



Replacing Cisco Nexus 5010 or 5020 cluster switches with Nexus 5596 cluster switches

You can replace Cisco Nexus 5010 or Nexus 5020 cluster switches with Nexus 5596 cluster switches in a cluster. This is a nondisruptive procedure.

Before you begin

If the Nexus 5010 or 5020 switches that are being replaced are not running NX-OS 5.2(1)N1(1), they should be upgraded first. The software release on the 5010 or 5020 switches should at least be NX-OS 5.0 (2)N1(1) for the installation, but the switches should be maintained at the same NX-OS level.

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

- The following cluster switches are supported:
 - Cisco Nexus 5596 cluster switch
 - Cisco Nexus 5010 or Nexus 5020 cluster switch
- The cluster switches use the following Inter-Switch Link (ISL) ports:
 - Cisco Nexus 5596: ports 1/41 to 1/48
 - Cisco Nexus 5020: ports 1/33 to 1/40
 - Cisco Nexus 5010: ports 1/13 to 1/20
- The switch models use the following node ports:
 - Cisco Nexus 5596: ports 1 to 40
 - Cisco Nexus 5020: 1/1 to 1/32
 - Cisco Nexus 5010 (18 node): 1/1 to 1/12, 2/1 to 2/6
 - Cisco Nexus 5010 (12 node): 1/1 to 1/12
- The ONTAP and NX-OS versions that are supported in this procedure are listed in the “Cisco Ethernet Switches” page. [Cisco Ethernet Switch](#)

About this task

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

- The 5010 or 5020 switches to be replaced by the 5596 switches are C1 and C2.
- The 5596 switches that will replace the 5010 or 5020 switches are S1 and S2.
- clus1 is the cluster logical interface (LIF) that is connected to cluster switch 1 (S1 and C1).
- clus2 is the cluster LIF 2 that is connected to cluster switch 2 (S2 and C2).
- The names of the SVMs are node1 and node2.

- The cluster ports that are used in this procedure as an example are e1a and e2a.
You can refer to the *Hardware Universe* for the actual cluster ports on your platforms.

This procedure has two parts:

- Steps [1](#) on page 2 to [15](#) on page 6 replace the second Cisco 5010 or 5020 switch with the second 5596 switch.
- Steps [16](#) on page 6 to [31](#) on page 9 replace the first Cisco 5010 or 5020 switch with the first 5596 switch.

Steps

1. Verify that the appropriate reference configuration files (RCF) and image are installed on the new 5596 switches as necessary for your requirements, and make any essential site customizations.

You should prepare both switches at this time. If you need to upgrade the RCF and image, follow these steps:

- 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. Shut down all of the ports on the 5596 switch designated as S2 that will be connected to the cluster nodes: **shutdown**

Example

The commands in this example disable ports 1 to 40 on S2:

```
S2# configure
Enter configuration commands, one per line. End with CNTL/Z.
S2(config)# interface ethernet 1/1-40
S2(config-if-range)# shutdown
S2(config-if-range)# exit
S2(config)# exit
S2#
```

- On all cluster nodes, migrate all of the cluster LIFs to the appropriate ports connected to the 5010 or 5020 switch designated as C2 that is being replaced: **network interface migrate**

You can verify the cluster health using by the `network interface show` command. The migration of a LIF can only be performed while you are logged in to the node or an SVM that the LIF is on.

Note: This step is optional. It lets you manually migrate the interface to the port connected to the surviving switch. If you skip this step, traffic will automatically fail over when you shut down the cluster ports in step 5 on page 3.

Example

The following example shows the traffic migrated from the cluster port on the 5010 or 5020 switch C2 to the 5596 switch S2 and verifies that they are migrated:

```
cluster::*> network interface migrate -vserver node1 -lif clus2 -
-dest-node node1 -dest-port e1a

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	e1a	true
	clus2	up/up	10.10.0.2/24	node1	e1a	false
node2	clus1	up/up	10.10.0.3/24	node2	e1a	true
	clus2	up/up	10.10.0.4/24	node2	e1a	false

4 entries were displayed.

- On all cluster nodes, shut down the cluster ports associated with the 5010 or 5020 switch C2 being replaced and verify the cluster health.

You can disable a port from any node in the cluster.

Example

The following example shows how to shut down port e2a on node1 and node2:

```
cluster::*> network port modify -node node1 -port e2a -up-admin false
cluster::*> network port modify -node node2 -port e2a -up-admin false
```

- Shut down the ISL ports on the active 5010 or 5020 switch C1 that is being replaced.

Example

The following example shows how to disable ISLs 33 to 40 on a 5020 switch:

