



Migrate switches

Cluster and storage switches

NetApp
April 05, 2024

Table of Contents

- Migrate switches. 1
 - Migration requirements for Cisco Nexus 3232C cluster switches 1
 - Migrate a CN1610 cluster switch to a Cisco Nexus 3232C cluster switch 3
 - Migrate from a Cisco Nexus 5596 cluster switch to a Cisco Nexus 3232C cluster switch. 37
 - Migrate from a two-node switchless cluster to a cluster with Cisco Nexus 3232C cluster switches 72

Migrate switches

Migration requirements for Cisco Nexus 3232C cluster switches

Before you migrate to Cisco Nexus 3232C cluster switches, review the configuration information, port connections, and cabling requirements.

CN1610 migrate requirements

The cluster switches support the following node connections:

- NetApp CN1610: ports 0/1 through 0/12 (10 GbE)
- Cisco Nexus 3232C: ports e1/1-30 (40 or 100 or 4x10GbE)

The cluster switches use the following inter-switch link (ISL) ports.

- NetApp CN1610: ports 0/13 through 0/16 (10 GbE)
- Cisco Nexus 3232C: ports 1/31-32 (100GbE)



You must use 4x10G breakout cables on the Cisco Nexus 3232C cluster switch.

The following table shows the cabling connections that are required at each stage as you make the transition from NetApp CN1610 switches to Cisco Nexus 3232C cluster switches:

Stage	Description	Required cables
Initial	CN1610 to CN1610 (SFP+ to SFP+)	4 SFP+ optical fiber or copper direct-attach cables
Transition	CN1610 to 3232C (QSFP to SFP+)	1 QSFP and 4 SFP+ optical fiber or copper breakout cables
Final	3232C to 3232C (QSFP to QSFP)	2 QSFP optical fiber or copper direct-attach cables

You must have downloaded the applicable reference configuration files (RCFs). The number of 10 GbE and 40/100 GbE ports are defined in the RCFs available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.

The ONTAP and NX-OS versions that are supported in this procedure are listed on the [Cisco Ethernet Switches page](#).

The ONTAP and FASTPATH versions that are supported in this procedure are listed on the [NetApp CN1601 and CN1610 Switches page](#).

CN5596 requirements

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.



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](#).

Migrate a CN1610 cluster switch to a Cisco Nexus 3232C cluster switch

To replace the existing CN1610 cluster switches in a cluster with Cisco Nexus 3232C cluster switches, you must perform a specific sequence of tasks.

Review requirements

Before migration, be sure to review [Migration requirements](#).



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

If necessary, refer to the following for more information:

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

Migrate the switches

About the examples

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 fiber cables: e4a and e4e. The [Hardware Universe](#) has information about the cluster fiber cables on your platforms.

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

- The nodes are n1, n2, n3, and n4.
- The command outputs might vary depending on different releases of ONTAP software.
- The CN1610 switches to be replaced are CL1 and CL2.
- The Nexus 3232C switches to replace the CN1610 switches are C1 and C2.
- n1_clus1 is the first cluster logical interface (LIF) that is connected to cluster switch 1 (CL1 or C1) for node n1.
- n1_clus2 is the first cluster LIF that is connected to cluster switch 2 (CL2 or C2) for node n1.
- n1_clus3 is the second LIF that is connected to cluster switch 2 (CL2 or C2) for node n1.
- n1_clus4 is the second LIF that is 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.

Step 1: Prepare for migration

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.



The 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
```

Show example

The following example displays 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	0/1	CN1610
	e0b	CL2	0/1	CN1610
	e0c	CL2	0/2	CN1610
	e0d	CL1	0/2	CN1610
n2	/cdp			
	e0a	CL1	0/3	CN1610
	e0b	CL2	0/3	CN1610
	e0c	CL2	0/4	CN1610
	e0d	CL1	0/4	CN1610

8 entries were displayed.

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

a. Display the cluster network port attributes:

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

Show example

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

Node: n1

Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Open	Health Status	Ignore 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) Admin/Open	Health Status	Ignore 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	-	

8 entries were displayed.

b. Display information about the logical interfaces:

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

Show example

```
cluster::*> network interface show -role cluster
(network interface show)
      Logical      Status      Network      Current      Current
Is
Vserver  Interface  Admin/Oper  Address/Mask  Node      Port
Home
-----
-----
Cluster
true      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

8 entries were displayed.
```

c. Display information about the discovered cluster switches:

```
system cluster-switch show
```

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	cluster-network	10.10.1.101
CN1610		
Serial Number: 01234567		
Is Monitored: true		
Reason:		
Software Version: 1.2.0.7		
Version Source: ISDP		
CL2	cluster-network	10.10.1.102
CN1610		
Serial Number: 01234568		
Is Monitored: true		
Reason:		
Software Version: 1.2.0.7		
Version Source: ISDP		

2 entries displayed.

4. Verify that the appropriate RCF and image are installed on the new 3232C 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, you must complete the following procedure:

- a. See the [Cisco Ethernet Switch](#) page on the NetApp Support Site.
 - b. Note your switch and the required software versions in the table on that page.
 - c. 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 at [Cisco® Cluster and Management Network Switch Reference Configuration File Download](#).
5. Migrate the LIFs associated with the second CN1610 switch that you plan to replace:

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

Show example

You must migrate each LIF individually as shown in the following example:

```
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
```

Show example

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

8 entries were displayed.
```

Step 2: Replace cluster switch CL2 with C2

1. 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
```

Show example

The following example shows the four cluster interconnect ports being shut down for node n1 and node n2:

```
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
```

2. Ping the remote cluster interfaces, and then perform a remote procedure call server check:

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

Show example

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 9000 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)
```

3. Shut down the ISL ports 13 through 16 on the active CN1610 switch CL1 using the appropriate command.

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

Show example

The following example shows ISL ports 13 through 16 being shut down 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)#
```

4. Build a temporary ISL between CL1 and C2:

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

Show example

The following example shows a temporary ISL being built between CL1 (ports 13-16) and C2 (ports e1/24/1-4) using the Cisco switchport mode trunk command:

```
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
```

5. Remove the cables that are attached to the CN1610 switch CL2 on all the nodes.

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

6. Remove four ISL cables from ports 13 to 16 on the CN1610 switch CL1.

You must attach the appropriate Cisco QSFP28 to SFP+ breakout cables connecting port 1/24 on the new Cisco 3232C switch C2 to ports 13 to 16 on the existing CN1610 switch CL1.



When reconnecting any cables to the new Cisco 3232C switch, the cables used must be either optical fiber or Cisco twinax cables.

7. Make the ISL dynamic by configuring the ISL interface 3/1 on the active CN1610 switch to disable the static mode.

This configuration matches with the ISL configuration on the 3232C switch C2 when the ISLs are brought up on both switches.

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

Show example

The following example shows the ISL interface 3/1 being configured 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) #
```

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

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

Show example

The following example shows ISL ports 13 through 16 being brought up on the port-channel interface 3/1:

```
(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) #
```

9. Verify that the ISLs are up on the CN1610 switch CL1.

The "Link State" should be Up, "Type" should be Dynamic, and the "Port Active" column should be True for ports 0/13 to 0/16.

Show example

The following example shows the ISLs being verified as up on the CN1610 switch CL1:

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

10. Verify that the ISLs are up on the 3232C switch C2:

```
show port-channel summary
```

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

Ports Eth1/24/1 through Eth1/24/4 should indicate (P), meaning that all four ISL ports are up in the port channel. Eth1/31 and Eth1/32 should indicate (D) as they are not connected.

Show example

The following example shows the ISLs being verified as up on the 3232C 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      Po1 (SU)       Eth      LACP      Eth1/31 (D)  Eth1/32 (D)
2      Po2 (SU)       Eth      LACP      Eth1/24/1 (P) Eth1/24/2 (P)
Eth1/24/3 (P)
                                   Eth1/24/4 (P)
```

11. Bring up all of the cluster interconnect ports that are connected to the 3232C switch C2 on all of the nodes:

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

Show example

The following example shows how to bring up the cluster interconnect ports connected to the 3232C switch C2:

```
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
```

12. Revert all of the migrated cluster interconnect LIFs that are connected to C2 on all of the nodes:

```
network interface revert -vserver cluster -lif lif-name
```

Show example

```
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
```

13. Verify that all of the cluster interconnect ports are reverted to their home ports:

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

Show example

The following example shows that the LIFs on clus2 are reverted to their home ports; 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
(network interface show)
      Logical      Status      Network      Current  Current  Is
Vserver Interface  Admin/Oper  Address/Mask  Node     Port     Home
-----
Cluster
true    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

8 entries were displayed.
```

14. Verify that all of the cluster ports are connected:

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

Show example

The following example shows the output verifying all of the cluster interconnects are up:

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

Node: n1

Port  IPspace  Broadcast  Link  MTU  Speed (Mbps)  Health  Ignore
Status  Domain      Health      Admin/Open  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  Link  MTU  Speed (Mbps)  Health  Ignore
Status  Domain      Health      Admin/Open  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    -

8 entries were displayed.
```

15. Ping the remote cluster interfaces and then perform a remote procedure call server check:

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

Show example

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)
```

16. Migrate the LIFs that are associated with the first CN1610 switch CL1:

```
network interface migrate -vserver cluster -lif lif-name -source-node node-name
```

Show example

You must migrate each cluster LIF individually to the appropriate cluster ports hosted on cluster switch C2 as shown in the following example:

```
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
```

