Replace switches

Cluster and storage switches

NetApp
April 25, 2024

Table of Contents

Replace switches

Replace a Cisco Nexus 3232C cluster switch

Follow these steps to replace a defective Cisco Nexus 3232C switch in a cluster. This is a non-disruptive procedure.

Review requirements

What you'll need

Make sure that the existing cluster and network configuration has the following characteristics:

- The Nexus 3232C cluster infrastructure are redundant and fully functional on both switches.

  The Cisco Ethernet Switches page has the latest RCF and NX-OS versions on your switches.

- All cluster ports must be in the **up** state.
- Management connectivity must exist on both switches.
- All cluster logical interfaces (LIFs) are in the **up** state and are not migrated.

The replacement Cisco Nexus 3232C switch has the following characteristics:

- Management network connectivity is functional.
- Console access to the replacement switch is in place.
- The appropriate RCF and NX-OS operating system image is loaded onto the switch.
- Initial customization of the switch is complete.

For more information

See the following:

- Cisco Ethernet Switch description page
- Hardware Universe

Replace the switch

About this task

This replacement procedure describes the following scenario:

- The cluster initially has four nodes connected to two Nexus 3232C cluster switches, CL1 and CL2.
- You plan to replace cluster switch CL2 with C2 (steps 1 to 21):
  - On each node, you migrate the cluster LIFs connected to cluster switch CL2 to cluster ports connected to cluster switch CL1.
  - You disconnect the cabling from all ports on cluster switch CL2 and reconnect the cabling to the same ports on the replacement cluster switch C2.
  - You revert the migrated cluster LIFs on each node.

About the examples

This replacement procedure replaces the second Nexus 3232C cluster switch CL2 with the new 3232C switch C2.

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

- The four nodes are n1, n2, n3, and n4.
- n1_clus1 is the first cluster logical interface (LIF) connected to cluster switch C1 for node n1.
- n1_clus2 is the first cluster LIF connected to cluster switch CL2 or C2 for node n1.
- n1_clus3 is the second LIF connected to cluster switch C2 for node n1.-
- n1_clus4 is the second LIF connected to cluster switch CL1, 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 examples in this replacement procedure use four nodes. Two of the nodes use four 10 GB cluster interconnect ports: e0a, e0b, e0c, and e0d. The other two nodes use two 40 GB cluster interconnect ports: e4a and e4e. See the Hardware Universe to verify the correct cluster ports for your platform.

Step 1: Display and migrate the cluster ports to switch

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

```
cluster::> network device-discovery show
            Local  Discovered
Node        Port   Device                Interface         Platform
----------- ------ ------------------- ----------------
----------------
n1         /cdp
            e0a    CL1                   Ethernet1/1/1    N3K-C3232C
            e0b    CL2                   Ethernet1/1/1    N3K-C3232C
            e0c    CL2                   Ethernet1/1/2    N3K-C3232C
            e0d    CL1                   Ethernet1/1/2    N3K-C3232C

n2         /cdp
            e0a    CL1                   Ethernet1/1/3    N3K-C3232C
            e0b    CL2                   Ethernet1/1/3    N3K-C3232C
            e0c    CL2                   Ethernet1/1/4    N3K-C3232C
            e0d    CL1                   Ethernet1/1/4    N3K-C3232C

n3         /cdp
            e4a    CL1                   Ethernet1/7      N3K-C3232C
            e4e    CL2                   Ethernet1/7      N3K-C3232C

n4         /cdp
            e4a    CL1                   Ethernet1/8      N3K-C3232C
            e4e    CL2                   Ethernet1/8      N3K-C3232C
```

3. 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
                                                        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 -
-

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  -
```

b. Display information about the logical interfaces (LIFs):

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

Show example

```
cluster::*> network interface show -role cluster
             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
e0a     true
            n3_clus2   up/up      10.10.0.10/24      n3
e0e     true
            n4_clus1   up/up      10.10.0.11/24      n4
e0a     true
            n4_clus2   up/up      10.10.0.12/24      n4
e0e     true
```

c. Display the discovered cluster switches:

```
system cluster-switch show
```

Show example

The following output example displays the cluster switches:

```
cluster::> system cluster-switch show
Switch                      Type               Address
Model
--------------------------- ------------------ ----------------
---------------
CL1                         cluster-network    10.10.1.101
NX3232C
        Serial Number: FOX000001
         Is Monitored: true
               Reason: None
     Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version 7.0(3)I6(1)
       Version Source: CDP

CL2                         cluster-network    10.10.1.102
NX3232C
        Serial Number: FOX000002
         Is Monitored: true
               Reason: None
     Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version 7.0(3)I6(1)
       Version Source: CDP
```

4. Verify that the appropriate RCF and image are installed on the new Nexus 3232C switch and make any necessary site customizations.
   a. Go to the NetApp Support Site.

      mysupport.netapp.com

   b. Go to the **Cisco Ethernet Switches** page and note the required software versions in the table.

      Cisco Ethernet Switches

   c. Download the appropriate version of the RCF.
   d. Click **CONTINUE** on the **Description** page, accept the license agreement, and then navigate to the **Download** page.
   e. Download the correct version of the image software from the **Cisco® Cluster and Management Network Switch Reference Configuration File Download** page.

      Cisco® Cluster and Management Network Switch Reference Configuration File Download

5. Migrate the cluster LIFs to the physical node ports connected to the replacement switch C2:

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

Show example

You must migrate all the cluster LIFs 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
cluster::*> network interface migrate -vserver Cluster -lif n3_clus2
-source-node n3 -destination-
node n3 -destination-port e4a
cluster::*> network interface migrate -vserver Cluster -lif n4_clus2
-source-node n4 -destination-
node n4 -destination-port e4a
```

