



---

# Replacing CN1610 cluster switches with Cisco® Nexus 5596 cluster switches

To replace the existing CN1610 cluster switches in a cluster with Cisco Nexus 5596UP or 5596T cluster switches, you must perform a specific sequence of tasks.

## Before you begin

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:
  - NetApp CN1610
  - Cisco Nexus 5596UP and 5596T
- The cluster switches use the following Inter-Switch Link (ISL) ports:
  - Cisco Nexus 5596UP and 5596T: ports 41 through 48
  - NetApp CN1610: ports 13 through 16
- NetApp twinax cables are not compatible with Cisco 5596 switches.  
If your current CN1610 configuration uses NetApp twinax cables for cluster-node-to-switch connections or ISL connections and you want to continue using twinax in your environment, you must procure Cisco twinax cables. Alternatively, you can use fiber cables for both the ISL connections and the cluster-node-to-switch connections.
- The ONTAP and NX-OS versions that are supported in this procedure are listed on the “Cisco Ethernet Switches” page. [Cisco Ethernet Switches](#)
- The ONTAP and FASTPATH versions that are supported in this procedure are listed on the “NetApp CN1601 and CN1610 Switches” page. [NetApp CN1601 and CN1610 Switches](#)

## About this task

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

- The command outputs might vary depending on different releases of ONTAP software.
- The CN1610 switches to be replaced are CL1 and CL2.
- The 5596UP or 5596T switches to replace the CN1610 switches are C1 and C2.
- clus1 is the first cluster logical interface (LIF) that is connected to cluster switch 1 (CL1 or C1).
- clus2 is the first cluster LIF that is connected to cluster switch 2 (CL2 or C2).
- clus3 is the second LIF that is connected to cluster switch 2 (CL2 or C2).
- clus4 is the second LIF that is connected to cluster switch 1 (CL1 or C1).

The names of the SVMs are node1, node2, node3, and node4.

**Note:** The examples in this procedure use four nodes. Two nodes use two cluster interconnect ports: e0a and e0b. The other two nodes use four cluster interconnect ports: e0a, e0b, e0c, and e0d. The *Hardware Universe* contains the latest information about the actual cluster ports for your platforms.

This procedure has two parts:

- Steps 1 on page 2 through 21 on page 9 replace the second CN1610 cluster switch with the second Cisco 5596UP or 5596T cluster switch.
- Steps 22 on page 9 through 39 on page 16 replace the first CN1610 cluster switch with the first Cisco 5596UP or 5596T cluster switch.

## Steps

### 1. Display information about the devices in your configuration: **network device-discovery show**

The following example displays how many cluster interconnect interfaces have been configured in each node for each cluster interconnect switch:

#### Example

```
cluster::> network device-discovery show
```

Node	Local Port	Discovered Device	Interface	Platform
node1	e0a	CL1	0/1	CN1610
	e0b	CL2	0/1	CN1610
node2	e0a	CL1	0/2	CN1610
	e0b	CL2	0/2	CN1610
node3	e0a	CL1	0/3	CN1610
	e0b	CL2	0/3	CN1610
	e0c	CL2	0/4	CN1610
	e0d	CL1	0/4	CN1610
node4	e0a	CL1	0/5	CN1610
	e0b	CL2	0/5	CN1610
	e0c	CL2	0/6	CN1610
	e0d	CL1	0/6	CN1610

12 entries were displayed

### 2. Determine the administrative or operational status for each cluster interface with the following commands:

- Display the cluster network port attributes: **network port show**

#### Example

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

Node	Port	Role	Link	MTU	Auto-Begot Admin/Open	Duplex Admin/Open	Speed (Mbps) Admin/Open
node1	e0a	cluster	up	9000	true/true	full/full	auto/10000
	e0b	cluster	up	9000	true/true	full/full	auto/10000
node2	e0a	cluster	up	9000	true/true	full/full	auto/10000
	e0b	cluster	up	9000	true/true	full/full	auto/10000
node3	e0a	cluster	up	9000	true/true	full/full	auto/10000
	e0b	cluster	up	9000	true/true	full/full	auto/10000
	e0c	cluster	up	9000	true/true	full/full	auto/10000
	e0d	cluster	up	9000	true/true	full/full	auto/10000
node4	e0a	cluster	up	9000	true/true	full/full	auto/10000

```

e0b cluster up 9000 true/true full/full auto/10000
e0c cluster up 9000 true/true full/full auto/10000
e0d cluster up 9000 true/true full/full auto/10000

