



Cisco® Nexus Switches

Migrating from a Cisco Nexus 5596 Switch to a Cisco Nexus 3232C Switch

September 2020 | 215-13129_B0
doccomments@netapp.com

Contents

Migrating from a Cisco Nexus 5596 switch to a Cisco Nexus 3232C switch.....	3
How to migrate from a Cisco Nexus 5596 cluster switch to a Cisco Nexus 3232C cluster switch.....	4
Copyright and trademark.....	18
Copyright.....	18
Trademark.....	18
Index.....	19

Migrating from a Cisco Nexus 5596 switch to a Cisco Nexus 3232C switch

You must be aware of certain configuration information, port connections and cabling requirements when you are replacing Cisco Nexus 5596 cluster switches with Cisco Nexus 3232C cluster switches.

- The following cluster switches are used as examples in this procedure:
 - Nexus 5596
 - Nexus 3232C
- The cluster switches use the following ports for connections to nodes:
 - Ports e1/1-40 (10 GbE): Nexus 5596
 - Ports e1/1-30 (10/40/100 GbE): Nexus 3232C
- The cluster switches use the following Inter-Switch Link (ISL) ports:
 - Ports e1/41-48 (10 GbE): Nexus 5596
 - Ports e1/31-32 (40/100 GbE): Nexus 3232C
- The *Hardware Universe* contains information about supported cabling to Nexus 3232C switches:
 - Nodes with 10 GbE cluster connections require QSFP to SFP+ optical fiber breakout cables or QSFP to SFP+ copper breakout cables.
 - Nodes with 40/100 GbE cluster connections require supported QSFP/QSFP28 optical modules with fiber cables or QSFP/QSFP28 copper direct-attach cables.
- The cluster switches use the appropriate ISL cabling:
 - Beginning: Nexus 5596 (SFP+ to SFP+)
 - 8x SFP+ fiber or copper direct-attach cables
 - Interim: Nexus 5596 to Nexus 3232C (QSFP to 4xSFP+ break-out)
 - 1x QSFP to SFP+ fiber break-out or copper break-out cables
 - Final: Nexus 3232C to Nexus 3232C (QSFP28 to QSFP28)
 - 2x QSFP28 fiber or copper direct-attach cables
- On Nexus 3232C switches, you can operate QSFP/QSFP28 ports in either 40/100 Gigabit Ethernet or 4 x10 Gigabit Ethernet modes.

By default, there are 32 ports in the 40/100 Gigabit Ethernet mode. These 40 Gigabit Ethernet ports are numbered in a 2-tuple naming convention. For example, the second 40 Gigabit Ethernet port is numbered as 1/2. The process of changing the configuration from 40 Gigabit Ethernet to 10 Gigabit Ethernet is called *breakout* and the process of changing the configuration from 10 Gigabit Ethernet to 40 Gigabit Ethernet is called *breakin*. When you break out a 40/100 Gigabit Ethernet port into 10 Gigabit Ethernet ports, the resulting ports are numbered using a 3-tuple naming convention. For example, the break-out ports of the second 40/100 Gigabit Ethernet port are numbered as 1/2/1, 1/2/2, 1/2/3, and 1/2/4.
- On the left side of Nexus 3232C switches are 2 SFP+ ports, called 1/33 and 1/34.
- You have configured some of the ports on Nexus 3232C switches to run at 10 GbE or 40/100 GbE.

Note: You can break out the first six ports into 4x10 GbE mode by using the `interface breakout module 1 port 1-6 map 10g-4x` command. Similarly, you can regroup the first six QSFP+ ports from breakout configuration by using the `no interface breakout module 1 port 1-6 map 10g-4x` command.

- You have done the planning, migration, and read the required documentation on 10 GbE and 40/100 GbE connectivity from nodes to Nexus 3232C cluster switches.
- The ONTAP and NX-OS versions supported in this procedure are on the *Cisco Ethernet Switches* page.
[Cisco Ethernet Switches](#)

How to migrate from a Cisco Nexus 5596 cluster switch to a Cisco Nexus 3232C cluster switch

To replace existing Cisco Nexus 5596 cluster switches in a cluster with Nexus 3232C cluster switches, you must perform a specific sequence of tasks.

About this task

The examples in this procedure describe replacing Cisco Nexus 5596 switches with Cisco Nexus 3232C switches. You can use these steps (with modifications) for other older Cisco switches (for example, 3132Q-V). The procedure also uses the following switch and node nomenclature:

- The command outputs might vary depending on different releases of ONTAP.
- The Nexus 5596 switches to be replaced are CL1 and CL2.
- The Nexus 3232C switches to replace the Nexus 5596 switches are C1 and C2.
- n1_clus1 is the first cluster logical interface (LIF) connected to cluster switch 1 (CL1 or C1) for node n1.
- n1_clus2 is the first cluster LIF connected to cluster switch 2 (CL2 or C2) for node n1.
- n1_clus3 is the second LIF connected to cluster switch 2 (CL2 or C2) for node n1.
- n1_clus4 is the second LIF connected to cluster switch 1 (CL1 or C1) for node n1.
- The number of 10 GbE and 40/100 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.
- The nodes are n1, n2, n3, and n4.