6. Verify the status of the cluster ports and their home designations:

```
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
e0a     true
            n1_clus2   up/up      10.10.0.2/24       n1
e0a     false
            n1_clus3   up/up      10.10.0.3/24       n1
e0d     false
            n1_clus4   up/up      10.10.0.4/24       n1
e0d     true
            n2_clus1   up/up      10.10.0.5/24       n2
e0a     true
            n2_clus2   up/up      10.10.0.6/24       n2
e0a     false
            n2_clus3   up/up      10.10.0.7/24       n2
e0d     false
            n2_clus4   up/up      10.10.0.8/24       n2
e0d     true
            n3_clus1   up/up      10.10.0.9/24       n3
e4a     true
            n3_clus2   up/up      10.10.0.10/24      n3
e4a     false
            n4_clus1   up/up      10.10.0.11/24      n4
e4a     true
            n4_clus2   up/up      10.10.0.12/24      n4
e4a     false
```

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

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

Show example

The following example shows the cluster interconnect ports are 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
cluster::*> network port modify -node n3 -port e4e -up-admin false
cluster::*> network port modify -node n4 -port e4e -up-admin false
```

8. 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
Cluster n3_clus1 n4        e0a    10.10.0.9
Cluster n3_clus2 n3        e0e    10.10.0.10
Cluster n4_clus1 n4        e0a    10.10.0.11
Cluster n4_clus2 n4        e0e    10.10.0.12
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 10.10.0.9
10.10.0.10 10.10.0.11
10.10.0.12 Cluster Vserver Id = 4294967293 Ping status:
....
Basic connectivity succeeds on 32 path(s)
Basic connectivity fails on 0 path(s) ................
Detected 9000 byte MTU on 32 path(s):
    Local 10.10.0.1 to Remote 10.10.0.5
    Local 10.10.0.1 to Remote 10.10.0.6
    Local 10.10.0.1 to Remote 10.10.0.7
    Local 10.10.0.1 to Remote 10.10.0.8
    Local 10.10.0.1 to Remote 10.10.0.9
    Local 10.10.0.1 to Remote 10.10.0.10
    Local 10.10.0.1 to Remote 10.10.0.11
    Local 10.10.0.1 to Remote 10.10.0.12
    Local 10.10.0.2 to Remote 10.10.0.5
    Local 10.10.0.2 to Remote 10.10.0.6
    Local 10.10.0.2 to Remote 10.10.0.7
    Local 10.10.0.2 to Remote 10.10.0.8
    Local 10.10.0.2 to Remote 10.10.0.9
    Local 10.10.0.2 to Remote 10.10.0.10
    Local 10.10.0.2 to Remote 10.10.0.11
    Local 10.10.0.2 to Remote 10.10.0.12
    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 10.10.0.3 to Remote 10.10.0.9
    Local 10.10.0.3 to Remote 10.10.0.10
    Local 10.10.0.3 to Remote 10.10.0.11
    Local 10.10.0.3 to Remote 10.10.0.12
    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
    Local 10.10.0.4 to Remote 10.10.0.9
    Local 10.10.0.4 to Remote 10.10.0.10
    Local 10.10.0.4 to Remote 10.10.0.11
    Local 10.10.0.4 to Remote 10.10.0.12
Larger than PMTU communication succeeds on 32 path(s) RPC status:
8 paths up, 0 paths down (tcp check)
8   paths up, 0 paths down (udp check)
```

Step 2: Migrate ISLs to switch CL1 and C2

1. Shut down the ports 1/31 and 1/32 on cluster switch CL1.

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

   Show example

```
(CL1)# configure
(CL1)(Config)# interface e1/31-32
(CL1)(config-if-range)# shutdown
(CL1)(config-if-range)# exit
(CL1)(Config)# exit
(CL1)#
```

2. Remove all the cables attached to the cluster switch CL2 and reconnect them to the replacement switch C2 for all the nodes.
3. Remove the inter-switch link (ISL) cables from ports e1/31 and e1/32 on cluster switch CL2 and reconnect them to the same ports on the replacement switch C2.
4. Bring up ISL ports 1/31 and 1/32 on the cluster switch CL1.

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

Show example

```
(CL1)# configure
(CL1)(Config)# interface e1/31-32
(CL1)(config-if-range)# no shutdown
(CL1)(config-if-range)# exit
(CL1)(Config)# exit
(CL1)#
```

5. Verify that the ISLs are up on CL1.

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

   Ports Eth1/31 and Eth1/32 should indicate `(P)`, which means that the ISL ports are up in the port-channel:

   Show example

```
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      Eth1/31(P)   Eth1/32(P)
```

6. Verify that the ISLs are up on cluster switch C2.

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

Show example

Ports Eth1/31 and Eth1/32 should indicate (P), which means that both ISL ports are up in the port-channel.

```
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)
```

7. On all nodes, bring up all the cluster interconnect ports connected to the replacement switch C2:

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

   Show example

```
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
cluster::*> network port modify -node n3 -port e4e -up-admin true
cluster::*> network port modify -node n4 -port e4e -up-admin true
```