```

- b. Display information about the logical interfaces: **network interface show**

**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	10.10.0.1/24	node1	e0a	true
	clus2	up/up	10.10.0.2/24	node1	e0b	true
node2	clus1	up/up	10.10.0.3/24	node2	e0a	true
	clus2	up/up	10.10.0.4/24	node2	e0b	true
node3	clus1	up/up	10.10.0.5/24	node3	e0a	true
	clus2	up/up	10.10.0.6/24	node3	e0b	true
	clus3	up/up	10.10.0.7/24	node3	e0c	true
	clus4	up/up	10.10.0.8/24	node3	e0d	true
node4	clus1	up/up	10.10.0.9/24	node4	e0a	true
	clus2	up/up	10.10.0.10/24	node4	e0b	true
	clus3	up/up	10.10.0.11/24	node4	e0c	true
	clus4	up/up	10.10.0.12/24	node4	e0d	true

12 entries were displayed.

- c. Display information about the discovered cluster switches: **system cluster-switch show**

**Example**

The following example displays the cluster switches that are known to the cluster, along with their management IP addresses:

```


cluster::> system cluster-switch show

```

Switch	Type	Address	Model
CL1 Serial Number: 01234567 Is Monitored: true Reason: Software Version: 1.2.0.7 Version Source: ISDP	cluster-network	10.10.0.101	CN1610
CL2 Serial Number: 01234568 Is Monitored: true Reason: Software Version: 1.2.0.7 Version Source: ISDP	cluster-network	10.10.0.102	CN1610

2 entries were displayed.

The following figure displays the cluster network topology utilizing two CN1610 switches:



- 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:

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

[Cisco Ethernet Switch](#)

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

- If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message: **system node autosupport invoke -node \* -type all -message MAINT=xh**

x is the duration of the maintenance window in hours.

**Note:** The message notifies 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
```

- Shut down ports 1 through 40 on the second 5596 switch C2, that is connected to the cluster nodes: **shutdown**

### Example

The following example shows ports 1 through 40 being shut down on the 5596 switch C2:

```
C2# configure
C2(config)# interface ethernet 1/1-40
C2(config-if-range)# shutdown
C2(config-if-range)# exit
C2(config)# exit
C2#
```

6. Migrate the LIFs that are associated with the second CN1610 switch to be replaced: **network interface migrate**

### Example

The following example shows the process for migrating the LIFs on node1 and node3, but you must ensure that LIF migration is done on all of the nodes:

```
cluster::*> network interface migrate -vserver node1 -lif clus2 -source-node node1 -dest-
node node1 -dest-port e0a
cluster::*> network interface migrate -vserver node3 -lif clus2 -source-node node3 -dest-
node node3 -dest-port e0a
cluster::*> network interface migrate -vserver node3 -lif clus3 -source-node node3 -dest-
node node3 -dest-port e0d
```

7. Verify the cluster's health: **network interface show**

### Example

The following example shows the result of the previous `network interface migrate` command:

```
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	e0a	true
	clus2	up/up	10.10.0.2/24	node1	e0a	false
node3	clus1	up/up	10.10.0.5/24	node3	e0a	true
	clus2	up/up	10.10.0.6/24	node3	e0a	false
	clus3	up/up	10.10.0.7/24	node3	e0d	false
	clus4	up/up	10.10.0.8/24	node3	e0d	true

8. Shut down the cluster interconnect ports that are physically connected to switch CL2: **network port modify**

### Example

The following example shows how to shut down the specified ports on node1 and node3:

```
cluster::*> network port modify -node node1 -port e0b -up-admin false
cluster::*> network port modify -node node3 -port e0b -up-admin false
cluster::*> network port modify -node node3 -port e0c -up-admin false
```