Note: The examples in this procedure use four nodes: Two nodes use four 10 GbE cluster interconnect ports: e0a, e0b, e0c, and e0d. The other two nodes use two 40 GbE cluster interconnect ports: e4a, e4e. The *Hardware Universe* lists the actual cluster ports on your platforms.

This procedure covers the following scenarios:

- The cluster starts with two nodes connected and functioning in a two Nexus 5596 cluster switches.
- The cluster switch CL2 to be replaced by C2 (steps 1 to 19):
 - Traffic on all cluster ports and LIFs on all nodes connected to CL2 are migrated onto the first cluster ports and LIFs connected to CL1.
 - Disconnect cabling from all cluster ports on all nodes connected to CL2, and then use supported break-out cabling to reconnect the ports to new cluster switch C2.
 - Disconnect cabling between ISL ports between CL1 and CL2, and then use supported break-out cabling to reconnect the ports from CL1 to C2.
 - Traffic on all cluster ports and LIFs connected to C2 on all nodes is reverted.
- The cluster switch CL1 to be replaced by C1 (steps 20 to 33):
 - Traffic on all cluster ports or LIFs on all nodes connected to CL1 are migrated onto the second cluster ports or LIFs connected to C2.
 - Disconnect cabling from all cluster port on all nodes connected to CL1 and reconnect, using supported break-out cabling, to new cluster switch C1.

- Disconnect cabling between ISL ports between CL1 and C2, and reconnect using supported cabling, from C1 to C2.
- Traffic on all cluster ports or LIFs connected to C1 on all nodes is reverted.
- Two FAS9000 nodes have been added to cluster with examples showing cluster details (steps 34 to 37).

Note: The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

Steps

1. 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 AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Display information about the devices in your configuration:

```
network device-discovery show
```

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

```
cluster::> network device-discovery show
Node      Local Port   Discovered Device      Interface      Platform
-----
n1        /cdp
          e0a      CL1          Ethernet1/1    N5K-C5596UP
          e0b      CL2          Ethernet1/1    N5K-C5596UP
          e0c      CL2          Ethernet1/2    N5K-C5596UP
          e0d      CL1          Ethernet1/2    N5K-C5596UP
n2        /cdp
          e0a      CL1          Ethernet1/3    N5K-C5596UP
          e0b      CL2          Ethernet1/3    N5K-C5596UP
          e0c      CL2          Ethernet1/4    N5K-C5596UP
          e0d      CL1          Ethernet1/4    N5K-C5596UP
8 entries were displayed.
```

3. Determine the administrative or operational status for each cluster interface:

- a. Display the network port attributes:

```
network port show -role cluster
```

The following example displays the network port attributes on nodes n1 and n2:

```
cluster::*> network port show -role cluster
(network port show)
Node: n1
Port      IPspace      Broadcast Domain  Link  MTU      Speed(Mbps)  Health  Ignore
Admin/Oper  Status
-----
e0a      Cluster      Cluster          up    9000    auto/10000  -      -
e0b      Cluster      Cluster          up    9000    auto/10000  -      -
e0c      Cluster      Cluster          up    9000    auto/10000  -      -
e0d      Cluster      Cluster          up    9000    auto/10000  -      -
Node: n2
Port      IPspace      Broadcast Domain  Link  MTU      Speed(Mbps)  Health  Ignore
Admin/Oper  Status
-----
e0a      Cluster      Cluster          up    9000    auto/10000  -      -
e0b      Cluster      Cluster          up    9000    auto/10000  -      -
e0c      Cluster      Cluster          up    9000    auto/10000  -      -
```

```
e0d Cluster Cluster up 9000 auto/10000 - -
8 entries were displayed.
```

- b. Display information about the logical interfaces:

network interface show -role cluster

The following example displays the general information about all of the LIFs on the cluster, including their current ports:

```
cluster::*> network interface show -role cluster
(network interface show)
Vserver Logical Status Network Current Current Is
Interface Admin/Oper Address/Mask Node Port Home
-----
Cluster
n1_clus1 up/up 10.10.0.1/24 n1 e0a true
n1_clus2 up/up 10.10.0.2/24 n1 e0b true
n1_clus3 up/up 10.10.0.3/24 n1 e0c true
n1_clus4 up/up 10.10.0.4/24 n1 e0d true
n2_clus1 up/up 10.10.0.5/24 n2 e0a true
n2_clus2 up/up 10.10.0.6/24 n2 e0b true
n2_clus3 up/up 10.10.0.7/24 n2 e0c true
n2_clus4 up/up 10.10.0.8/24 n2 e0d true
8 entries were displayed.
```

