface ethernet 1/1-40
cs-new(config-if-range)# shutdown
cs-new(config-if-range)# exit
cs-new(config)# exit
```

4. 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 ONTAP commands in this procedure.

5. Migrate clus1 to port e2a on the console of each node: **network interface migrate**

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

Repeat this step if you have more than one cluster interconnect port connected to cs1 on the same node.

6. Verify that the migration took place: **network interface show -role cluster**

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. View the status of all node members: **cluster show**

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

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

8. Shut down the cluster port e1a on all nodes: **network port modify**

Example

The following example shows port e1a on nodex being shut down.

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

9. Shut down the ISL ports 41 through 48 on the Nexus 5596 cluster switch cs2.

Example

The following example shows the ISL ports 41 through 48 on cs2 are shut down:

```

cs2# configure
cs2(config)# interface ethernet 1/41-48
cs2(config-if-range)# shutdown
cs2(config-if-range)# exit
cs2(config)# exit

```

10. Disconnect all cables from the cluster ports and ISLs on the Nexus 5596 cs1 cluster switch.
11. Reconnect the cables to the appropriate ports on the Nexus 5596 cs new switch.
12. Bring up the ISLs 41 through 48 between the cs-new and cs2 Nexus 5596 switches and verify the port channel operation status.

Example

The following example shows ISL ports 41 through 48 are opened:

```

cs2# configure
cs2(config)# interface ethernet 1/41-48
cs2(config-if-range)# no shutdown
cs2(config-if-range)# exit
cs2(config)# exit

```

13. Verify that the port-channel members have a status of (P) on both the cs-new and cs2 switches: **show port-channel summary**

Example

The following example shows the output for the Nexus 5596 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      Po1(SU)    Eth       LACP      Eth1/41(P)  Eth1/42(P)  Eth1/43(P)
                                         Eth1/44(P)  Eth1/45(P)  Eth1/46(P)
                                         Eth1/48(P)  Eth1/48(P)

```

The following example shows output for the new Nexus 5596 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      Po1(SU)    Eth       LACP      Eth1/41(P)  Eth1/42(P)  Eth1/43(P)
                                         Eth1/44(P)  Eth1/45(P)  Eth1/46(P)
                                         Eth1/47(P)  Eth1/48(P)

```

- Bring up the ports on the Nexus 5596 cs-new switch that are associated with the cluster nodes.

Example

The following example shows ports 1 through 40 brought up on the Nexus 5596 cs-new switch:

```
cs-new # configure
cs-new(config-if-range)# interface ethernet 1/1-40
cs-new(config-if-range)# no shut
cs-new(config-if-range)# exit
cs-new(config)# exit
```

- Enable the first cluster port e1a on all nodes: **network port modify**

Example

The following example shows port e1a brought up on node1:

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

Repeat this step if you have more than one cluster interconnect port connected to cs1 on the same node.

- Verify on all nodes that the cluster ports are up: **network port show -role cluster**

```
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.

- Revert clus1 (which was previously migrated) to e1a on all nodes: **network interface revert**

Example

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

Repeat this step if you have more than one cluster interconnect port connected to cs1 on the same node.

- Verify that all cluster LIFs are up, operational and display true in the “Is Home” column: **network interface show**

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 are e1a and e2a in the following example:

```
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.
```

19. Display information about all nodes in a cluster: **cluster show**

Example

The following 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**. 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
```

20. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message: **system node autosupport invoke -node * -type all -message MAINT=END**

Example

The command reenables automatic case creation:

```
cluster::*> system node autosupport invoke -node * -type all -message MAINT=END
```

Related information

[Cisco Ethernet Switch web page](#)

[Hardware Universe](#)

[NetApp KB Article 1010449: How to suppress automatic case creation during scheduled maintenance windows](#)

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

docomments@netapp.com

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

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

Active IQ, AltaVault, Arch Design, ASUP, AutoSupport, Campaign Express, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Element, Fitness, Flash Accel, Flash Cache, Flash Pool, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, Fueled by SolidFire, GetSuccessful, Helix Design, LockVault, Manage ONTAP, MetroCluster, MultiStore, NetApp, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANscreen, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, SolidFire Helix, StorageGRID, SyncMirror, Tech OnTap, Unbound Cloud, and WAFI and other names are trademarks or registered trademarks of NetApp, 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. A current list of NetApp trademarks is available on the web.

<http://www.netapp.com/us/legal/netapptmlist.aspx>