Step 3: Replace cluster switch CL1 with C1

1. Verify the cluster's status:

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

Show example

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

```
cluster::*> network interface show -role cluster
(network interface show)
      Logical      Status      Network      Current  Current  Is
Vserver Interface  Admin/Oper  Address/Mask  Node      Port
Home
-----
-----
Cluster
false      n1_clus1      up/up      10.10.0.1/24  n1        e0b
true       n1_clus2      up/up      10.10.0.2/24  n1        e0b
true       n1_clus3      up/up      10.10.0.3/24  n1        e0c
false      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.
```

2. Shut down the node ports that are connected to CL1 on all of the nodes:

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

Show example

The following example shows specific 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
```

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

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

Show example

The following example shows ISLs 24, 31, and 32 being shut down on the active 3232C switch C2:

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

4. Remove the cables that are attached to the CN1610 switch CL1 on all of the nodes.

Using the appropriate cabling, you must reconnect the disconnected ports on all the nodes to the Nexus 3232C switch C1.

5. Remove the QSFP28 cables from Nexus 3232C C2 port e1/24.

You must connect ports e1/31 and e1/32 on C1 to ports e1/31 and e1/32 on C2 using supported Cisco QSFP28 optical fiber or direct-attach cables.

6. Restore the configuration on port 24 and remove the temporary port-channel 2 on C2:

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

Show example

The following example shows the running-configuration file being copied to the startup-configuration file:

```
C2# configure
C2(config)# no interface breakout module 1 port 24 map 10g-4x
C2(config)# no interface port-channel 2
C2(config-if)# interface e1/24
C2(config-if)# description 100GbE/40GbE Node Port
C2(config-if)# spanning-tree port type edge
Edge port type (portfast) should only be enabled on ports connected
to a single
host. Connecting hubs, concentrators, switches, bridges, etc... to
this
interface when edge port type (portfast) is enabled, can cause
temporary bridging loops.
Use with CAUTION

Edge Port Type (Portfast) has been configured on Ethernet 1/24 but
will only
have effect when the interface is in a non-trunking mode.

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.
```

7. Bring up ISL ports 31 and 32 on C2, the active 3232C switch.

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

Show example

The following example shows ISLs 31 and 32 being brought upon the 3232C switch C2:

```
C2# configure
C2(config)# interface ethernet 1/31-32
C2(config-if-range)# no shutdown
C2(config-if-range)# exit
C2(config)# exit
C2# copy running-config startup-config
[] 100%
Copy Complete.
```

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

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

Show example

The following example shows the ISL connections being verified. Ports Eth1/31 and Eth1/32 indicate (P), meaning that both the ISL ports are 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      Po1(SU)        Eth       LACP      Eth1/31(P)  Eth1/32(P)
```

```
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      Po1(SU)        Eth       LACP      Eth1/31(P)  Eth1/32(P)
```

9. Bring up all of the cluster interconnect ports connected to the new 3232C switch C1 on all of the nodes:

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

Show example

The following example shows all of the cluster interconnect ports connected to the new 3232C switch C1 being brought up:

```
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
```

10. Verify the status of the cluster node port:

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

Show example

The following example shows output that verifies that the cluster interconnect ports on nodes n1 and n2 on the new 3232C switch C1 are up:

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

Node: n1

Port  IPspace  Broadcast  Link  MTU  Speed (Mbps)  Health  Ignore
Status  Domain      Status      Admin/Open  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  Link  MTU  Speed (Mbps)  Health  Ignore
Status  Domain      Status      Admin/Open  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    -

8 entries were displayed.
```

Step 4: Complete the procedure

1. Revert all of the migrated cluster interconnect LIFs that were originally connected to C1 on all of the nodes:

```
network interface revert -server cluster -lif lif-name
```

Show example

You must migrate each LIF individually as shown in the following example:

```
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
```

2. Verify that the interface is now home:

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

Show example

The following example shows the status of cluster interconnect interfaces is up and "Is Home" for nodes n1 and n2:

```
cluster::*> network interface show -role cluster
(network interface show)
      Logical      Status      Network      Current  Current  Is
Vserver Interface  Admin/Oper  Address/Mask  Node     Port    Home
-----
Cluster
true      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

8 entries were displayed.
```

3. Ping the remote cluster interfaces and then perform a remote procedure call server check:

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

Show example

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 9000 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)
3  paths up, 0 paths down (udp check)
```