9. Ping the remote cluster interfaces, and perform an RPC server check: **cluster ping-cluster**

### Example

The following example shows how to ping the remote cluster interfaces:

```

cluster::*> cluster ping-cluster -node node1
Host is node1
Getting addresses from network interface table...
Local = 10.10.0.1      10.10.0.2
Remote = 10.10.0.3 10.10.0.4 10.10.0.5 10.10.0.6 10.10.0.7 10.10.0.8 10.10.0.9 10.10.0.10
10.10.0.11      10.10.0.12
Ping status:
....
Basic connectivity succeeds on 20 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 20 path(s):
  Local 10.10.0.1 to Remote 10.10.0.3
  Local 10.10.0.1 to Remote 10.10.0.4
  Local 10.10.0.1 to Remote 10.10.0.5
  Local 10.10.0.1 to Remote 10.10.0.6
  Local 10.10.0.1 to Remote 10.10.0.7
  Local 10.10.0.1 to Remote 10.10.0.8
  Local 10.10.0.1 to Remote 10.10.0.9
  Local 10.10.0.1 to Remote 10.10.0.10
  Local 10.10.0.1 to Remote 10.10.0.11
  Local 10.10.0.1 to Remote 10.10.0.12
  Local 10.10.0.2 to Remote 10.10.0.3
  Local 10.10.0.2 to Remote 10.10.0.4
  Local 10.10.0.2 to Remote 10.10.0.5
  Local 10.10.0.2 to Remote 10.10.0.6
  Local 10.10.0.2 to Remote 10.10.0.7
  Local 10.10.0.2 to Remote 10.10.0.8
  Local 10.10.0.2 to Remote 10.10.0.9
  Local 10.10.0.2 to Remote 10.10.0.10
  Local 10.10.0.2 to Remote 10.10.0.11
  Local 10.10.0.2 to Remote 10.10.0.12

Larger than PMTU communication succeeds on 20 path(s)
RPC status:
10 paths up, 0 paths down (tcp check)
10 paths up, 0 paths down (udp check)

```

10. Shut down ISLs 13 through 16 on the active CN1610 switch CL1: **shutdown**

#### Example

The following example shows how to shut down ISLs 13 through 16 on the CN1610 switch CL1:

```

(CL1)# configure
(CL1)(Config)# interface 0/13-0/16
(CL1)(Interface 0/13-0/16)# shutdown
(CL1)(Interface 0/13-0/16)# exit
(CL1)(Config)# exit
(CL1)#

```

11. Remove all of the cables from the CN1610 switch CL2, and attach them to the new Cisco 5596 switch C2.

For the ISL ports, reconnect only the four ports on the Cisco 5596 switch C2 (ports 45 to 48).

**Attention:** When reconnecting any cables to the new Cisco 5596 switch, the cables used must now be either fiber cables or Cisco twinax cables.

12. To make the ISL dynamic, configure the ISL interface 3/1 on the active CN1610 switch CL1 to disable the static mode.

This configuration matches with the ISL configuration on the 5596 switch C2 when the ISLs are brought up on both switches in step [10](#) on page 6.

### Example

The following example shows the configuration of the ISL interface 3/1 using the `no port-channel static` command to make the ISL dynamic:

```
(CL1)# configure
(CL1)(Config)# interface 3/1
(CL1)(Interface 3/1)# no port-channel static
(CL1)(Interface 3/1)# exit
(CL1)(Config)# exit
(CL1)#
```

13. Bring up ISLs 13 through 16 on the active CN1610 switch CL1.

### Example

The following example illustrates the process of bringing up ISLs 13 through 16 on the interface 3/1:

```
(CL1)#
(CL1)# configure
(CL1)(Config)# interface 0/13-0/16,3/1
(CL1)(Interface 0/13-0/16,3/1)# no shutdown
(CL1)(Interface 0/13-0/16,3/1)# exit
(CL1)(Config)# exit
(CL1)#
```