Step 3: Revert all LIFs to originally assigned ports

1. Revert all the migrated cluster interconnect LIFs on all the nodes:

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

Show example

You must revert all the cluster interconnect LIFs individually as shown in the following 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
Cluster::*> network interface revert -vserver cluster -lif n3_clus2
Cluster::*> network interface revert -vserver cluster -lif n4_clus2
```

2. Verify that the cluster interconnect ports are now reverted to their home:

   ```
   network interface show
   ```

Show example

The following example shows that all the LIFs have been successfully reverted because the ports listed under the `Current Port` column have a status of `true` in the `Is Home` column. If a port has a value of `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
             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
```

3. Verify that the cluster ports are connected:

   `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
--------- ------------ ---------------- ---- ---- -----------
-------- -----
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  -
-
```

4. 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
Cluster n3_clus1 n3        e0a    10.10.0.9
Cluster n3_clus2 n3        e0e    10.10.0.10
Cluster n4_clus1 n4        e0a    10.10.0.11
Cluster n4_clus2 n4        e0e    10.10.0.12
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 10.10.0.9
10.10.0.10 10.10.0.11 10.10.0.12
Cluster Vserver Id = 4294967293 Ping status:
....
Basic connectivity succeeds on 32 path(s)
Basic connectivity fails on 0 path(s) ................
Detected 1500 byte MTU on 32 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 10.10.0.1 to Remote 10.10.0.9
    Local 10.10.0.1 to Remote 10.10.0.10
    Local 10.10.0.1 to Remote 10.10.0.11
    Local 10.10.0.1 to Remote 10.10.0.12
    Local 10.10.0.2 to Remote 10.10.0.5
    Local 10.10.0.2 to Remote 10.10.0.6
    Local 10.10.0.2 to Remote 10.10.0.7
    Local 10.10.0.2 to Remote 10.10.0.8
    Local 10.10.0.2 to Remote 10.10.0.9
    Local 10.10.0.2 to Remote 10.10.0.10
    Local 10.10.0.2 to Remote 10.10.0.11
    Local 10.10.0.2 to Remote 10.10.0.12
    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 10.10.0.3 to Remote 10.10.0.9
    Local 10.10.0.3 to Remote 10.10.0.10
    Local 10.10.0.3 to Remote 10.10.0.11
    Local 10.10.0.3 to Remote 10.10.0.12
    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
    Local 10.10.0.4 to Remote 10.10.0.9
    Local 10.10.0.4 to Remote 10.10.0.10
    Local 10.10.0.4 to Remote 10.10.0.11
    Local 10.10.0.4 to Remote 10.10.0.12
Larger than PMTU communication succeeds on 32 path(s) RPC status:
8 paths up, 0 paths down (tcp check)
8   paths up, 0 paths down (udp check)
```

Step 4: Verify all ports and LIF are correctly migrated

1. Display the information about the devices in your configuration by entering the following commands:

   You can execute the following commands in any order:

   - `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
            Local  Discovered
Node        Port   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

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 -

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  -

cluster::*> network interface show -role cluster

            Logical    Status     Network            Current
Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node
Port    Home
----------- ---------- ---------- ------------------ -------------
------- ----
Cluster
            nm1_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

cluster::*> system cluster-switch show
Switch                      Type               Address
Model
--------------------------- ------------------ ----------------
---------------
CL1                          cluster-network   10.10.1.101
NX3232C
             Serial Number: FOX000001
              Is Monitored: true
                    Reason: None
          Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version 7.0(3)I6(1)
            Version Source: CDP
CL2                          cluster-network   10.10.1.102
NX3232C
             Serial Number: FOX000002
              Is Monitored: true
                    Reason: None
          Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version 7.0(3)I6(1)
            Version Source: CDP

C2                          cluster-network    10.10.1.103
NX3232C
             Serial Number: FOX000003
```

```
                Is Monitored: true
                      Reason: None
            Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version 7.0(3)I6(1)
            Version Source: CDP 3 entries were displayed.
```

2. Delete the replaced cluster switch CL2 if it has not been removed automatically:

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

3. Verify that the proper cluster switches are monitored:

   ```
   system cluster-switch show
   ```

   Show example

   The following example shows the cluster switches are monitored because the `Is Monitored` state is `true`.

   ```
   cluster::> system cluster-switch show
   Switch                      Type               Address
   Model
   --------------------------- ------------------ ----------------
   ---------------
   CL1                         cluster-network    10.10.1.101
   NX3232C
               Serial Number: FOX000001
                Is Monitored: true
                      Reason: None
            Software Version: Cisco Nexus Operating System (NX-OS)
   Software, Version 7.0(3)I6(1)
              Version Source: CDP

   C2                          cluster-network    10.10.1.103
   NX3232C
               Serial Number: FOX000002
                Is Monitored: true
                      Reason: None
            Software Version: Cisco Nexus Operating System (NX-OS)
   Software, Version 7.0(3)I6(1)
              Version Source: CDP
   ```

4. 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:
CL1
C2

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

Enter the switch name: CL1
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.

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

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

Replace a Cisco Nexus 3232C storage switch

Follow these steps to replace a defective Cisco Nexus 3232C storage switch. This is a non-disruptive procedure.

Review requirements

The existing network configuration must have the following characteristics:

- The Cisco Ethernet Switches page has the latest RCF and NX-OS versions on your switches.
- Management connectivity must exist on both switches.

> Make sure that all troubleshooting steps have been completed to confirm that your switch needs replacing.

The replacement Cisco Nexus 3232C switch must have the following characteristics:

- Management network connectivity must be functional.
- Console access to the replacement switch must be in place.
- The appropriate RCF and NX-OS operating system image must be loaded onto the switch.
- Initial customization of the switch must be complete.

Replace the switch

This procedure replaces the second Nexus 3232C storage switch S2 with the new 3232C switch NS2. The two nodes are node1 and node2.

Step 1: Confirm the switch to be replaced is S2

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. Check on the health status of the storage node ports to make sure that there is connection to storage switch S1:

   ```
   storage port show -port-type ENET
   ```

Show example

