NetApp

Migrate switches

Cluster and storage switches

NetApp
April 05, 2024

Table of Contents

Migrate switches

Migrate CN1610 cluster switches to BES-53248 cluster switches

To migrate the CN1610 cluster switches in a cluster to Broadcom-supported BES-53248 cluster switches, review the migration requirements and then follow the migration procedure.

The following cluster switches are supported:

- CN1610
- BES-53248

Review requirements

Verify that your configuration meets the following requirements:

- Some of the ports on BES-53248 switches are configured to run at 10GbE.
- The 10GbE connectivity from nodes to BES-53248 cluster switches have been planned, migrated, and documented.
- The cluster is fully functioning (there should be no errors in the logs or similar issues).
- Initial customization of the BES-53248 switches is complete, so that:
    - BES-53248 switches are running the latest recommended version of EFOS software.
    - Reference Configuration Files (RCFs) have been applied to the switches.
    - Any site customization, such as DNS, NTP, SMTP, SNMP, and SSH, are configured on the new switches.

Node connections

The cluster switches support the following node connections:

- NetApp CN1610: ports 0/1 through 0/12 (10GbE)
- BES-53248: ports 0/1-0/16 (10GbE/25GbE)

> ℹ️ Additional ports can be activated by purchasing port licenses.

ISL ports

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

- NetApp CN1610: ports 0/13 through 0/16 (10GbE)
- BES-53248: ports 0/55-0/56 (100GbE)

The *NetApp Hardware Universe* contains information about ONTAP compatibility, supported EFOS firmware, and cabling to BES-53248 cluster switches.

ISL cabling

The appropriate ISL cabling is as follows:

- **Beginning:** For CN1610 to CN1610 (SFP+ to SFP+), four SFP+ optical fiber or copper direct-attach cables.
- **Final:** For BES-53248 to BES-53248 (QSFP28 to QSFP28), two QSFP28 optical transceivers/fiber or copper direct-attach cables.

Migrate the switches

Follow this procedure to migrate CN1610 cluster switches to BES-53248 cluster switches.

About the examples

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

- The examples use two nodes, each deploying two 10 GbE cluster interconnect ports: `e0a` and `e0b`.
- The command outputs might vary depending on different releases of ONTAP software.
- The CN1610 switches to be replaced are `CL1` and `CL2`.
- The BES-53248 switches to replace the CN1610 switches are `cs1` and `cs2`.
- The nodes are `node1` and `node2`.
- The switch CL2 is replaced by cs2 first, followed with CL1 by cs1.
- The BES-53248 switches are pre-loaded with the supported versions of Reference Configuration File (RCF) and Ethernet Fabric OS (EFOS) with ISL cables connected on ports 55 and 56.
- The cluster LIF names are `node1_clus1` and `node1_clus2` for node1, and `node2_clus1` and `node2_clus2` for node2.

About this task

This procedure covers the following scenario:

- The cluster starts with two nodes connected to two CN1610 cluster switches.
- CN1610 switch CL2 is replaced by BES-53248 switch cs2:
  - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
  - Disconnect the cables from all cluster ports on all nodes connected to CL2, and then use supported cables to reconnect the ports to the new cluster switch cs2.
- CN1610 switch CL1 is replaced by BES-53248 switch cs1:
  - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
  - Disconnect the cables from all cluster ports on all nodes connected to CL1, and then use supported cables to reconnect the ports to the new cluster switch cs1.

> (i) No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure migrates all of the cluster LIFs to the operational partner switch while performing the steps on the target switch.

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
   ```

   where 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.

   The following command suppresses automatic case creation for two hours:

   ```
   cluster1::*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Change the privilege level to advanced, entering **y** when prompted to continue:

   ```
   set -privilege advanced
   ```

   The advanced prompt (*>) appears.