14. Verify that the ISLs are up on the CN1610 switch CL1: `show port-channel`

### Example

```
(CL1)# show port-channel 3/1
Local Interface..... 3/1
Channel Name..... ISL-LAG
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr   Device/      Port      Port
Ports Timeout    Speed     Active
-----
0/13  actor/long    10G Full  True
      partner/long
0/14  actor/long    10G Full  True
      partner/long
0/15  actor/long    10G Full  True
      partner/long
0/16  actor/long    10G Full  True
      partner/long
```

15. Verify that the ISLs are up on the 5596 switch C2: `show port-channel summary`

### Example

```
C2# 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(D)  Eth1/42(D)  Eth1/43(D)
              Eth1/44(D)  Eth1/45(P)  Eth1/46(P)
              Eth1/47(P)  Eth1/48(P)
```

16. Bring up ports 1 through 40 on the 5596 switch C2 that are associated with the cluster nodes: **no shutdown**

**Example**

The following example shows how to bring up ports 1 through 40 on the 5596 switch C2:

```
C2#
C2# configure
C2(config)# interface ethernet 1/1-40
C2(config-if-range)# no shutdown
C2(config-if-range)# exit
C2(config)# exit
C2#
```

17. Bring up all of the cluster interconnect ports that are connected to the 5596 switch C2 on all of the nodes: **network port modify**

**Example**

The following example shows how to bring up the specified ports on nodes 1 and 3:

```
cluster::*> network port modify -node node1 -port e0b -up-admin true
cluster::*> network port modify -node node3 -port e0b -up-admin true
cluster::*> network port modify -node node3 -port e0c -up-admin true
```

18. Revert all of the cluster interconnect LIFs that are connected to C2 on all of the nodes: **network interface revert**

**Example**

The following example shows how to revert the LIFs on clus2 and clus3 to their home ports:

```
cluster::*> network interface revert -vserver node1 -lif clus2
cluster::*> network interface revert -vserver node3 -lif clus2
cluster::*> network interface revert -vserver node3 -lif clus3
```

19. Verify that all of the cluster interconnect ports are reverted to their home ports: **network interface show**

**Example**

The following example shows that the LIFs on clus2 are reverted to their home ports, and shows that the LIFs are successfully reverted if the ports in the “Current Port” column have a status of **true** in the “Is Home” column. If the Is Home value is **false**, then the LIF is not reverted.

```
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	e0a	true
	clus2	up/up	10.10.0.2/24	node1	e0b	true
node3	clus1	up/up	10.10.0.5/24	node3	e0a	true
	clus2	up/up	10.10.0.6/24	node3	e0b	true
	clus3	up/up	10.10.0.7/24	node3	e0c	true
	clus4	up/up	10.10.0.8/24	node3	e0d	true



20. Verify that the port is connected: `network port show`

**Example**

The following example shows the result of the previous `network port modify` command, verifying that all of the cluster interconnects are `up`:

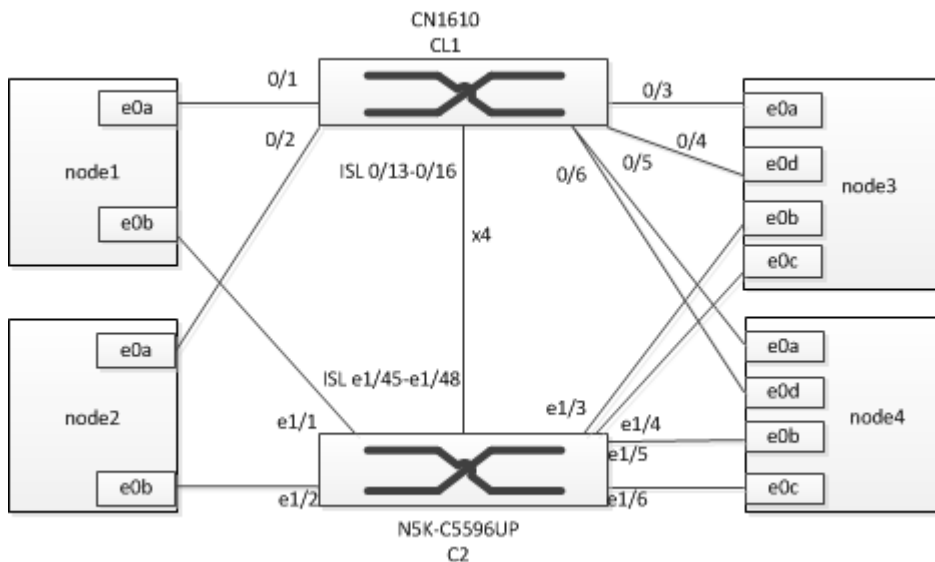
```
cluster::*> network port show -role cluster
```

Node	Port	Role	Link	MTU	Auto-Begot Admin/Open	Duplex Admin/Open	Speed (Mbps) Admin/Open
-----							
node1	e0a	cluster	up	9000	true/true	full/full	auto/10000
	e0b	cluster	up	9000	true/true	full/full	auto/10000
node2	e0a	cluster	up	9000	true/true	full/full	auto/10000
	e0b	cluster	up	9000	true/true	full/full	auto/10000
node3	e0a	cluster	up	9000	true/true	full/full	auto/10000
	e0b	cluster	up	9000	true/true	full/full	auto/10000
	e0c	cluster	up	9000	true/true	full/full	auto/10000
	e0d	cluster	up	9000	true/true	full/full	auto/10000
node4	e0a	cluster	up	9000	true/true	full/full	auto/10000
	e0b	cluster	up	9000	true/true	full/full	auto/10000
	e0c	cluster	up	9000	true/true	full/full	auto/10000
	e0d	cluster	up	9000	true/true	full/full	auto/10000

21. Ping the remote cluster interfaces and perform an RPC server check: `cluster ping-cluster`

**Example**

The following figure displays the cluster network topology after a CN1610 switch is replaced with a Nexus 5596 switch:



22. To begin replacing the first CN1610 switch with the first 5596 switch C1, use the `shutdown` command to shut down ports 1 through 40 on C1.

## Example