```
storage::*> storage port show -port-type ENET
                                          Speed                      VLAN
Node                Port Type  Mode    (Gb/s) State    Status      ID
------------------ ---- ----- ------- ------ -------- --------- ----
node1
                    e3a  ENET  storage    100 enabled  online      30
                    e3b  ENET  storage      0 enabled  offline     30
                    e7a  ENET  storage      0 enabled  offline     30
                    e7b  ENET  storage      0 enabled  offline     30
node2
                    e3a  ENET  storage    100 enabled  online      30
                    e3b  ENET  storage      0 enabled  offline     30
                    e7a  ENET  storage      0 enabled  offline     30
                    e7b  ENET  storage      0 enabled  offline     30
```

3. Verify that storage switch S1 is available:

```
network device-discovery show
```

Show example

```
storage::*> network device-discovery show
Node/       Local  Discovered
Protocol    Port   Device (LLDP: ChassisID)  Interface
Platform
----------- ------ ------------------------- -----------------
----------------
node1/cdp
            e3a    S1                        Ethernet1/1
NX3232C
            e4a    node2                     e4a               AFF-
A700
            e4e    node2                     e4e               AFF-
A700
node1/lldp
            e3a    S1                        Ethernet1/1       -
            e4a    node2                     e4a               -
            e4e    node2                     e4e               -
node2/cdp
            e3a    S1                        Ethernet1/2
NX3232C
            e4a    node1                     e4a               AFF-
A700
            e4e    node1                     e4e               AFF-
A700
node2/lldp
            e3a    S1                        Ethernet1/2       -
            e4a    node1                     e4a               -
            e4e    node1                     e4e               -
```

4. Run the `show lldp neighbors` command on the working switch to confirm that you can see both nodes and all shelves:

   `show lldp neighbors`

Show example

```
S1# show lldp neighbors
Capability codes:
  (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
  (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other
Device ID               Local Intf      Hold-time  Capability  Port
ID
node1                   Eth1/1          121        S           e3a
node2                   Eth1/2          121        S           e3a
SHFGD2008000011         Eth1/5          121        S           e0a
SHFGD2008000011         Eth1/6          120        S           e0a
SHFGD2008000022         Eth1/7          120        S           e0a
SHFGD2008000022         Eth1/8          120        S           e0a
```

Step 2: Configure cabling

1. Verify the shelf ports in the storage system:

```
storage shelf port show -fields remote-device,remote-port
```

Show example

```
storage::*> storage shelf port show -fields remote-device,remote-
port

shelf  id  remote-port  remote-device
-----  --  -----------  -------------
3.20   0   Ethernet1/5  S1
3.20   1   -            -
3.20   2   Ethernet1/6  S1
3.20   3   -            -
3.30   0   Ethernet1/7  S1
3.20   1   -            -
3.30   2   Ethernet1/8  S1
3.20   3   -            -
```

2. Remove all cables attached to storage switch S2.
3. Reconnect all cables to the replacement switch NS2.

Step 3: Verify all device configurations on switch NS2

1. Verify the health status of the storage node ports:

```
storage port show -port-type ENET
```

Show example

```
storage::*> storage port show -port-type ENET
                                          Speed
VLAN
Node                Port Type  Mode    (Gb/s) State    Status
ID
------------------ ---- ----- ------- ------ -------- ------------
---
node1
                    e3a  ENET  storage    100 enabled  online
30
                    e3b  ENET  storage      0 enabled  offline
30
                    e7a  ENET  storage      0 enabled  offline
30
                    e7b  ENET  storage    100 enabled  online
30
node2
                    e3a  ENET  storage    100 enabled  online
30
                    e3b  ENET  storage      0 enabled  offline
30
                    e7a  ENET  storage      0 enabled  offline
30
                    e7b  ENET  storage    100 enabled  online
30
```

2. Verify that both switches are available:

```
network device-discovery show
```

Show example

```
storage::*> network device-discovery show
Node/       Local  Discovered
Protocol    Port   Device (LLDP: ChassisID)  Interface
Platform
----------- ------ ------------------------- ----------------
--------
node1/cdp
            e3a    S1                        Ethernet1/1
NX3232C
            e4a    node2                     e4a               AFF-
A700
            e4e    node2                     e4e               AFF-
A700
            e7b    NS2                       Ethernet1/1
NX3232C
node1/lldp
            e3a    S1                        Ethernet1/1       -
            e4a    node2                     e4a               -
            e4e    node2                     e4e               -
            e7b    NS2                       Ethernet1/1       -
node2/cdp
            e3a    S1                        Ethernet1/2
NX3232C
            e4a    node1                     e4a               AFF-
A700
            e4e    node1                     e4e               AFF-
A700
            e7b    NS2                       Ethernet1/2
NX3232C
node2/lldp
            e3a    S1                        Ethernet1/2       -
            e4a    node1                     e4a               -
            e4e    node1                     e4e               -
            e7b    NS2                       Ethernet1/2       -
```

3. Verify the shelf ports in the storage system:

```
storage shelf port show -fields remote-device,remote-port
```

Show example

```
storage::*> storage shelf port show -fields remote-device,remote-
port
shelf id remote-port remote-device
----- -- ----------- -------------
3.20  0  Ethernet1/5 S1
3.20  1  Ethernet1/5 NS2
3.20  2  Ethernet1/6 S1
3.20  3  Ethernet1/6 NS2
3.30  0  Ethernet1/7 S1
3.20  1  Ethernet1/7 NS2
3.30  2  Ethernet1/8 S1
3.20  3  Ethernet1/8 NS2
```

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

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

Replace Cisco Nexus 3232C cluster switches with switchless connections

You can migrate from a cluster with a switched cluster network to one where two nodes are directly connected for ONTAP 9.3 and later.

Review requirements

Guidelines

Review the following guidelines:

- Migrating to a two-node switchless cluster configuration is a nondisruptive operation. Most systems have two dedicated cluster interconnect ports on each node, but you can also use this procedure for systems with a larger number of dedicated cluster interconnect ports on each node, such as four, six or eight.
- You cannot use the switchless cluster interconnect feature with more than two nodes.
- If you have an existing two-node cluster that uses cluster interconnect switches and is running ONTAP 9.3 or later, you can replace the switches with direct, back-to-back connections between the nodes.

What you'll need

- A healthy cluster that consists of two nodes connected by cluster switches. The nodes must be running the same ONTAP release.
- Each node with the required number of dedicated cluster ports, which provide redundant cluster interconnect connections to support your system configuration. For example, there are two redundant ports for a system with two dedicated cluster interconnect ports on each node.

Migrate the switches

About this task

The following procedure removes the cluster switches in a two-node cluster and replaces each connection to the switch with a direct connection to the partner node.

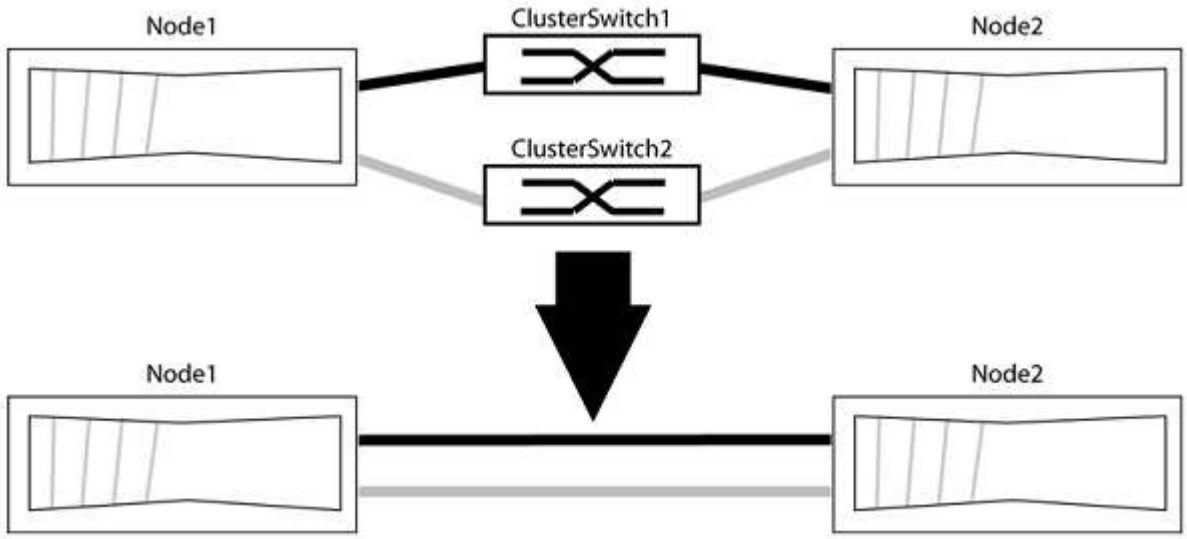

About the examples

The examples in the following procedure show nodes that are using "e0a" and "e0b" as cluster ports. Your nodes might be using different cluster ports as they vary by system.

Step 1: Prepare for migration

1. Change the privilege level to advanced, entering `y` when prompted to continue:

   ```
   set -privilege advanced
   ```

   The advanced prompt `*>` appears.

2. ONTAP 9.3 and later supports automatic detection of switchless clusters, which is enabled by default.

   You can verify that detection of switchless clusters is enabled by running the advanced privilege command:

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

   Show example

   The following example output shows if the option is enabled.

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

   If "Enable Switchless Cluster Detection" is `false`, contact NetApp support.

3. 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=<number_of_hours>h
   ```

   where `h` is the duration of the maintenance window in hours. The message notifies technical support of this maintenance task so that they can suppress automatic case creation during the maintenance window.

   In the following example, the command suppresses automatic case creation for two hours:

   Show example

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

Step 2: Configure ports and cabling

1. Organize the cluster ports on each switch into groups so that the cluster ports in group1 go to cluster switch1 and the cluster ports in group2 go to cluster switch2. These groups are required later in the procedure.
2. Identify the cluster ports and verify link status and health:

   ```
   network port show -ipspace Cluster
   ```

   In the following example for nodes with cluster ports "e0a" and "e0b", one group is identified as “node1:e0a” and “node2:e0a” and the other group as “node1:e0b” and “node2:e0b”. Your nodes might be using different cluster ports because they vary by system.