Step 2: Configure ports and cabling

1. On the new switches, confirm that the ISL is cabled and healthy between switches cs1 and cs2:

   ```
   show port-channel
   ```

Show example

The following example shows that the ISL ports are **up** on switch cs1:

```
(cs1)# show port-channel 1/1
Local Interface................................ 1/1
Channel Name................................... Cluster-ISL
Link State..................................... Up
Admin Mode..................................... Enabled
Type........................................... Dynamic
Port channel Min-links......................... 1
Load Balance Option............................ 7
(Enhanced hashing mode)

Mbr    Device/        Port      Port
Ports  Timeout        Speed     Active
------ -------------  --------- -------
0/55   actor/long     100G Full  True
       partner/long
0/56   actor/long     100G Full  True
       partner/long
(cs1) #
```

The following example shows that the ISL ports are **up** on switch cs2:

```
(cs2)# show port-channel 1/1
Local Interface................................ 1/1
Channel Name................................... Cluster-ISL
Link State..................................... Up
Admin Mode..................................... Enabled
Type........................................... Dynamic
Port channel Min-links......................... 1
Load Balance Option............................ 7
(Enhanced hashing mode)

Mbr    Device/        Port      Port
Ports  Timeout        Speed     Active
------ -------------  --------- -------
0/55   actor/long     100G Full  True
       partner/long
0/56   actor/long     100G Full  True
       partner/long
```

2. Display the cluster ports on each node that is connected to the existing cluster switches:

```
network device-discovery show -protocol cdp
```

Show example

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

```
cluster1::*> network device-discovery show -protocol cdp
Node/       Local  Discovered
Protocol    Port   Device (LLDP: ChassisID)  Interface
Platform
----------- ------ ------------------------- ----------------
----------------
node2       /cdp
            e0a    CL1                        0/2
CN1610
            e0b    CL2                        0/2
CN1610
node1       /cdp
            e0a    CL1                        0/1
CN1610
            e0b    CL2                        0/1
CN1610
```

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

   a. Verify that all the cluster ports are `up` with a `healthy` status:

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

Show example

```
cluster1::*> 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
```

b. Verify that all the cluster interfaces (LIFs) are on their home ports:

```
network interface show -vserver Cluster
```

Show example

```
cluster1::*> network interface show -vserver Cluster

          Logical    Status     Network                 Current
Current Is
Vserver      Interface  Admin/Oper Address/Mask        Node
Port     Home
----------- ---------- ---------- ------------------
------------- ------- ----
Cluster
          node1_clus1  up/up     169.254.209.69/16  node1
e0a     true
          node1_clus2  up/up     169.254.49.125/16  node1
e0b     true
          node2_clus1  up/up     169.254.47.194/16  node2
e0a     true
          node2_clus2  up/up     169.254.19.183/16  node2
e0b     true
```

4. Verify that the cluster displays information for both cluster switches:

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command: `system switch ethernet show -is-monitoring -enabled-operational true`

```
cluster1::*> system switch ethernet show -is-monitoring-enabled
-operational true
Switch                            Type             Address        Model
----------------------------      ---------------  -------------  --------
CL1                               cluster-network  10.10.1.101    CN1610
      Serial Number: 01234567
       Is Monitored: true
             Reason:
   Software Version: 1.3.0.3
     Version Source: ISDP

CL2                               cluster-network  10.10.1.102    CN1610
      Serial Number: 01234568
       Is Monitored: true
             Reason:
   Software Version: 1.3.0.3
     Version Source: ISDP
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command: `system cluster-switch show -is-monitoring -enabled-operational true`

```
cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                           Type              Address        Model
---------------------------      ----------------  -------------  --------
CL1                              cluster-network   10.10.1.101    CN1610
      Serial Number: 01234567
       Is Monitored: true
             Reason:
   Software Version: 1.3.0.3
     Version Source: ISDP

CL2                              cluster-network   10.10.1.102    CN1610
      Serial Number: 01234568
       Is Monitored: true
             Reason:
   Software Version: 1.3.0.3
     Version Source: ISDP
cluster1::*>
```

5. Disable auto-revert on the cluster LIFs.

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto
-revert false
```