- c. Display information about the discovered cluster switches:

system cluster-switch show

The following example shows the active cluster switches:

```
cluster::*> system cluster-switch show
Switch Type Address Model
-----
CL1 cluster-network 10.10.1.101 NX5596
Serial Number: 01234567
Is Monitored: true
Reason:
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version
7.1(1)N1(1)
Version Source: CDP
CL2 cluster-network 10.10.1.102 NX5596
Serial Number: 01234568
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.
```

4. Verify that the appropriate RCF and image are installed on the new 3232C switches as necessary for your requirements, and make the essential site customizations, such as users and passwords, network addresses, and other customizations.

Note: You must prepare both switches at this time.

If you need to upgrade the RCF and image, you must complete the following steps:

- a. Go to the *Cisco Ethernet Switches* page on the NetApp Support Site.
[Cisco Ethernet Switches](#)
- b. Note your switch and the required software versions in the table on that page.
- c. Download the appropriate version of the RCF.
- d. Click **CONTINUE** on the **Description** page, accept the license agreement, and then follow the instructions on the **Download** page to download the RCF.
- e. Download the appropriate version of the image software.

See the *ONTAP 8.x or later Cluster and Management Network Switch Reference Configuration Files* Download page, and then click the appropriate version.

To find the correct version, see the *ONTAP 8.x or later Cluster Network Switch Download page*.

- 5. Migrate the LIFs associated with the second Nexus 5596 switch to be replaced:

```
network interface migrate -vserver Cluster -lif lif-name -source-node source-node-name - destination-node node-name -destination-port destination-port-name
```

The following example shows the LIFs being migrated for nodes n1 and n2; LIF migration must be done on all of the nodes:

```
cluster::*> network interface migrate -vserver Cluster -lif n1_clus2 -source-node n1 - destination-node n1 -destination-port e0a
cluster::*> network interface migrate -vserver Cluster -lif n1_clus3 -source-node n1 - destination-node n1 -destination-port e0d
cluster::*> network interface migrate -vserver Cluster -lif n2_clus2 -source-node n2 - destination-node n2 -destination-port e0a
cluster::*> network interface migrate -vserver Cluster -lif n2_clus3 -source-node n2 - destination-node n2 -destination-port e0d
```

- 6. Verify the cluster's health:

```
network interface show -role cluster
```

The following example shows the current status of each cluster:

```
cluster::*> network interface show -role cluster
(network interface show)
Vserver      Logical      Status      Network      Current      Current      Is
Interface    Admin/Oper   Address/Mask Node          Port         Home
-----
Cluster
      n1_clus1   up/up       10.10.0.1/24 n1           e0a          true
      n1_clus2   up/up       10.10.0.2/24 n1           e0a          false
      n1_clus3   up/up       10.10.0.3/24 n1           e0d          false
      n1_clus4   up/up       10.10.0.4/24 n1           e0d          true
      n2_clus1   up/up       10.10.0.5/24 n2           e0a          true
      n2_clus2   up/up       10.10.0.6/24 n2           e0a          false
      n2_clus3   up/up       10.10.0.7/24 n2           e0d          false
      n2_clus4   up/up       10.10.0.8/24 n2           e0d          true
8 entries were displayed.
```

- 7. Shut down the cluster interconnect ports that are physically connected to switch CL2:

```
network port modify -node node-name -port port-name -up-admin false
```

The following commands shut down the specified ports on n1 and n2, but the ports must be shut down on all nodes:

```
cluster::*> network port modify -node n1 -port e0b -up-admin false
cluster::*> network port modify -node n1 -port e0c -up-admin false
cluster::*> network port modify -node n2 -port e0b -up-admin false
cluster::*> network port modify -node n2 -port e0c -up-admin false
```

- 8. Ping the remote cluster interfaces and perform an RPC server check:

```
cluster ping-cluster -node node-name
```