```
C1# configure
C1(config)# interface ethernet 1/1-40
C1(config-if-range)# shutdown
C1(config-if-range)# exit
C1(config)# exit
C1#
```

23. On each node in the cluster, migrate the interfaces that are associated with the first CN1610 switch CL1, to be replaced: **network interface migrate**

## Example

The following example migrates the LIF clus1 on an SVM named node1 to the specified port on node1, and LIFs clus1 and clus4 on an SVM named node3 to the specified ports on node3:

```
cluster::*> network interface migrate -vserver node1 -lif clus1 -source-node node1 -dest-node node1 -
dest-port e0b
cluster::*> network interface migrate -vserver node3 -lif clus1 -source-node node3 -dest-node node3 -
dest-port e0b
cluster::*> network interface migrate -vserver node3 -lif clus4 -source-node node3 -dest-node node3 -
dest-port e0c
```

24. Verify the cluster's status: **network interface show**

## Example

The following example shows that the LIF clus1 on an SVM named node1 was migrated to the specified port on node1:

```
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	e0b	false
	clus2	up/up	10.10.0.2/24	node1	e0b	true

25. Shut down the node ports that are connected to CL1 on all of the nodes: **network port modify**

## Example

The following example shows how to shut down the specified ports on nodes - node1 and node3:

```
cluster::*> network port modify -node node1 -port e0a -up-admin false
cluster::*> network port modify -node node3 -port e0a -up-admin false
cluster::*> network port modify -node node3 -port e0d -up-admin false
```