6. On cluster switch CL2, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs:

```
(CL2)# configure
(CL2)(Config)# interface 0/1-0/16
(CL2)(Interface 0/1-0/16)# shutdown
(CL2)(Interface 0/1-0/16)# exit
(CL2)(Config)# exit
(CL2)#
```

7. Verify that the cluster LIFs have failed over to the ports hosted on cluster switch CL1. This might take a few seconds.

```
network interface show -vserver Cluster
```

Show example

```
cluster1::*> network interface show -vserver Cluster
            Logical      Status      Network                 Current
Current Is
Vserver     Interface    Admin/Oper Address/Mask            Node
Port     Home
----------- ------------ ---------- ------------------ -------------
------- ----
Cluster
            node1_clus1  up/up       169.254.209.69/16  node1
e0a      true
            node1_clus2  up/up       169.254.49.125/16  node1
e0a      false
            node2_clus1  up/up       169.254.47.194/16  node2
e0a      true
            node2_clus2  up/up       169.254.19.183/16  node2
e0a      false
```

8. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
cluster1::*> cluster show
Node        Health  Eligibility    Epsilon
---------- ------- ------------- -------
node1       true    true           false
node2       true    true           false
```

9. Move all cluster node connection cables from the old CL2 switch to the new cs2 switch.

10. Confirm the health of the network connections moved to cs2:

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

Show example

```
cluster1::*> 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
```

All cluster ports that were moved should be `up`.

11. Check neighbor information on the cluster ports:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local   Discovered
Protocol       Port    Device (LLDP: ChassisID)   Interface
Platform
-----------    ------  ------------------------   ----------------
----------------
node2          /cdp
               e0a     CL1                        0/2
CN1610
               e0b     cs2                        0/2              BES-
53248
node1          /cdp
               e0a     CL1                        0/1
CN1610
               e0b     cs2                        0/1              BES-
53248
```

12. Confirm the switch port connections are healthy from switch cs2's perspective:

```
cs2# show port all
cs2# show isdp neighbors
```

13. On cluster switch CL1, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs:

```
(CL1)# configure
(CL1)(Config)# interface 0/1-0/16
(CL1)(Interface 0/1-0/16)# shutdown
(CL1)(Interface 0/13-0/16)# exit
(CL1)(Config)# exit
(CL1)#
```

All cluster LIFs failover to the cs2 switch.

14. Verify that the cluster LIFs have failed over to the ports hosted on switch cs2. This might take a few seconds:

```
network interface show -vserver Cluster
```

Show example

```
cluster1::*> network interface show -vserver Cluster
            Logical      Status     Network              Current
Current Is
Vserver     Interface    Admin/Oper Address/Mask         Node
Port     Home
----------- ------------ ---------- ------------------ -------------
------- ----
Cluster
            node1_clus1  up/up      169.254.209.69/16  node1
e0b     false
            node1_clus2  up/up      169.254.49.125/16  node1
e0b     true
            node2_clus1  up/up      169.254.47.194/16  node2
e0b     false
            node2_clus2  up/up      169.254.19.183/16  node2
e0b     true
```

15. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
cluster1::*> cluster show
Node        Health  Eligibility   Epsilon
----------  ------- ------------- -------
node1       true    true          false
node2       true    true          false
```

16. Move the cluster node connection cables from CL1 to the new cs1 switch.

17. Confirm the health of the network connections moved to cs1:

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

Show example

```
cluster1::*> 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
```

All cluster ports that were moved should be `up`.

18. Check neighbor information on the cluster ports:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/       Local  Discovered
Protocol    Port   Device (LLDP: ChassisID)  Interface
Platform
----------- ------ ------------------------- ----------------
----------------
node1       /cdp
            e0a    cs1                        0/1              BES-
53248
            e0b    cs2                        0/1              BES-
53248
node2       /cdp
            e0a    cs1                        0/2              BES-
53248
            e0b    cs2                        0/2              BES-
53248
```