The following example shows node n1 being pinged and the RPC status indicated afterward:

```
cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1      e0a  10.10.0.1
Cluster n1_clus2 n1      e0b  10.10.0.2
Cluster n1_clus3 n1      e0c  10.10.0.3
Cluster n1_clus4 n1      e0d  10.10.0.4
Cluster n2_clus1 n2      e0a  10.10.0.5
Cluster n2_clus2 n2      e0b  10.10.0.6
Cluster n2_clus3 n2      e0c  10.10.0.7
Cluster n2_clus4 n2      e0d  10.10.0.8

Local = 10.10.0.1 10.10.0.2 10.10.0.3 10.10.0.4
Remote = 10.10.0.5 10.10.0.6 10.10.0.7 10.10.0.8
Cluster Vserver Id = 4294967293
```

```
Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 16 path(s):
  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 10.10.0.2 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 10.10.0.3 to Remote 10.10.0.5
  Local 10.10.0.3 to Remote 10.10.0.6
  Local 10.10.0.3 to Remote 10.10.0.7
  Local 10.10.0.3 to Remote 10.10.0.8
  Local 10.10.0.4 to Remote 10.10.0.5
  Local 10.10.0.4 to Remote 10.10.0.6
  Local 10.10.0.4 to Remote 10.10.0.7
  Local 10.10.0.4 to Remote 10.10.0.8
Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
4 paths up, 0 paths down (udp check)
```

9. Shut down ISLs 41 through 48 on CL1, the active Nexus 5596 switch using the Cisco shutdown command.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-OS Command References](#).

The following example shows ISLs 41 through 48 being shut down on the Nexus 5596 switch CL1:

```
(CL1)# configure
(CL1)(Config)# interface e1/41-48
(CL1)(config-if-range)# shutdown
(CL1)(config-if-range)# exit
(CL1)(Config)# exit
(CL1)#
```

10. Build a temporary ISL between CL1 and C2 using the appropriate Cisco commands.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-OS Command References](#).

The following example shows a temporary ISL being set up between CL1 and C2:

```
C2# configure
C2(config)# interface port-channel 2
C2(config-if)# switchport mode trunk
C2(config-if)# spanning-tree port type network
C2(config-if)# mtu 9216
C2(config-if)# interface breakout module 1 port 24 map 10g-4x
C2(config)# interface e1/24/1-4
C2(config-if-range)# switchport mode trunk
C2(config-if-range)# mtu 9216
C2(config-if-range)# channel-group 2 mode active
C2(config-if-range)# exit
C2(config-if)# exit
```

11. On all nodes, remove all cables attached to the Nexus 5596 switch CL2.

With supported cabling, reconnect disconnected ports on all nodes to the Nexus 3232C switch C2.

12. Remove all the cables from the Nexus 5596 switch CL2.

Attach the appropriate Cisco QSFP to SFP+ break-out cables connecting port 1/24 on the new Cisco 3232C switch, C2, to ports 45 to 48 on existing Nexus 5596, CL1.

13. Bring up ISLs ports 45 through 48 on the active Nexus 5596 switch CL1.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-OS Command References](#).

The following example shows ISLs ports 45 through 48 being brought up:

```
(CL1)# configure
(CL1)(Config)# interface e1/45-48
(CL1)(config-if-range)# no shutdown
(CL1)(config-if-range)# exit
(CL1)(Config)# exit
(CL1)#
```

14. Verify that the ISLs are up on the Nexus 5596 switch CL1.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-OS Command References](#).

The following example shows Ports eth1/45 through eth1/48 indicating (P), meaning that the ISL ports are up in the port-channel.

```
CL1# 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(D) Eth1/42(D) Eth1/43(D)
                               Eth1/44(D) Eth1/45(P) Eth1/46(P)
                               Eth1/47(P) Eth1/48(P)
```

15. Verify that interfaces eth1/45-48 already have channel-group 1 mode active in their running configuration.

16. On all nodes, bring up all the cluster interconnect ports connected to the 3232C switch C2:

network port modify -node node-name -port port-name -up-admin true

The following example shows the specified ports being brought up on nodes n1 and n2:

```
cluster::*> network port modify -node n1 -port e0b -up-admin true
cluster::*> network port modify -node n1 -port e0c -up-admin true
cluster::*> network port modify -node n2 -port e0b -up-admin true
cluster::*> network port modify -node n2 -port e0c -up-admin true
```

17. On all nodes, revert all of the migrated cluster interconnect LIFs connected to C2:

network interface revert -vserver Cluster -lif lif-name

The following example shows the migrated cluster LIFs being reverted to their home ports:

```
cluster::*> network interface revert -vserver Cluster -lif n1_clus2
cluster::*> network interface revert -vserver Cluster -lif n1_clus3
cluster::*> network interface revert -vserver Cluster -lif n2_clus2
cluster::*> network interface revert -vserver Cluster -lif n2_clus3
```