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

5. 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

Show example

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

```
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
n2	/cdp			
	e0d	C1	Ethernet1/1/2	N3K-C3232C
	e0a	C1	Ethernet1/1/3	N3K-C3232C
	e0b	C2	Ethernet1/1/3	N3K-C3232C
n3	/cdp			
	e0c	C2	Ethernet1/1/4	N3K-C3232C
	e0d	C1	Ethernet1/1/4	N3K-C3232C
	e4a	C1	Ethernet1/7	N3K-C3232C
n4	/cdp			
	e4e	C2	Ethernet1/7	N3K-C3232C
	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

		Broadcast		Speed (Mbps)		Health	
Ignore							
Port	IPspace	Domain	Link	MTU	Admin/Open	Status	
Health	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

		Broadcast			Speed (Mbps)	Health
Ignore						
Port	IPspace	Domain	Link	MTU	Admin/Open	Status
Health	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: n3

		Broadcast			Speed (Mbps)	Health
Ignore						
Port	IPspace	Domain	Link	MTU	Admin/Open	Status
Health	Status					
-----	-----	-----	-----	-----	-----	-----
e4a	cluster	cluster	up	9000	auto/40000	-
e4e	cluster	cluster	up	9000	auto/40000	-

Node: n4

		Broadcast			Speed (Mbps)	Health
Ignore						
Port	IPspace	Domain	Link	MTU	Admin/Open	Status
Health	Status					
-----	-----	-----	-----	-----	-----	-----
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)

	Logical	Status	Network	Current	Current
Is					
Vserver	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)I6(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)I6(1)
    Version Source: CDP
CL1                cluster-network  10.10.1.101    CN1610

    Serial Number: 01234567
    Is Monitored: true
    Reason:
    Software Version: 1.2.0.7
    Version Source: ISDP
CL2                cluster-network  10.10.1.102
CN1610

    Serial Number: 01234568
    Is Monitored: true
    Reason:
    Software Version: 1.2.0.7
    Version Source: ISDP 4 entries were displayed.

```

6. Remove the replaced CN1610 switches if they are not automatically removed:

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

Show example

You must delete both devices individually as shown in the following example:

```

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

```

7. Verify that the proper cluster switches are monitored:

```
system cluster-switch show
```

Show example

The following example shows cluster switches C1 and C2 are being monitored:

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

Switch Model	Type	Address

C1 NX3232C	cluster-network	10.10.1.103
Serial Number: FOX000001		
Is Monitored: true		
Reason:		
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version		
7.0(3)I6(1)		
Version Source: CDP		
C2 NX3232C	cluster-network	10.10.1.104
Serial Number: FOX000002		
Is Monitored: true		
Reason:		
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version		
7.0(3)I6(1)		
Version Source: CDP		

2 entries were displayed.

8. 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
```

Show example

```
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>

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



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

9. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

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

Migrate from a Cisco Nexus 5596 cluster switch to a Cisco Nexus 3232C cluster switch

Follow this procedure to migrate an existing Cisco Nexus 5596 cluster switches in a cluster with Nexus 3232C cluster switches.

Review requirements

Before migration, be sure to review [Migration requirements](#).



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

For more information, see:

- [Cisco Ethernet Switch description page](#)
- [Hardware Universe](#)

Migrate the switch

About the examples

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.

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.

Scenarios

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 CL2 to be replaced by C2.
 - 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.

Step 1: Prepare for migration

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.



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
```

Show example

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
```

Show example

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

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health      Health      Speed (Mbps)
Port        IPspace      Broadcast Domain Link MTU  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

Ignore

Health      Health      Speed (Mbps)
Port        IPspace      Broadcast Domain Link MTU  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.
```

b. Display information about the logical interfaces:

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

Show example

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

c. Display information about the discovered cluster switches:

```
system cluster-switch show
```

Show example

The following example shows the active cluster switches:

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

Switch Model	Type	Address
CL1 NX5596	cluster-network	10.10.1.101
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 NX5596	cluster-network	10.10.1.102
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.



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 vservice-name -lif lif-name -source-node
source-node-name - destination-node node-name -destination-port destination-
port-name
```

Show example

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
```

Show example

The following example shows the current status of each cluster:

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

Step 2: Configure ports

1. 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
```

Show example

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
```

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

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

Show example

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)
```

3. 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](#).

Show example

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) #
```

4. 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](#).

Show example

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
```

5. 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.

6. 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.

7. 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](#).

Show example

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) #
```

8. 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](#).

Show example

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)
```

9. Verify that interfaces eth1/45-48 already have `channel-group 1 mode active` in their running configuration.
10. 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
```

Show example

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
```

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

```
network interface revert -vserver Cluster -lif lif-name
```

Show example

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
```

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

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

Show example

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

13. Verify that the clustered ports are connected:

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

Show example

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
```

```
Ignore
```

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	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
```

```
Ignore
```

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	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.
```

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

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

Show example

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)
```

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

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

Show example

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
```

16. Verify the cluster's status:

```
network interface show
```

Show example

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
```

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

17. 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
```

Show example

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
```

18. 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](#).

Show example

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#
```

19. 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.

20. 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.

21. 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](#).