26. Ping the remote cluster interfaces, and perform an RPC server check: **cluster ping-cluster**

## Example

The following example shows how to ping the remote cluster interfaces:

```
cluster::*> cluster ping-cluster -node node1
Host is node1
Getting addresses from network interface table...
Local = 10.10.0.1          10.10.0.2
Remote = 10.10.0.3 10.10.0.4 10.10.0.5 10.10.0.6 10.10.0.7 10.10.0.8 10.10.0.9 10.10.0.10
10.10.0.11          10.10.0.12
Ping status:
....
```

```

Basic connectivity succeeds on 20 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 20 path(s):
  Local 10.10.0.1 to Remote 10.10.0.3
  Local 10.10.0.1 to Remote 10.10.0.4
  Local 10.10.0.1 to Remote 10.10.0.5
  Local 10.10.0.1 to Remote 10.10.0.6
  Local 10.10.0.1 to Remote 10.10.0.7
  Local 10.10.0.1 to Remote 10.10.0.8
  Local 10.10.0.1 to Remote 10.10.0.9
  Local 10.10.0.1 to Remote 10.10.0.10
  Local 10.10.0.1 to Remote 10.10.0.11
  Local 10.10.0.1 to Remote 10.10.0.12
  Local 10.10.0.2 to Remote 10.10.0.3
  Local 10.10.0.2 to Remote 10.10.0.4
  Local 10.10.0.2 to Remote 10.10.0.5
  Local 10.10.0.2 to Remote 10.10.0.6
  Local 10.10.0.2 to Remote 10.10.0.7
  Local 10.10.0.2 to Remote 10.10.0.8
  Local 10.10.0.2 to Remote 10.10.0.9
  Local 10.10.0.2 to Remote 10.10.0.10
  Local 10.10.0.2 to Remote 10.10.0.11
  Local 10.10.0.2 to Remote 10.10.0.12

Larger than PMTU communication succeeds on 20 path(s)
RPC status:
10 paths up, 0 paths down (tcp check)
10 paths up, 0 paths down (udp check)

```

27. Shut down ISLs 41 through 48 on the active 5596 switch C2: **shutdown**

#### Example

The following example shows how to shut down ISLs 41 through 48 on the 5596 switch C2:

```

C2# configure
C2(config)# interface ethernet 1/41-48
C2(config-if-range)# shutdown
C2(config-if-range)# exit
C2(config)# exit
C2#

```

**Note:** Note that even though there are only four physical connections at this stage, there will be eight by the time the new 5596 switch C1 is in place so that all eight connections will be used.

28. Remove the cables from the CN1610 switch CL1, and attach them to the first Cisco 5596 switch C1.

#### Example

For the ISL ports, connect all of the eight ports 41 through 48 to the second 5596 switch C2.

**Attention:** When reconnecting any cables to the new Cisco 5596 switch, the cables used must now be either fiber cables or Cisco twinax cables. You can see the “Before you begin” section for cabling caveats.

29. Bring up ISLs 41 through 48 on C2, the active 5596 switch: **no shutdown**

#### Example

The following example shows how to bring up ISLs 41 through 48 on the 5596 switch C2:

```
C2# configure
C2(config)# interface ethernet 1/41-48
C2(config-if-range)# no shutdown
```

30. Verify that the ISL connections are **up** on the 5596 switch C2.

#### Example

The following example verifies that the ISL connections are **up** on switch C2:

```
C2# 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)
```

31. Bring up ports 1 through 40 on the 5596 switch C1.

#### Example

The following example shows how to bring up ports 1 through 40 on the 5596 switch C1:

```
C1# configure
C1(config)# interface ethernet 1/1-40
C1(config-if-range)# no shutdown
C1(config-if-range)# exit
C1(config)# exit
C1#
```

32. Bring up all of the cluster interconnect ports connected to the new 5596 switch C1 on all of the nodes: **network port modify**

#### Example

The following example shows how to bring up all of the cluster interconnect ports for node1 and node3 on the 5596 switch C1:

```
cluster::*> network port modify -node node1 -port e0a -up-admin true
cluster::*> network port modify -node node3 -port e0a -up-admin true
cluster::*> network port modify -node node3 -port e0d -up-admin true
```

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

#### Example

The following example verifies that all of the cluster interconnect ports on node1 and node3 on the new 5596 switch C1 are **up**:

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

Node  Port  Role  Link  MTU  Auto-Begot  Duplex  Speed (Mbps)
Admin/Open  Admin/Open  Admin/Open
```