18. Verify all the cluster interconnect ports are now reverted to their home:

network interface show -role cluster

The following example shows that the LIFs on clus2 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**, the LIF has not been reverted.

```
cluster::*> network interface show -role cluster
(network interface show)
Vserver      Logical      Status      Network      Current      Current      Is
-----      -
Interface    Admin/Oper   Address/Mask Node          Port         Home
-----
Cluster
n1_clus1     up/up        10.10.0.1/24 n1           e0a          true
n1_clus2     up/up        10.10.0.2/24 n1           e0b          true
n1_clus3     up/up        10.10.0.3/24 n1           e0c          true
n1_clus4     up/up        10.10.0.4/24 n1           e0d          true
n2_clus1     up/up        10.10.0.5/24 n2           e0a          true
n2_clus2     up/up        10.10.0.6/24 n2           e0b          true
n2_clus3     up/up        10.10.0.7/24 n2           e0c          true
n2_clus4     up/up        10.10.0.8/24 n2           e0d          true
8 entries were displayed.
```

19. Verify that the clustered ports are connected:

network port show -role cluster

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

```
cluster::*> network port show -role cluster
(network port show)
Node: n1
Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health Ignore
-----      -
Admin/Oper  Status      Status
-----
e0a       Cluster      Cluster      up   9000 auto/10000 - -
e0b       Cluster      Cluster      up   9000 auto/10000 - -
e0c       Cluster      Cluster      up   9000 auto/10000 - -
e0d       Cluster      Cluster      up   9000 auto/10000 - -
Node: n2
Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health Ignore
-----      -
Admin/Oper  Status      Status
-----
e0a       Cluster      Cluster      up   9000 auto/10000 - -
e0b       Cluster      Cluster      up   9000 auto/10000 - -
e0c       Cluster      Cluster      up   9000 auto/10000 - -
e0d       Cluster      Cluster      up   9000 auto/10000 - -
8 entries were displayed.
```

20. Ping the remote cluster interfaces and perform an RPC server check:

cluster ping-cluster -node node-name

The following example shows node n1 being pinged and the RPC status indicated afterward:

```
cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1 e0a 10.10.0.1
Cluster n1_clus2 n1 e0b 10.10.0.2
Cluster n1_clus3 n1 e0c 10.10.0.3
Cluster n1_clus4 n1 e0d 10.10.0.4
Cluster n2_clus1 n2 e0a 10.10.0.5
Cluster n2_clus2 n2 e0b 10.10.0.6
Cluster n2_clus3 n2 e0c 10.10.0.7
Cluster n2_clus4 n2 e0d 10.10.0.8

Local = 10.10.0.1 10.10.0.2 10.10.0.3 10.10.0.4
Remote = 10.10.0.5 10.10.0.6 10.10.0.7 10.10.0.8
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
```

```
.....
Detected 1500 byte MTU on 16 path(s):
  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 10.10.0.2 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 10.10.0.3 to Remote 10.10.0.5
  Local 10.10.0.3 to Remote 10.10.0.6
  Local 10.10.0.3 to Remote 10.10.0.7
  Local 10.10.0.3 to Remote 10.10.0.8
  Local 10.10.0.4 to Remote 10.10.0.5
  Local 10.10.0.4 to Remote 10.10.0.6
  Local 10.10.0.4 to Remote 10.10.0.7
  Local 10.10.0.4 to Remote 10.10.0.8
Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
4 paths up, 0 paths down (udp check)
```

- 21. On each node in the cluster, migrate the interfaces associated with the first Nexus 5596 switch, CL1, to be replaced:

```
network interface migrate -vserver Cluster -lif lif-name -source-node source-node-name - destination-node destination-node-name -destination-port destination-port-name
```

The following example shows the ports or LIFs being migrated on nodes n1 and n2:

```
cluster::*> network interface migrate -vserver Cluster -lif n1_clus1 -source-node n1 -
destination-node n1 -destination-port e0b
cluster::*> network interface migrate -vserver Cluster -lif n1_clus4 -source-node n1 -
destination-node n1 -destination-port e0c
cluster::*> network interface migrate -vserver Cluster -lif n2_clus1 -source-node n2 -
destination-node n2 -destination-port e0b
cluster::*> network interface migrate -vserver Cluster -lif n2_clus4 -source-node n2 -
destination-node n2 -destination-port e0c
```

- 22. Verify the cluster's status:

```
network interface show
```

The following example shows that the required cluster LIFs have been migrated to appropriate cluster ports hosted on cluster switch, C2:

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

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
Cluster	n1_clus1	up/up	10.10.0.1/24	n1	e0b	false
	n1_clus2	up/up	10.10.0.2/24	n1	e0b	true
	n1_clus3	up/up	10.10.0.3/24	n1	e0c	true
	n1_clus4	up/up	10.10.0.4/24	n1	e0c	false
	n2_clus1	up/up	10.10.0.5/24	n2	e0b	false
	n2_clus2	up/up	10.10.0.6/24	n2	e0b	true
	n2_clus3	up/up	10.10.0.7/24	n2	e0c	true
	n2_clus4	up/up	10.10.0.8/24	n2	e0c	false

8 entries were displayed.