```
C1# configure
Enter configuration commands, one per line. End with CNTL/Z.
C1(config)# interface ethernet 1/33-40
C1(config-if-range)# shutdown
C1(config-if-range)# exit
C1(config)# exit
C1#
```

The following example shows how to disable ISLs 13 to 20 on a 5010 switch:

```

C1# configure
Enter configuration commands, one per line.  End with CNTL/Z.
C1(config)# interface ethernet 1/13-20
C1(config-if-range)# shutdown
C1(config-if-range)# exit
C1(config)# exit
C1#

```

- Remove all of the cables from the Cisco 5010 or 5020 switch C2 and attach them to S2, the new 5596 switch.

Attach the cables from the C2 switch to the S2 switch as follows:

- Connect the C2 ports 13 to 20 from a 5010 switch to the S2 ports 41 to 48 on the 5596 switch.
- Connect the C2 ports 33 to 40 from a 5020 switch to the S2 ports 41 to 48 on the 5596 switch.

- Bring up the ISLs on C1, the active Cisco 5010 or 5020 switch.

Example

The following example shows how to enable ISLs 33 to 40 on a 5020 switch:

```

C1# configure
Enter configuration commands, one per line.  End with CNTL/Z.
C1(config)# interface ethernet 1/33-40
C1(config-if-range)# no shutdown
C1(config-if-range)# exit
C1(config)# exit
C1#

```

The following example shows how to enable ISLs 13 to 20 on a 5010 switch:

```

C1# configure
Enter configuration commands, one per line.  End with CNTL/Z.
C1(config)# interface ethernet 1/13-20
C1(config-if-range)# no shutdown
C1(config-if-range)# exit
C1(config)# exit
C1#

```

There are now eight physical connections on the 5010 or 5020 switch C1 ISLs to match the eight ISL physical connections on the 5596 switch S2.

- Verify that the connections are up on both of the switches: **show port-channel summary**

Example

The following example shows the status of the ISLs displayed on a Cisco 5020 switch C1:

The output from a 5020 here displays the status of the ISLs on the C1 switch (you will see similar output on a 5010 switch):

```

C1# 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-Channel	Type	Protocol	Member Ports

```
-----
1      Pol(SU)      Eth      LACP      Eth1/33(P)  Eth1/34(P)  Eth1/35(P)
                                     Eth1/36(P)  Eth1/37(P)  Eth1/38(P)
                                     Eth1/39(P)  Eth1/40(P)
-----
```

The output from the 5596 switch S2 here displays the status of the ISLs:

```
S2# 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      Eth1/41(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 new 5596 switch S2 that are associated with the cluster nodes.

Example

The following example shows how to bring up ports 1 to 40 on the S2 switch:

```
S2# configure
Enter configuration commands, one per line.  End with CNTL/Z.
S2(config)# interface ethernet 1/1-40
S2(config-if-range)# no shutdown
S2(config-if-range)# exit
S2(config)# exit
S2#
```

- On a single node, bring up the cluster node port associated with the new 5596 switch S2: **network port modify**

Example

The following example shows how to bring up port e2a connected to the new S2 switch:

```
cluster::*> network port modify -node node1 -port e2a -up-admin true
```

- Verify the status of the cluster node port: **network port show**

Example

The following example shows port e2a verified on the new 5596 switch S2 as being up:

```
cluster::*> network port show -role cluster
Node  Port  Role      Link MTU  Auto-Negot Duplex      Speed (Mbps)
Admin/Oper Admin/Oper Admin/Oper
-----
node1
     e1a  cluster  up    9000  true/true  full/full  auto/10000
     e2a  cluster  up    9000  true/true  full/full  auto/10000
```

- On the same node, revert cluster LIF clus2 (associated with port e2a) to its home port: **network interface revert**

```
cluster::*> network interface revert -vserver node1 -lif clus2
```

- Verify that the interface is now reverted to its home port: **network interface show**

Example

The following example shows the status of clus2 displayed as up and port e2a displayed as home (true):

```
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	10.10.10.1/24	node1	e1a	true
	clus2	up/up	10.10.10.2/24	node1	e2a	true

- If bringing up the first node is successful, bring up the ports and revert the cluster interfaces on the remaining nodes using steps 11 on page 5 to 14 on page 6.

Before proceeding to step 16 on page 6, remember to verify that the cluster is healthy and that all nodes are in membership.

- To begin the second part of the replacement process, if you need to update the RCF and image on the S1 switch, install the appropriate versions on the switch and make any necessary site customizations, as shown in step 1 on page 2.
- Shut down all of the ports on S1, the new 5596 switch that will be connected to the cluster nodes.