Show example

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.
```

22. 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](#).

Show example

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
```

23. 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

Show example

```
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      Po1 (SU)      Eth     LACP      Eth1/31 (P)  Eth1/32 (P)
```

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

```
network port modify
```

Show example

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
```

25. Verify the status of the cluster node port:

```
network port show
```

Show example

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

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  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

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  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.
```

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

```
network interface revert -server Cluster -lif lif-name
```

Show example

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
```

27. Verify that the interface is home:

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

Show example

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

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

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

Show example

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)
```

29. 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.

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`

Show example

```
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
```

```
Ignore
```

Health	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper	Speed(Mbps)	Health 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

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	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: n3

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----		----	----	-----	
-----	-----						
e4a	Cluster	Cluster		up	9000	auto/40000	-
-							
e4e	Cluster	Cluster		up	9000	auto/40000	-
-							

Node: n4

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	----	-----	-----	
-----	-----						
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)
```

	Logical	Status	Network	Current
Current Is				
Vserver	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 Model	Type	Address

C1 NX3232C	cluster-network	10.10.1.103
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 NX3232C	cluster-network	10.10.1.104
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 NX5596	cluster-network	10.10.1.101
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 NX5596	cluster-network	10.10.1.102
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.
```

30. 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
```

Show example

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

Step 3: Complete the procedure

1. Verify that the proper cluster switches are monitored:

```
system cluster-switch show
```

Show example

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

Switch Model	Type	Address

C1 NX3232C	cluster-network	10.10.1.103
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 NX3232C	cluster-network	10.10.1.104
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.

2. 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
```

Show example

```
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::*>
```



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

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

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

Migrate from a two-node switchless cluster to a cluster with Cisco Nexus 3232C cluster switches

If you have a two-node *switchless* cluster, you can migrate to a two-node *switched* cluster that includes Cisco Nexus 3232C cluster network switches. This is a nondisruptive procedure.

Review requirements

Migration requirements

Before migration, be sure to review [Migration requirements](#).

What you'll need

Ensure that:

- Ports are available for node connections. The cluster switches use the Inter-Switch Link (ISL) ports e1/31-32.
- You have appropriate cables for cluster connections:
 - The nodes with 10 GbE cluster connections require QSFP optical modules with breakout fiber cables or QSFP to SFP+ copper breakout cables.
 - The 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 require the appropriate ISL cabling: 2x QSFP28 fiber or copper direct-attach cables.
- The configurations are properly set up and functioning.

The two nodes must be connected and functioning in a two-node switchless cluster setting.

- All cluster ports are in the **up** state.
- The Cisco Nexus 3232C cluster switch are supported.
- The existing cluster network configuration has the following:
 - A redundant and fully functional Nexus 3232C cluster infrastructure on both switches
 - The latest RCF and NX-OS versions on your switches
 - Management connectivity on both switches
 - Console access to both switches
 - All cluster logical interfaces (LIFs) in the **up** state without having been migrated
 - Initial customization of the switch
 - All ISL ports enabled and cabled

Migrate the switches

About the examples

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

- Nexus 3232C cluster switches, C1 and C2.

- The nodes are n1 and n2.

The examples in this procedure use two nodes, each utilizing two 40 GbE cluster interconnect ports e4a and e4e. The [Hardware Universe](#) has details about the cluster ports on your platforms.

- n1_clus1 is the first cluster logical interface (LIF) to be connected to cluster switch C1 for node n1.
- n1_clus2 is the first cluster LIF to be connected to cluster switch C2 for node n1.
- n2_clus1 is the first cluster LIF to be connected to cluster switch C1 for node n2.
- n2_clus2 is the second cluster LIF to be connected to cluster switch C2 for node n2.
- 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 procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

Step 1: Display and migrate physical and logical ports

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.



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

2. Determine the administrative or operational status for each cluster interface:
 - a. Display the network port attributes:

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

Show example

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health      Health      Speed (Mbps)
Port        IPspace      Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
-----
e4a         Cluster      Cluster      up    9000 auto/40000 -
e4e         Cluster      Cluster      up    9000 auto/40000 -
-
Node: n2

Ignore

Health      Health      Speed (Mbps)
Port        IPspace      Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
-----
e4a         Cluster      Cluster      up    9000 auto/40000 -
e4e         Cluster      Cluster      up    9000 auto/40000 -
4 entries were displayed.
```

- b. Display information about the logical interfaces and their designated home nodes:

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

Show example

```
cluster::*> network interface show -role cluster
(network interface show)
      Logical      Status      Network      Current
Current Is
Vserver      Interface  Admin/Oper  Address/Mask  Node
Port      Home
-----
Cluster
      n1_clus1      up/up      10.10.0.1/24      n1
e4a      true
      n1_clus2      up/up      10.10.0.2/24      n1
e4e      true
      n2_clus1      up/up      10.10.0.3/24      n2
e4a      true
      n2_clus2      up/up      10.10.0.4/24      n2
e4e      true