- 23. On all the nodes, shut down the node ports that are connected to CL1:

```
network port modify -node node-name -port port-name -up-admin false
```

The following example shows the specified ports being shut down on nodes n1 and n2:

```
cluster::*> network port modify -node n1 -port e0a -up-admin false
cluster::*> network port modify -node n1 -port e0d -up-admin false
```

```
cluster::*> network port modify -node n2 -port e0a -up-admin false  
cluster::*> network port modify -node n2 -port e0d -up-admin false
```

24. Shut down ISL 24, 31 and 32 on the active 3232C switch C2.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-OS Command References](#).

The following example shows ISLs being shutdown:

```
C2# configure  
C2(Config)# interface e1/24/1-4  
C2(config-if-range)# shutdown  
C2(config-if-range)# exit  
C2(config)# interface 1/31-32  
C2(config-if-range)# shutdown  
C2(config-if-range)# exit  
C2(config-if)# exit  
C2#
```

25. On all nodes, remove all cables attached to the Nexus 5596 switch CL1.

With supported cabling, reconnect disconnected ports on all nodes to the Nexus 3232C switch C1.

26. Remove the QSFP breakout cable from Nexus 3232C C2 ports e1/24.

Connect ports e1/31 and e1/32 on C1 to ports e1/31 and e1/32 on C2 using supported Cisco QSFP optical fiber or direct-attach cables.

27. Restore the configuration on port 24 and remove the temporary Port Channel 2 on C2.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-OS Command References](#).

The following example shows the configuration on port m24 being restored using the appropriate Cisco commands:

```
C2# configure  
C2(config)# no interface breakout module 1 port 24 map 10g-4x  
C2(config)# no interface port-channel 2  
C2(config-if)# int e1/24  
C2(config-if)# description 40GbE Node Port  
C2(config-if)# spanning-tree port type edge  
C2(config-if)# spanning-tree bpduguard enable  
C2(config-if)# mtu 9216  
C2(config-if-range)# exit  
C2(config)# exit  
C2# copy running-config startup-config  
[#####] 100%  
Copy Complete.
```

28. Bring up ISL ports 31 and 32 on C2, the active 3232C switch, by entering the following Cisco command:

```
no shutdown
```

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-OS Command References](#).

The following example shows the Cisco commands

```
switchname configure  
brought up on the 3232C switch C2:
```

```
C2# configure  
C2(config)# interface ethernet 1/31-32  
C2(config-if-range)# no shutdown
```

29. Verify that the ISL connections are up on the 3232C switch C2.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Ports eth1/31 and eth1/32 should indicate (P) meaning that both ISL ports up in the port-channel

```
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-          Type   Protocol  Member Ports
Channel
-----
1     Pol(SU)       Eth    LACP      Eth1/31(P) Eth1/32(P)
```

- On all nodes, bring up all the cluster interconnect ports connected to the new 3232C switch C1:

network port modify

The following example shows all the cluster interconnect ports being brought up for n1 and n2 on the 3232C switch C1:

```
cluster::*> network port modify -node n1 -port e0a -up-admin true
cluster::*> network port modify -node n1 -port e0d -up-admin true
cluster::*> network port modify -node n2 -port e0a -up-admin true
cluster::*> network port modify -node n2 -port e0d -up-admin true
```

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

The following example shows verifies that all cluster interconnect ports on all nodes on the new 3232C switch C1 are up:

```
cluster::*> network port show -role cluster
(network port show)
Node: n1
-----
Port      IPspace    Broadcast Domain Link MTU   Speed(Mbps) Health Ignore
Admin/Oper Status  Status
-----
e0a      Cluster    Cluster          up  9000 auto/10000 -      -
e0b      Cluster    Cluster          up  9000 auto/10000 -      -
e0c      Cluster    Cluster          up  9000 auto/10000 -      -
e0d      Cluster    Cluster          up  9000 auto/10000 -      -
Node: n2
-----
Port      IPspace    Broadcast Domain Link MTU   Speed(Mbps) Health Ignore
Admin/Oper Status  Status
-----
e0a      Cluster    Cluster          up  9000 auto/10000 -      -
e0b      Cluster    Cluster          up  9000 auto/10000 -      -
e0c      Cluster    Cluster          up  9000 auto/10000 -      -
e0d      Cluster    Cluster          up  9000 auto/10000 -      -
8 entries were displayed.
```