Example

The following example shows ports 1 to 40 being shut down on the new 5596 switch S1:

```
S1# configure
Enter configuration commands, one per line. End with CNTL/Z.
S1(config)# interface ethernet 1/1-40
S1(config-if-range)# shutdown
S1(config-if-range)# exit
S1(config)# exit
S1#
```

- On all cluster nodes, migrate the LIF clus1 connected to C1, the 5010 or 5020 switch: **network interface migrate**
Note that the migration of a LIF can only be performed while you are logged in to the node or an SVM that the LIF is on.

Example

The following example shows the LIF clus1 on an SVM named node1 being migrated to port e2a on node1:

```
cluster::*> network interface migrate -vserver node1 -lif clus1 -dest-node node1 -dest-port e2a
```

- Verify the health of the cluster: **network interface show**

Example

The following example shows that the ports have been migrated:

```
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    10.10.0.3/24    node2    e2a    false
      clus2      up/up    10.10.0.4/24    node2    e2a    true
4 entries were displayed.

```

20. On all of the nodes, shut down the node ports that are associated with cluster LIF clus2: **network port modify**

Example

The following example shows port e1a being shut down on node1:

```
cluster::*> network port modify -node node1 -port e1a -up-admin false
```

21. Shut down the ISLs on S2, the active 5596 switch that is already connected to the ports on the 5010 or 5020 switch C1 that is being replaced: **shutdown**

Example

The following example shows how to disable ISLs 41 to 48 on the active 5596 switch S2:

```

S2# configure
Enter configuration commands, one per line.  End with CNTL/Z.
S2(config)# interface ethernet 1/41-48
S2(config-if-range)# shutdown
S2(config-if-range)# exit
S2(config)# exit
S2#

```

22. Remove the cables from the 5010 or 5020 switch C1 and attach them to S1, the new 5596 switch:
- Connect the C1 ports 13 to 20 from a 5010 switch to the S1 ports 41 to 48 on the 5596 switch.
 - Connect the C1 ports 33 to 40 from a 5020 switch to the S1 ports 41 to 48 on the 5596 switch.
23. Bring up the ISLs on S2, the active 5596 switch: **no shutdown**

Example

The following example shows how to enable the ISLs being brought up on switch S2:

```

S2# configure
Enter configuration commands, one per line.  End with CNTL/Z.
S2(config)# interface ethernet 1/41-48
S2(config-if-range)# no shutdown
S2(config-if-range)# exit
S2(config)# exit
S2#

```

24. Verify that the ISLs are up on both 5596 switches S1 and S2: **show port-channel summary**

The output should be identical on each switch and all ports should be up (P).

Example

The following example shows the status of the ISLs displayed on switch S2:

```

S2# 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-Channel	Type	Protocol	Member Ports		
1	Pol(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)	

25. Bring up the node-facing ports (cluster ports) on the 5596 switch S1 that are associated with the cluster nodes.

Example

The following example shows how to enable ports 1 to 40 brought up on switch S1:

```
S1# configure
Enter configuration commands, one per line. End with CNTL/Z.
S1(config)# interface ethernet 1/1-40
S1(config-if-range)# no shutdown
S1(config-if-range)# exit
S1(config)# exit
S1#
```

26. On a single node, bring up the cluster port e1a that is connected to the 5596 switch S1: **network port modify**

Example

The following example shows port e1a brought up on node1:

```
cluster::*> network port modify -node node1 -port e1a -up-admin true
```

27. Verify the status of the cluster node port: **network port show**

Example

The following example shows port e1a verified on node1:

```
cluster::*>network port show -role cluster
Node  Port  Role      Link MTU  Auto-Negot  Duplex      Speed (Mbps)
-----  -----  -----  -----  -----  -----  -----
node1  e1a    cluster  up  9000    true/true   full/full   auto/10000
       e2a    cluster  up  9000    true/true   full/full   auto/10000
```

28. On that same node, revert cluster LIF clus1 (associated with port e1a), to its home port: **network interface revert**

Example

```
cluster::*> network interface revert -vserver node1 -lif clus1
```

29. Verify that the interface is now home: **network interface show**

Depending on the version of ONTAP that you are using, the LIF might revert as soon as the port is brought up. The `network interface show` command can verify if it is already reverted as `clus1` and `clus2` will both appear as `true` under the “Is Home” column and correspond with the correct ports.

Example

The following example shows the status of `clus1` displayed as up and port `e1a` displayed as home:


```
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      10.10.10.1/24  node1      e1a          true
             clus2      up/up      10.10.10.2/24  node1      e2a          true
```

30. If bringing up the first node is successful, bring up the ports and revert the cluster interfaces on the remaining nodes using steps [26](#) on page 8 to [30](#) on page 9.
31. 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](#)

[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.

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.

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., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

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