19. Confirm the switch port connections are healthy from switch cs1's perspective:

```
cs1# show port all
cs1# show isdp neighbors
```

20. Verify that the ISL between cs1 and cs2 is still operational:

```
show port-channel
```

Show example

The following example shows that the ISL ports are **up** on switch cs1:

```
(cs1)# show port-channel 1/1
Local Interface................................ 1/1
Channel Name................................... Cluster-ISL
Link State..................................... Up
Admin Mode..................................... Enabled
Type........................................... Dynamic
Port channel Min-links......................... 1
Load Balance Option............................ 7
(Enhanced hashing mode)

Mbr     Device/        Port        Port
Ports   Timeout        Speed       Active
------  -------------  ---------   -------
0/55    actor/long     100G Full   True
        partner/long
0/56    actor/long     100G Full   True
        partner/long
(cs1) #
```

The following example shows that the ISL ports are **up** on switch cs2:

```
(cs2)# show port-channel 1/1
Local Interface................................ 1/1
Channel Name................................... Cluster-ISL
Link State..................................... Up
Admin Mode..................................... Enabled
Type........................................... Dynamic
Port channel Min-links......................... 1
Load Balance Option............................ 7
(Enhanced hashing mode)

Mbr     Device/        Port        Port
Ports   Timeout        Speed       Active
------  -------------  ---------   -------
0/55    actor/long     100G Full   True
        partner/long
0/56    actor/long     100G Full   True
        partner/long
```

21. Delete the replaced CN1610 switches from the cluster's switch table, if they are not automatically removed:

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command: `system switch ethernet delete -device`
device-name

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

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command: `system cluster-switch delete -device`
device-name

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

Step 3: Verify the configuration

1. Enable auto-revert on the cluster LIFs.

   ```
   cluster1::*> network interface modify -vserver Cluster -lif * -auto
   -revert true
   ```

2. Verify that the cluster LIFs have reverted to their home ports (this might take a minute):

   `network interface show -vserver Cluster`

   If the cluster LIFs have not reverted to their home port, manually revert them:

   `network interface revert -vserver Cluster -lif *`

3. Verify that the cluster is healthy:

   `cluster show`

4. Ping the remote cluster interfaces to verify connectivity:

   `cluster ping-cluster -node <name>`

Show example

```
cluster1::*> cluster ping-cluster -node node2
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69  node1      e0a
Cluster node1_clus2 169.254.49.125  node1      e0b
Cluster node2_clus1 169.254.47.194  node2      e0a
Cluster node2_clus2 169.254.19.183  node2      e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
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 4 path(s):
    Local 169.254.47.194 to Remote 169.254.209.69
    Local 169.254.47.194 to Remote 169.254.49.125
    Local 169.254.19.183 to Remote 169.254.209.69
    Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

5. To set up log collection, run the following command for each switch. You are prompted to enter the switch name, username, and password for log collection.

```
system switch ethernet log setup-password
```

```
cluster1::*> system switch ethernet log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
cs1
cs2

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

Enter the switch name: cs1
Would you like to specify a user other than admin for log
collection? {y|n}: n

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

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

Enter the switch name: cs2
Would you like to specify a user other than admin for log
collection? {y|n}: n

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

6. To start log collection, run the following command, replacing DEVICE with the switch used in the previous command. This starts both types of log collection: the detailed **Support** logs and an hourly collection of **Periodic** data.

```
system switch ethernet log modify -device <switch-name> -log-request true
```

```
cluster1::*> system switch ethernet log modify -device cs1 -log
-request true

Do you want to modify the cluster switch log collection
configuration?
{y|n}: [n] y

Enabling cluster switch log collection.

cluster1::*> system switch ethernet log modify -device cs2 -log
-request true

Do you want to modify the cluster switch log collection
configuration?
{y|n}: [n] y

Enabling cluster switch log collection.
```

Wait for 10 minutes and then check that the log collection completes:

```
system switch ethernet log show
```

> ⚠️ If any of these commands return an error or if the log collection does not complete, contact NetApp support.