- On all nodes, revert the specific cluster LIFs to their home ports:

network interface revert -server Cluster -lif lif-name

The following example shows the specific cluster LIFs being reverted to their home ports on nodes n1 and n2:

```
cluster::*> network interface revert -vserver Cluster -lif n1_clus1
cluster::*> network interface revert -vserver Cluster -lif n1_clus4
```

```
cluster::*> network interface revert -vserver Cluster -lif n2_clus1
cluster::*> network interface revert -vserver Cluster -lif n2_clus4
```

33. Verify that the interface is home:

```
network interface show -role cluster
```

The following example shows the status of cluster interconnect interfaces are **up** and **Is Home** for n1 and n2:

```
cluster::*> network interface show -role cluster
(network interface show)
Vserver      Logical      Status      Network      Current      Current      Is
Interface    Admin/Oper   Address/Mask Node          Port         Home
-----
Cluster
n1_clus1     up/up       10.10.0.1/24 n1           e0a          true
n1_clus2     up/up       10.10.0.2/24 n1           e0b          true
n1_clus3     up/up       10.10.0.3/24 n1           e0c          true
n1_clus4     up/up       10.10.0.4/24 n1           e0d          true
n2_clus1     up/up       10.10.0.5/24 n2           e0a          true
n2_clus2     up/up       10.10.0.6/24 n2           e0b          true
n2_clus3     up/up       10.10.0.7/24 n2           e0c          true
n2_clus4     up/up       10.10.0.8/24 n2           e0d          true
8 entries were displayed.
```

34. Ping the remote cluster interfaces and perform an RPC server check:

```
cluster ping-cluster -node node-name
```

The following example shows node n1 being pinged and the RPC status indicated afterward:

```
cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1 e0a 10.10.0.1
Cluster n1_clus2 n1 e0b 10.10.0.2
Cluster n1_clus3 n1 e0c 10.10.0.3
Cluster n1_clus4 n1 e0d 10.10.0.4
Cluster n2_clus1 n2 e0a 10.10.0.5
Cluster n2_clus2 n2 e0b 10.10.0.6
Cluster n2_clus3 n2 e0c 10.10.0.7
Cluster n2_clus4 n2 e0d 10.10.0.8

Local = 10.10.0.1 10.10.0.2 10.10.0.3 10.10.0.4
Remote = 10.10.0.5 10.10.0.6 10.10.0.7 10.10.0.8
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 16 path(s):
  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 10.10.0.2 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 10.10.0.3 to Remote 10.10.0.5
  Local 10.10.0.3 to Remote 10.10.0.6
  Local 10.10.0.3 to Remote 10.10.0.7
  Local 10.10.0.3 to Remote 10.10.0.8
  Local 10.10.0.4 to Remote 10.10.0.5
  Local 10.10.0.4 to Remote 10.10.0.6
  Local 10.10.0.4 to Remote 10.10.0.7
  Local 10.10.0.4 to Remote 10.10.0.8
Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
4 paths up, 0 paths down (udp check)
```

35. Expand the cluster by adding nodes to the Nexus 3232C cluster switches.

The following examples show nodes n3 and n4 have 40 GbE cluster ports connected to ports e1/7 and e1/8 respectively on both the Nexus 3232C cluster switches, and both nodes have joined the cluster. The 40 GbE cluster interconnect ports used are e4a and e4e.