```

-----
node1
  e0a  cluster  up    9000  true/true  full/full  auto/10000
  e0b  cluster  up    9000  true/true  full/full  auto/10000
node3
  e0a  cluster  up    9000  true/true  full/full  auto/10000
  e0b  cluster  up    9000  true/true  full/full  auto/10000
  e0c  cluster  up    9000  true/true  full/full  auto/10000
  e0d  cluster  up    9000  true/true  full/full  auto/10000
-----

```

34. Revert the specified cluster LIFs clus1 and clus4 to their home ports on all of the nodes: **network interface revert**

#### Example

The following example shows how to revert the specified cluster LIFs to their home ports:

```

cluster::*> network interface revert -vserver node1 -lif clus1
cluster::*> network interface revert -vserver node3 -lif clus1
cluster::*> network interface revert -vserver node3 -lif clus4

```

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

#### Example

The following example shows that the status of the cluster interconnect interfaces is **up** and is home for node1 and node3:

```

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	e0a	true
	clus2	up/up	10.10.0.2/24	node1	e0b	true
node3	clus1	up/up	10.10.0.5/24	node1	e0a	true
	clus2	up/up	10.10.0.6/24	node1	e0b	true
	clus3	up/up	10.10.0.7/24	node1	e0c	true
	clus4	up/up	10.10.0.8/24	node1	e0d	true

36. Ping the remote cluster interfaces and perform an RPC server check: **cluster ping-cluster**

#### Example

The following example shows how to ping the remote cluster interfaces:

```

cluster::*> cluster ping-cluster -node node1
Host is node1
Getting addresses from network interface table...
Local = 10.10.0.1      10.10.0.2
Remote = 10.10.0.3 10.10.0.4 10.10.0.5 10.10.0.6 10.10.0.7 10.10.0.8 10.10.0.9 10.10.0.10
10.10.0.11      10.10.0.12
Ping status:
....
Basic connectivity succeeds on 20 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 20 path(s):
  Local 10.10.0.1 to Remote 10.10.0.3
  Local 10.10.0.1 to Remote 10.10.0.4
  Local 10.10.0.1 to Remote 10.10.0.5
  Local 10.10.0.1 to Remote 10.10.0.6
  Local 10.10.0.1 to Remote 10.10.0.7
  Local 10.10.0.1 to Remote 10.10.0.8
  Local 10.10.0.1 to Remote 10.10.0.9
  Local 10.10.0.1 to Remote 10.10.0.10
  Local 10.10.0.1 to Remote 10.10.0.11
  Local 10.10.0.1 to Remote 10.10.0.12
  Local 10.10.0.2 to Remote 10.10.0.3

```

```

Local 10.10.0.2 to Remote 10.10.0.4
Local 10.10.0.2 to Remote 10.10.0.5
Local 10.10.0.2 to Remote 10.10.0.6
Local 10.10.0.2 to Remote 10.10.0.7
Local 10.10.0.2 to Remote 10.10.0.8
Local 10.10.0.2 to Remote 10.10.0.9
Local 10.10.0.2 to Remote 10.10.0.10
Local 10.10.0.2 to Remote 10.10.0.11
Local 10.10.0.2 to Remote 10.10.0.12

```

```

Larger than PMTU communication succeeds on 20 path(s)
RPC status:
10 paths up, 0 paths down (tcp check)
10 paths up, 0 paths down (udp check)

```

37. Use the following commands to display information about the devices in your configuration:

- network device-discovery show
- network port show -role cluster
- network interface show -role cluster
- system cluster-switch show

### Example

```

cluster::*> network device-discovery show

```

Node	Local Port	Discovered Device	Interface	Platform
node1	e0a	C1	Ethernet1/1	N5K-C5596UP
	e0b	C2	Ethernet1/1	N5K-C5596UP
node2	e0a	C1	Ethernet1/2	N5K-C5596UP
	e0b	C2	Ethernet1/2	N5K-C5596UP
node3	e0a	C1	Ethernet1/3	N5K-C5596UP
	e0b	C2	Ethernet1/3	N5K-C5596UP
	e0c	C2	Ethernet1/4	N5K-C5596UP
	e0d	C1	Ethernet1/4	N5K-C5596UP
node4	e0a	C1	Ethernet1/5	N5K-C5596UP
	e0b	C2	Ethernet1/5	N5K-C5596UP
	e0c	C2	Ethernet1/6	N5K-C5596UP
	e0d	C1	Ethernet1/6	N5K-C5596UP