   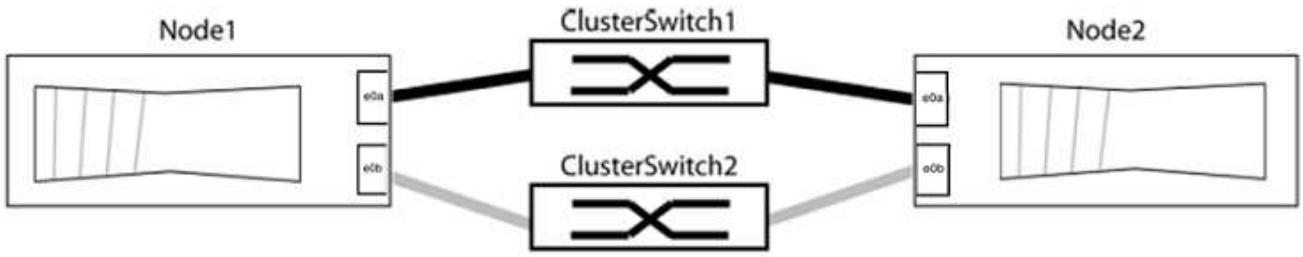

   Verify that the ports have a value of `up` for the “Link” column and a value of `healthy` for the “Health Status” column.

Show example

```
cluster::> network port show -ipspace Cluster
Node: node1

Ignore
                                                  Speed(Mbps) Health
Health
Port  IPspace   Broadcast Domain Link  MTU   Admin/Oper  Status
Status
----- --------- ---------------- ----- ----- ----------- -------
-------
e0a   Cluster   Cluster          up    9000  auto/10000  healthy
false
e0b   Cluster   Cluster          up    9000  auto/10000  healthy
false

Node: node2

Ignore
                                                  Speed(Mbps) Health
Health
Port  IPspace   Broadcast Domain Link  MTU   Admin/Oper  Status
Status
----- --------- ---------------- ----- ----- ----------- -------
-------
e0a   Cluster   Cluster          up    9000  auto/10000  healthy
false
e0b   Cluster   Cluster          up    9000  auto/10000  healthy
false
4 entries were displayed.
```

3. Confirm that all the cluster LIFs are on their home ports.

   Verify that the “is-home” column is `true` for each of the cluster LIFs:

   `network interface show -vserver Cluster -fields is-home`

Show example

```
cluster::*> net int show -vserver Cluster -fields is-home
(network interface show)
vserver  lif          is-home
-------- ------------ --------
Cluster  node1_clus1  true
Cluster  node1_clus2  true
Cluster  node2_clus1  true
Cluster  node2_clus2  true
4 entries were displayed.
```

If there are cluster LIFs that are not on their home ports, revert those LIFs to their home ports:

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

4. Disable auto-revert for the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

5. Verify that all ports listed in the previous step are connected to a network switch:

```
network device-discovery show -port cluster_port
```

The “Discovered Device” column should be the name of the cluster switch that the port is connected to.

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to cluster switches "cs1" and "cs2".

```
cluster::> network device-discovery show -port e0a|e0b
  (network device-discovery show)
Node/     Local  Discovered
Protocol  Port   Device (LLDP: ChassisID)  Interface  Platform
--------- ------ ------------------------- ---------- ----------
node1/cdp
          e0a    cs1                       0/11       BES-53248
          e0b    cs2                       0/12       BES-53248
node2/cdp
          e0a    cs1                       0/9        BES-53248
          e0b    cs2                       0/9        BES-53248
4 entries were displayed.
```

6. Verify the cluster connectivity:

```
cluster ping-cluster -node local
```

7. Verify that the cluster is healthy:

```
cluster ring show
```

All units must be either master or secondary.

8. Set up the switchless configuration for the ports in group 1.

> To avoid potential networking issues, you must disconnect the ports from group1 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.

a. Disconnect all the cables from the ports in group1 at the same time.

In the following example, the cables are disconnected from port "e0a" on each node, and cluster traffic continues through the switch and port "e0b" on each node:

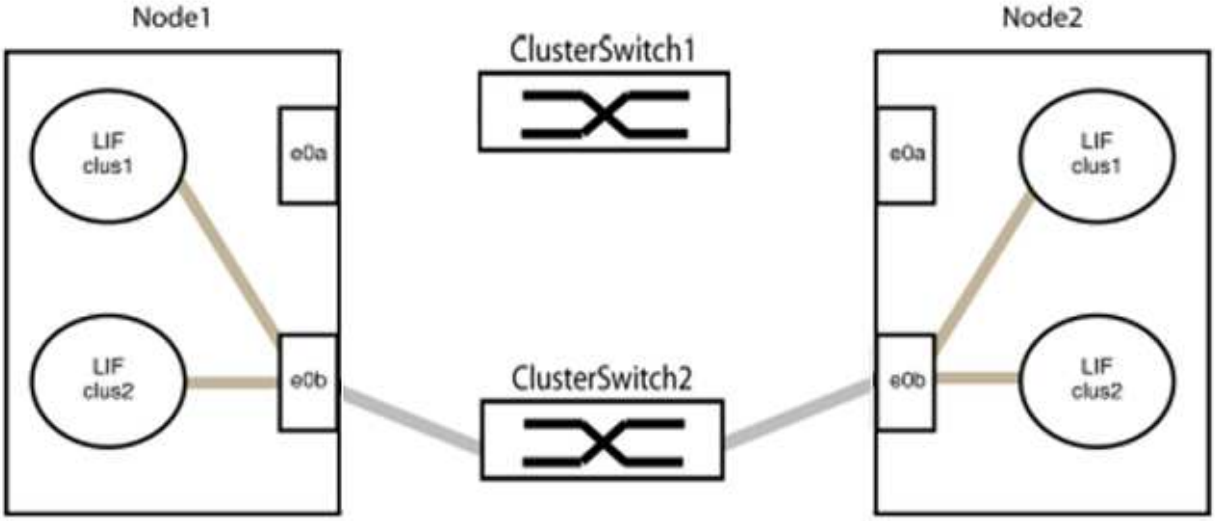

b. Cable the ports in group1 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2:

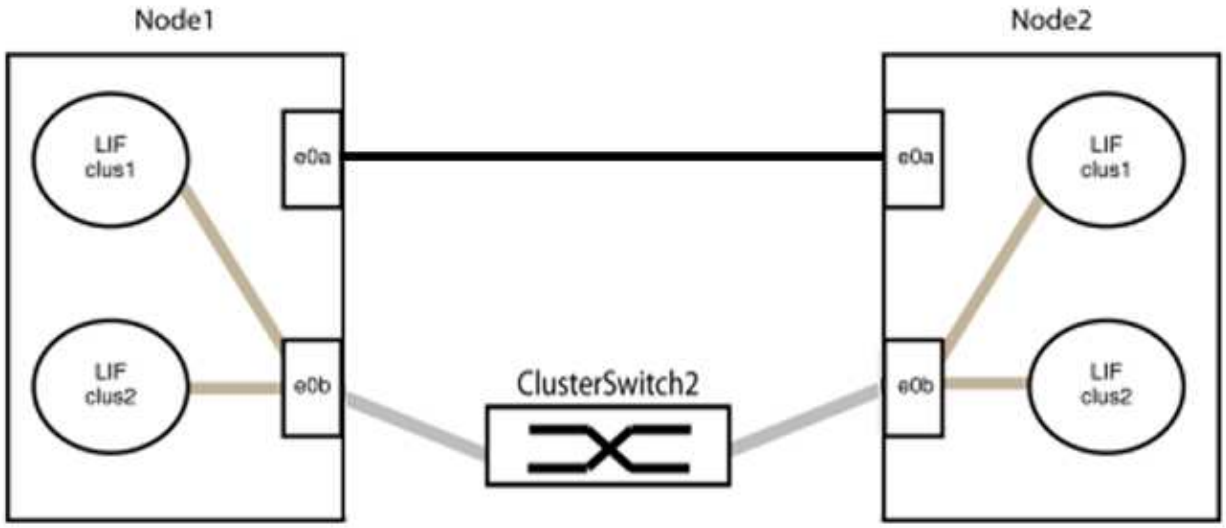

9. The switchless cluster network option transitions from `false` to `true`. This might take up to 45 seconds. Confirm that the switchless option is set to `true`:

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

   The following example shows that the switchless cluster is enabled:

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

10. Verify that the cluster network is not disrupted:

    ```
    cluster ping-cluster -node local
    ```

    Before proceeding to the next step, you must wait at least two minutes to confirm a working back-to-back connection on group 1.

11. Set up the switchless configuration for the ports in group 2.

    To avoid potential networking issues, you must disconnect the ports from group2 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.

    a. Disconnect all the cables from the ports in group2 at the same time.

       In the following example, the cables are disconnected from port "e0b" on each node, and cluster traffic continues through the direct connection between the "e0a" ports:

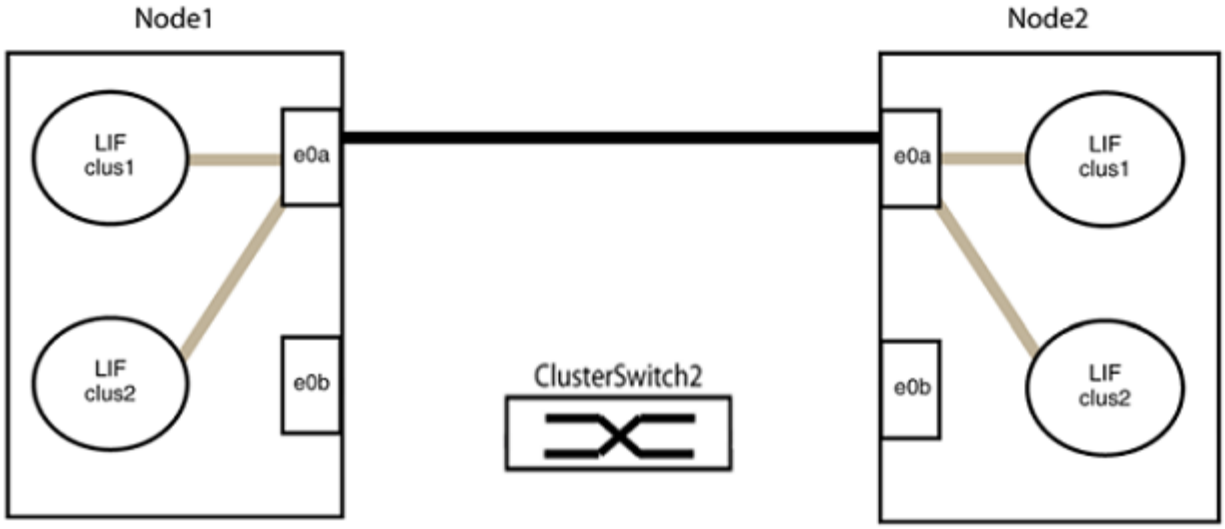

b. Cable the ports in group2 back-to-back.