36. Display the 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

```
cluster::> network device-discovery show
Node      Local Port   Discovered Device      Interface      Platform
-----
n1        /cdp
e0a      C1        Ethernet1/1/1  N3K-C3232C
e0b      C2        Ethernet1/1/1  N3K-C3232C
e0c      C2        Ethernet1/1/2  N3K-C3232C
e0d      C1        Ethernet1/1/2  N3K-C3232C
n2        /cdp
e0a      C1        Ethernet1/1/3  N3K-C3232C
e0b      C2        Ethernet1/1/3  N3K-C3232C
e0c      C2        Ethernet1/1/4  N3K-C3232C
e0d      C1        Ethernet1/1/4  N3K-C3232C
n3        /cdp
e4a      C1        Ethernet1/7    N3K-C3232C
e4e      C2        Ethernet1/7    N3K-C3232C
n4        /cdp
e4a      C1        Ethernet1/8    N3K-C3232C
e4e      C2        Ethernet1/8    N3K-C3232C
12 entries were displayed.
```

```
cluster::*> network port show -role cluster
(network port show)
Node: n1
Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health Ignore
Admin/Oper Status Health
-----
e0a      Cluster      Cluster      up  9000    auto/10000 - -
e0b      Cluster      Cluster      up  9000    auto/10000 - -
e0c      Cluster      Cluster      up  9000    auto/10000 - -
e0d      Cluster      Cluster      up  9000    auto/10000 - -
Node: n2
Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health Ignore
Admin/Oper Status Health
-----
e0a      Cluster      Cluster      up  9000    auto/10000 - -
e0b      Cluster      Cluster      up  9000    auto/10000 - -
e0c      Cluster      Cluster      up  9000    auto/10000 - -
e0d      Cluster      Cluster      up  9000    auto/10000 - -
Node: n3
Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health Ignore
Admin/Oper Status Health
-----
e4a      Cluster      Cluster      up  9000    auto/40000 - -
e4e      Cluster      Cluster      up  9000    auto/40000 - -
Node: n4
Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health Ignore
Admin/Oper Status Health
-----
e4a      Cluster      Cluster      up  9000    auto/40000 - -
e4e      Cluster      Cluster      up  9000    auto/40000 - -
12 entries were displayed.
```

```
cluster::*> network interface show -role cluster
(network interface show)
-----
Vserver      Logical      Status      Network      Current      Current      Is
Interface    Admin/Oper  Address/Mask Node          Port         Home
-----
Cluster
n1_clus1     up/up       10.10.0.1/24 n1           e0a          true
n1_clus2     up/up       10.10.0.2/24 n1           e0b          true
n1_clus3     up/up       10.10.0.3/24 n1           e0c          true
n1_clus4     up/up       10.10.0.4/24 n1           e0d          true
n2_clus1     up/up       10.10.0.5/24 n2           e0a          true
n2_clus2     up/up       10.10.0.6/24 n2           e0b          true
n2_clus3     up/up       10.10.0.7/24 n2           e0c          true
n2_clus4     up/up       10.10.0.8/24 n2           e0d          true
n3_clus1     up/up       10.10.0.9/24 n3           e4a          true
n3_clus2     up/up       10.10.0.10/24 n3          e4e          true
n4_clus1     up/up       10.10.0.11/24 n4          e4a          true
n4_clus2     up/up       10.10.0.12/24 n4          e4e          true
12 entries were displayed.
```

```
cluster::*> system cluster-switch show
-----
Switch      Type          Address      Model
-----
C1          cluster-network 10.10.1.103  NX3232C
  Serial Number: FOX000001
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software, Version
                    7.0(3)I4(1)
  Version Source: CDP
C2          cluster-network 10.10.1.104  NX3232C
  Serial Number: FOX000002
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software, Version
                    7.0(3)I4(1)
  Version Source: CDP
CL1        cluster-network 10.10.1.101  NX5596
  Serial Number: 01234567
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software, Version
                    7.1(1)N1(1)
  Version Source: CDP
CL2        cluster-network 10.10.1.102  NX5596
  Serial Number: 01234568
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software, Version
                    7.1(1)N1(1)
  Version Source: CDP
4 entries were displayed.
```

- 37. Remove the replaced Nexus 5596 by using the system cluster-switch delete command, if it is not automatically removed:

```
system cluster-switch delete -device switch-name
```

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

- 38. Verify that the proper cluster switches are monitored:

```
system cluster-switch show
```

```
cluster::> system cluster-switch show
-----
Switch      Type          Address      Model
-----
C1          cluster-network 10.10.1.103  NX3232C
```

```
Serial Number: FOX000001
Is Monitored: true
Reason:
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version
7.0(3)I4(1)
Version Source: CDP

C2                cluster-network      10.10.1.104      NX3232C
Serial Number: FOX000002
Is Monitored: true
Reason:
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version
7.0(3)I4(1)
Version Source: CDP

2 entries were displayed.
```

39. Enable the cluster switch health monitor log collection feature for collecting switch-related log files:

```
system cluster-switch log setup-password
```

```
system cluster-switch log enable-collection
```

```
cluster::*> system cluster-switch log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
C1
C2

cluster::*> system cluster-switch log setup-password

Enter the switch name: C1
RSA key fingerprint is e5:8b:c6:dc:e2:18:18:09:36:63:d9:63:dd:03:d9:cc
Do you want to continue? {y|n}::[n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system cluster-switch log setup-password

Enter the switch name: C2
RSA key fingerprint is 57:49:86:a1:b9:80:6a:61:9a:86:8e:3c:e3:b7:1f:b1
Do you want to continue? {y|n}:: [n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster::*> system cluster-switch log enable-collection

Do you want to enable cluster log collection for all nodes in the cluster?
{y|n}: [n] y

Enabling cluster switch log collection.

cluster::*>
```

Note: If any of these commands return an error, contact NetApp support.

40. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

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

Related information

[Cisco Ethernet Switch description page](#)

[Hardware Universe](#)

Copyright and trademark

Copyright

Copyright © 2020 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

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>

Index

Numerics

3232C cluster switches
replacing 5596 cluster switches with [3](#)

C

Cisco Nexus cluster switches
replacing 5596 switches with 3232C switches [3](#), [4](#)
cluster switches
replacing 5596 cluster switches with 3232C switches [4](#)