4 entries were displayed.
```

- c. Verify that switchless cluster detection is enabled using the advanced privilege command:

```
network options detect-switchless-cluster show`
```

Show example

The output in the following example shows that switchless cluster detection is enabled:

```
cluster::*> network options detect-switchless-cluster show
Enable Switchless Cluster Detection: true
```

3. Verify that the appropriate RCFs and image are installed on the new 3232C switches and make any necessary site customizations such as adding users, passwords, and network addresses.

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

[Cisco Cluster and Management Network Switch Reference Configuration File download page](#)

4. Click **CONTINUE** on the **Description** page, accept the license agreement, and then follow the instructions on the **Download** page to download the RCF.
5. On Nexus 3232C switches C1 and C2, disable all node-facing ports C1 and C2, but do not disable the ISL ports e1/31-32.

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

Show example

The following example shows ports 1 through 30 being disabled on Nexus 3232C cluster switches C1 and C2 using a configuration supported in RCF NX3232_RCF_v1.0_24p10g_24p100g.txt:

```
C1# copy running-config startup-config
[] 100% Copy complete.
C1# configure
C1(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C1(config-if-range)# shutdown
C1(config-if-range)# exit
C1(config)# exit
C2# copy running-config startup-config
[] 100% Copy complete.
C2# configure
C2(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C2(config-if-range)# shutdown
C2(config-if-range)# exit
C2(config)# exit
```

6. Connect ports 1/31 and 1/32 on C1 to the same ports on C2 using supported cabling.
7. Verify that the ISL ports are operational on C1 and C2:

```
show port-channel summary
```

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

Show example

The following example shows the Cisco `show port-channel summary` command being used to verify the ISL ports are operational on C1 and C2:

```
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

-----
-----
      Port-
Group Channel          Type   Protocol  Member Ports
-----
-----
1      Po1(SU)         Eth    LACP      Eth1/31(P)  Eth1/32(P)

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      Po1(SU)         Eth    LACP      Eth1/31(P)  Eth1/32(P)
```

8. Display the list of neighboring devices on the switch.

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

Show example

The following example shows the Cisco command `show cdp neighbors` being used to display the neighboring devices on the switch:

```
C1# show cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-
Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute
Device-ID          Local Intrfce  Hldtme Capability  Platform
Port ID
C2                  Eth1/31       174    R S I s          N3K-C3232C
Eth1/31
C2                  Eth1/32       174    R S I s          N3K-C3232C
Eth1/32
Total entries displayed: 2
C2# show cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-
Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute
Device-ID          Local Intrfce  Hldtme Capability  Platform
Port ID
C1                  Eth1/31       178    R S I s          N3K-C3232C
Eth1/31
C1                  Eth1/32       178    R S I s          N3K-C3232C
Eth1/32
Total entries displayed: 2
```

9. Display the cluster port connectivity on each node:

```
network device-discovery show
```

Show example

The following example shows the cluster port connectivity displayed for a two-node switchless cluster configuration:

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

	Local	Discovered		
Node	Port	Device	Interface	Platform
n1	/cdp			
	e4a	n2	e4a	FAS9000
	e4e	n2	e4e	FAS9000
n2	/cdp			
	e4a	n1	e4a	FAS9000
	e4e	n1	e4e	FAS9000

10. Migrate the n1_clus1 and n2_clus1 LIFs to the physical ports of their destination nodes:

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

Show example

You must execute the command for each local node as shown in the following example:

```
cluster::*> network interface migrate -vserver cluster -lif n1_clus1  
-source-node n1  
-destination-node n1 -destination-port e4e  
cluster::*> network interface migrate -vserver cluster -lif n2_clus1  
-source-node n2  
-destination-node n2 -destination-port e4e
```

Step 2: Shut down the reassigned LIFs and disconnect the cables

1. Verify the cluster interfaces have successfully migrated:

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

Show example

The following example shows the "Is Home" status for the n1_clus1 and n2_clus1 LIFs has become "false" after the migration is completed:

```
cluster::*> network interface show -role cluster
(network interface show)
Current Is Logical Status Network Current
Vserver Interface Admin/Oper Address/Mask Node
Port Home
-----
Cluster
e4e n1_clus1 up/up 10.10.0.1/24 n1
false
e4e n1_clus2 up/up 10.10.0.2/24 n1
true
e4e n2_clus1 up/up 10.10.0.3/24 n2
false
e4e n2_clus2 up/up 10.10.0.4/24 n2
true
4 entries were displayed.
```

2. Shut down cluster ports for the n1_clus1 and n2_clus1 LIFs, which were migrated in step 9:

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

Show example

You must execute the command for each port as shown in the following example:

```
cluster::*> network port modify -node n1 -port e4a -up-admin false
cluster::*> network port modify -node n2 -port e4a -up-admin false
```

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

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

Show example

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          e4a    10.10.0.1
Cluster n1_clus2 n1          e4e    10.10.0.2
Cluster n2_clus1 n2          e4a    10.10.0.3
Cluster n2_clus2 n2          e4e    10.10.0.4
Local = 10.10.0.1 10.10.0.2
Remote = 10.10.0.3 10.10.0.4
Cluster Vserver Id = 4294967293 Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s) .....
Detected 9000 byte MTU on 32 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 10.10.0.2 to Remote 10.10.0.3
    Local 10.10.0.2 to Remote 10.10.0.4
Larger than PMTU communication succeeds on 4 path(s) RPC status:
1 paths up, 0 paths down (tcp check)
1 paths up, 0 paths down (ucp check)
```

4. Disconnect the cable from e4a on node n1.

You can refer to the running configuration and connect the first 40 GbE port on the switch C1 (port 1/7 in this example) to e4a on n1 using cabling supported for Nexus 3232C switches.

Step 3: Enable the cluster ports

1. Disconnect the cable from e4a on node n2.

You can refer to the running configuration and connect e4a to the next available 40 GbE port on C1, port 1/8, using supported cabling.

2. Enable all node-facing ports on C1.

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

Show example

The following example shows ports 1 through 30 being enabled on Nexus 3232C cluster switches C1 and C2 using the configuration supported in RCF NX3232_RCF_v1.0_24p10g_26p100g.txt:

```
C1# configure
C1(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C1(config-if-range)# no shutdown
C1(config-if-range)# exit
C1(config)# exit
```

3. Enable the first cluster port, e4a, on each node:

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

Show example

```
cluster::*> network port modify -node n1 -port e4a -up-admin true
cluster::*> network port modify -node n2 -port e4a -up-admin true
```

4. Verify that the clusters are up on both nodes:

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

Show example

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Speed(Mbps) Health
Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e4a       Cluster      Cluster      up    9000 auto/40000 -
e4e       Cluster      Cluster      up    9000 auto/40000 -
-

Node: n2

Ignore

Speed(Mbps) Health
Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e4a       Cluster      Cluster      up    9000 auto/40000 -
e4e       Cluster      Cluster      up    9000 auto/40000 -

4 entries were displayed.
```

5. For each node, revert all of the migrated cluster interconnect LIFs:

```
network interface revert -vserver cluster -lif lif-name
```

Show example

You must revert each LIF to its home port individually as shown in the following example:

```
cluster::*> network interface revert -vserver cluster -lif n1_clus1
cluster::*> network interface revert -vserver cluster -lif n2_clus1
```

6. Verify that all the LIFs are now reverted to their home ports:

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

The `Is Home` column should display a value of `true` for all of the ports listed in the `Current Port` column. If the displayed value is `false`, the port has not been reverted.

Show example

```
cluster::*> network interface show -role cluster
(network interface show)
      Logical      Status      Network      Current
Current Is
Vserver      Interface  Admin/Oper Address/Mask      Node
Port      Home
-----
Cluster
      n1_clus1    up/up      10.10.0.1/24      n1
e4a      true
      n1_clus2    up/up      10.10.0.2/24      n1
e4e      true
      n2_clus1    up/up      10.10.0.3/24      n2
e4a      true
      n2_clus2    up/up      10.10.0.4/24      n2
e4e      true
4 entries were displayed.
```

Step 4: Enable the reassigned LIFs

1. Display the cluster port connectivity on each node:

```
network device-discovery show
```

Show example

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

	Local	Discovered		
Node	Port	Device	Interface	Platform

n1	/cdp			
	e4a	C1	Ethernet1/7	N3K-C3232C
	e4e	n2	e4e	FAS9000
n2	/cdp			
	e4a	C1	Ethernet1/8	N3K-C3232C
	e4e	n1	e4e	FAS9000

2. Migrate clus2 to port e4a on the console of each node:

```
network interface migrate cluster -lif lif-name -source-node source-node-name
-destination-node destination-node-name -destination-port destination-port-
name
```

Show example

You must migrate each LIF to its home port individually as shown in the following example:

```
cluster::*> network interface migrate -vserver cluster -lif n1_clus2
-source-node n1
-destination-node n1 -destination-port e4a
cluster::*> network interface migrate -vserver cluster -lif n2_clus2
-source-node n2
-destination-node n2 -destination-port e4a
```

3. Shut down cluster ports clus2 LIF on both nodes:

```
network port modify
```

Show example

The following example shows the specified ports being set to `false`, shutting the ports down on both nodes:

```
cluster::*> network port modify -node n1 -port e4e -up-admin false
cluster::*> network port modify -node n2 -port e4e -up-admin false
```

4. Verify the cluster LIF status:

```
network interface show
```

Show example

```
cluster::*> network interface show -role cluster
(network interface show)
      Logical      Status      Network      Current
Current Is
Vserver      Interface  Admin/Oper  Address/Mask      Node
Port      Home
-----
Cluster
      n1_clus1      up/up      10.10.0.1/24      n1
e4a      true
      n1_clus2      up/up      10.10.0.2/24      n1
e4a      false
      n2_clus1      up/up      10.10.0.3/24      n2
e4a      true
      n2_clus2      up/up      10.10.0.4/24      n2
e4a      false
4 entries were displayed.
```

5. Disconnect the cable from e4e on node n1.

You can refer to the running configuration and connect the first 40 GbE port on switch C2 (port 1/7 in this example) to e4e on node n1, using the appropriate cabling for the Nexus 3232C switch model.

6. Disconnect the cable from e4e on node n2.

You can refer to the running configuration and connect e4e to the next available 40 GbE port on C2, port 1/8, using the appropriate cabling for the Nexus 3232C switch model.

7. Enable all node-facing ports on C2.

Show example

The following example shows ports 1 through 30 being enabled on Nexus 3132Q-V cluster switches C1 and C2 using a configuration supported in RCF NX3232C_RCF_v1.0_24p10g_26p100g.txt:

```
C2# configure
C2(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C2(config-if-range)# no shutdown
C2(config-if-range)# exit
C2(config)# exit
```

8. Enable the second cluster port, e4e, on each node:

```
network port modify
```

Show example

The following example shows the second cluster port e4e being brought up on each node:

```
cluster::*> network port modify -node n1 -port e4e -up-admin true
cluster::*> *network port modify -node n2 -port e4e -up-admin true*s
```

9. For each node, revert all of the migrated cluster interconnect LIFs: `network interface revert`

Show example

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

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

10. Verify that all of the cluster interconnect ports are now reverted to their home ports:

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

The `Is Home` column should display a value of `true` for all of the ports listed in the `Current Port` column. If the displayed value is `false`, the port has not been reverted.

Show example

```
cluster::*> network interface show -role cluster
(network interface show)
      Logical      Status      Network      Current
Current Is
Vserver      Interface  Admin/Oper Address/Mask      Node
Port      Home
-----
Cluster
e4a      true      n1_clus1      up/up      10.10.0.1/24      n1
e4e      true      n1_clus2      up/up      10.10.0.2/24      n1
e4a      true      n2_clus1      up/up      10.10.0.3/24      n2
e4e      true      n2_clus2      up/up      10.10.0.4/24      n2
4 entries were displayed.
```

11. Verify that all of the cluster interconnect ports are in the up state:

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

12. Display the cluster switch port numbers through which each cluster port is connected to each node:

```
network device-discovery show
```

Show example

```
cluster::*> network device-discovery show
      Local   Discovered
Node      Port   Device      Interface      Platform
-----
n1      /cdp
e4a      C1      Ethernet1/7      N3K-C3232C
e4e      C2      Ethernet1/7      N3K-C3232C
n2      /cdp
e4a      C1      Ethernet1/8      N3K-C3232C
e4e      C2      Ethernet1/8      N3K-C3232C
```

13. Display discovered and monitored cluster switches:

```
system cluster-switch show
```

Show example

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

Switch Model	Type	Address

C1 NX3232CV Serial Number: FOX000001 Is Monitored: true Reason: Software Version: Cisco Nexus Operating System (NX-OS) Software, Version 7.0(3)I6(1) Version Source: CDP	cluster-network	10.10.1.101
C2 NX3232CV Serial Number: FOX000002 Is Monitored: true Reason: Software Version: Cisco Nexus Operating System (NX-OS) Software, Version 7.0(3)I6(1) Version Source: CDP 2 entries were displayed.	cluster-network	10.10.1.102

14. Verify that switchless cluster detection changed the switchless cluster option to disabled:

```
network options switchless-cluster show
```

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

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

Show example

```
cluster::*> cluster ping-cluster -node n1
Host is n1 Getting addresses from network interface table...
Cluster n1_clus1 n1          e4a    10.10.0.1
Cluster n1_clus2 n1          e4e    10.10.0.2
Cluster n2_clus1 n2          e4a    10.10.0.3
Cluster n2_clus2 n2          e4e    10.10.0.4
Local = 10.10.0.1 10.10.0.2
Remote = 10.10.0.3 10.10.0.4
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s) .....
Detected 9000 byte MTU on 32 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 10.10.0.2 to Remote 10.10.0.3
    Local 10.10.0.2 to Remote 10.10.0.4
Larger than PMTU communication succeeds on 4 path(s) RPC status:
1 paths up, 0 paths down (tcp check)
1 paths up, 0 paths down (ucp check)
```

16. 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
```

Show example

```
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>

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



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

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

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

Copyright information

Copyright © 2024 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, 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) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.