   In the following example, "e0a" on node1 is connected to "e0a" on node2 and "e0b" on node1 is connected to "e0b" on node2:

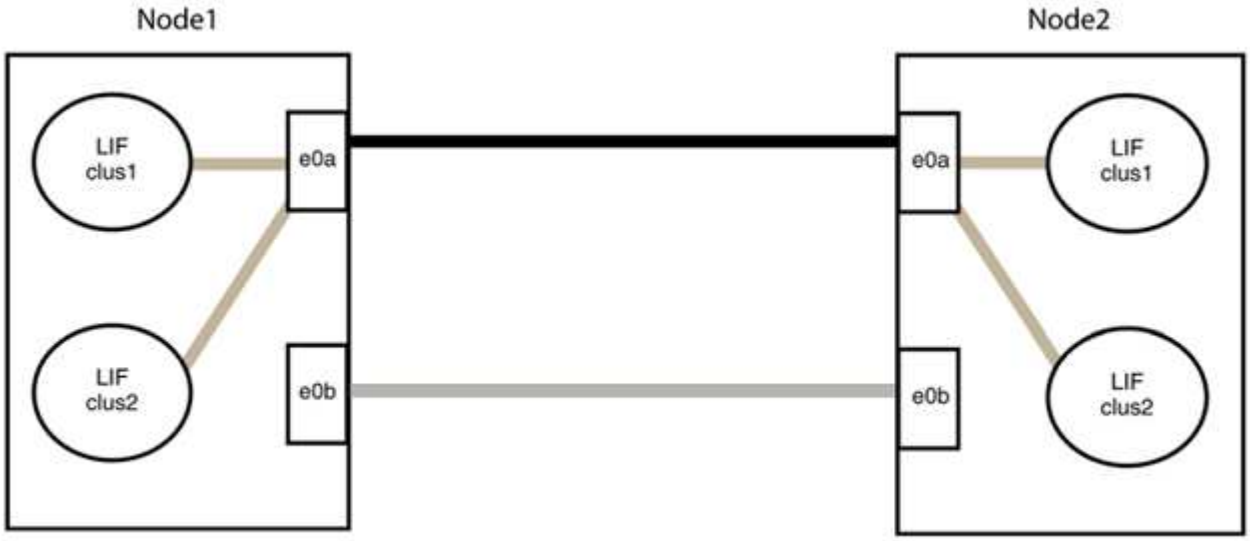

Step 3: Verify the configuration

1. Verify that the ports on both nodes are correctly connected:

   `network device-discovery show -port` *`cluster_port`*

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to the corresponding port on the cluster partner:

```
cluster::> net device-discovery show -port e0a|e0b
  (network device-discovery show)
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
---------- ------ ------------------------- ---------- ----------
node1/cdp
           e0a    node2                     e0a        AFF-A300
           e0b    node2                     e0b        AFF-A300
node1/lldp
           e0a    node2 (00:a0:98:da:16:44) e0a        -
           e0b    node2 (00:a0:98:da:16:44) e0b        -
node2/cdp
           e0a    node1                     e0a        AFF-A300
           e0b    node1                     e0b        AFF-A300
node2/lldp
           e0a    node1 (00:a0:98:da:87:49) e0a        -
           e0b    node1 (00:a0:98:da:87:49) e0b        -
8 entries were displayed.
```

2. Re-enable auto-revert for the cluster LIFs:

   `network interface modify -vserver Cluster -lif * -auto-revert true`

3. Verify that all LIFs are home. This might take a few seconds.

   `network interface show -vserver Cluster -lif` *`lif_name`*

Show example

The LIFs have been reverted if the “Is Home” column is `true`, as shown for `node1_clus2` and `node2_clus2` in the following example:

```
cluster::> network interface show -vserver Cluster -fields curr-
port,is-home
vserver  lif           curr-port is-home
-------- ------------- --------- -------
Cluster  node1_clus1   e0a       true
Cluster  node1_clus2   e0b       true
Cluster  node2_clus1   e0a       true
Cluster  node2_clus2   e0b       true
4 entries were displayed.
```

If any cluster LIFS have not returned to their home ports, revert them manually from the local node:

`network interface revert -vserver Cluster -lif` *`lif_name`*

4. Check the cluster status of the nodes from the system console of either node:

   `cluster show`

   Show example

   The following example shows epsilon on both nodes to be `false`:

   ```
   Node  Health  Eligibility Epsilon
   ----- ------- ----------- --------
   node1 true    true        false
   node2 true    true        false
   2 entries were displayed.
   ```

5. Confirm connectivity between the cluster ports:

   `cluster ping-cluster local`

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

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

   For more information, see NetApp KB Article 1010449: How to suppress automatic case creation during scheduled maintenance windows.

7. Change the privilege level back to admin:

```
set -privilege admin
```