12 entries were displayed

```

cluster::*> network port show -role cluster

```

Node	Port	Role	Link	MTU	Auto-Begot Admin/Open	Duplex Admin/Open	Speed (Mbps) Admin/Open
node1	e0a	cluster	up	9000	true/true	full/full	auto/10000
	e0b	cluster	up	9000	true/true	full/full	auto/10000
node2	e0a	cluster	up	9000	true/true	full/full	auto/10000
	e0b	cluster	up	9000	true/true	full/full	auto/10000
node3	e0a	cluster	up	9000	true/true	full/full	auto/10000
	e0b	cluster	up	9000	true/true	full/full	auto/10000
	e0c	cluster	up	9000	true/true	full/full	auto/10000
	e0d	cluster	up	9000	true/true	full/full	auto/10000
node4	e0a	cluster	up	9000	true/true	full/full	auto/10000
	e0b	cluster	up	9000	true/true	full/full	auto/10000
	e0c	cluster	up	9000	true/true	full/full	auto/10000
	e0d	cluster	up	9000	true/true	full/full	auto/10000

```
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	e0a	true
	clus2	up/up	10.10.0.2/24	node1	e0b	true
node2	clus1	up/up	10.10.0.3/24	node2	e0a	true
	clus2	up/up	10.10.0.4/24	node2	e0b	true
node3	clus1	up/up	10.10.0.5/24	node3	e0a	true
	clus2	up/up	10.10.0.6/24	node3	e0b	true
	clus3	up/up	10.10.0.7/24	node3	e0c	true
	clus4	up/up	10.10.0.8/24	node3	e0d	true
node4	clus1	up/up	10.10.0.9/24	node4	e0a	true
	clus2	up/up	10.10.0.10/24	node4	e0b	true
	clus3	up/up	10.10.0.11/24	node4	e0c	true
	clus4	up/up	10.10.0.12/24	node4	e0d	true

12 entries were displayed.

```
cluster::> system cluster-switch show
```

Switch	Type	Address	Model
C1	cluster-network	10.10.0.103	NX5596
Serial Number: FOX000001			
Is Monitored: true			
Reason:			
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version 7.1(1)N1(1)			
Version Source: CDP			
C2	cluster-network	10.10.0.104	NX5596
Serial Number: FOX000002			
Is Monitored: true			
Reason:			
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version 7.1(1)N1(1)			
Version Source: CDP			
CL1	cluster-network	10.10.0.101	CN1610
Serial Number: 01234567			
Is Monitored: true			
Reason:			
Software Version: 1.2.0.7			
Version Source: ISDP			
CL2	cluster-network	10.10.0.102	CN1610
Serial Number: 01234568			
Is Monitored: true			
Reason:			
Software Version: 1.2.0.7			
Version Source: ISDP			

4 entries were displayed.

38. Remove the replaced CN1610 switch: **system cluster-switch delete**

### Example

```
cluster::> system cluster-switch delete -device CL1
cluster::> system cluster-switch delete -device CL2
```

39. Verify that the proper cluster switches are being monitored.


## Example

```
cluster::> system cluster-switch show
```

Switch	Type	Address	Model
C1	cluster-network	10.10.0.103	NX5596
Serial Number: FOX000001			
Is Monitored: true			
Reason:			
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version 7.1(1)N1(1)			
Version Source: CDP			
C2	cluster-network	10.10.0.104	NX5596
Serial Number: FOX000002			
Is Monitored: true			
Reason:			
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version 7.1(1)N1(1)			
Version Source: CDP			

2 entries were displayed.

The following figure displays the cluster network topology after both the CN1610 switches are replaced by two Nexus 5596 switches:



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

## Example

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

## Related information

[NetApp CN1601 and CN1610 description page](#)  
[Cisco Ethernet Switch description page](#)  
[Hardware Universe](#)



## 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](mailto: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., 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, 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>