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

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

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

Migrate to a switched NetApp cluster environment

If you have an existing two-node *switchless* cluster environment, you can migrate to a two-node *switched* cluster environment using Broadcom-supported BES-53248 cluster switches, which enables you to scale beyond two nodes in the cluster.

The migration process works for all cluster node ports using optical or Twinax ports, but it is not supported on this switch if nodes are using onboard 10GBASE-T RJ45 ports for the cluster network ports.

Review requirements

Review the following requirements for the cluster environment.

- Be aware that most systems require two dedicated cluster-network ports on each controller.

- Make sure that the BES-53248 cluster switch is set up as described in Replace requirements before starting this migration process.

- For the two-node switchless configuration, ensure that:

  - The two-node switchless configuration is properly set up and functioning.

  - The nodes are running ONTAP 9.5P8 and later. Support for 40/100 GbE cluster ports starts with EFOS firmware version 3.4.4.6 and later.

  - All cluster ports are in the **up** state.

  - All cluster logical interfaces (LIFs) are in the **up** state and on their home ports.

- For the Broadcom-supported BES-53248 cluster switch configuration, ensure that:

  - The BES-53248 cluster switch is fully functional on both switches.

  - Both switches have management network connectivity.

  - There is console access to the cluster switches.

  - BES-53248 node-to-node switch and switch-to-switch connections are using Twinax or fiber cables.

    The *NetApp Hardware Universe* contains information about ONTAP compatibility, supported EFOS firmware, and cabling to BES-53248 switches.

- Inter-Switch Link (ISL) cables are connected to ports 0/55 and 0/56 on both BES-53248 switches.

- Initial customization of both the BES-53248 switches is complete, so that:

  - BES-53248 switches are running the latest version of software.

  - BES-53248 switches have optional port licenses installed, if purchased.

  - Reference Configuration Files (RCFs) are applied to the switches.

- Any site customization (SMTP, SNMP, and SSH) are configured on the new switches.

Port group speed constraints
- The 48 10/25GbE (SFP28/SFP+) ports are combined into 12 x 4-port groups as follows: Ports 1-4, 5-8, 9-12, 13-16, 17-20, 21-24, 25-28, 29-32, 33-36, 37-40, 41-44, and 45-48.

- The SFP28/SFP+ port speed must be the same (10GbE or 25GbE) across all ports in the 4-port group.

- If speeds in a 4-port group are different, the switch ports will not operate correctly.

Migrate to the cluster environment

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

- The names of the BES-53248 switches are `cs1` and `cs2`.

- The names of the cluster SVMs are `node1` and `node2`.

- The names of the LIFs are `node1_clus1` and `node1_clus2` on node 1, and `node2_clus1` and `node2_clus2` on node 2 respectively.

- The `cluster1::*>` prompt indicates the name of the cluster.

- The cluster ports used in this procedure are `e0a` and `e0b`.

  The *NetApp Hardware Universe* contains the latest information about the actual cluster ports for your platforms.

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
   ```

   where 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.

   The following command suppresses automatic case creation for two hours:

   ```
   cluster1::*> system node autosupport invoke -node \* -type all -message
   MAINT=2h
   ```

2. Change the privilege level to advanced, entering **y** when prompted to continue:

   ```
   set -privilege advanced
   ```

   The advanced prompt (`*>`) appears.

Step 2: Configure ports and cabling

1. Disable all activated node-facing ports (not ISL ports) on both the new cluster switches cs1 **and** cs2.

   > ⓘ You must not disable the ISL ports.

   The following example shows that node-facing ports 1 through 16 are disabled on switch cs1:

   ```
   (cs1)# configure
   (cs1)(Config)# interface 0/1-0/16
   (cs1)(Interface 0/1-0/16)# shutdown
   (cs1)(Interface 0/1-0/16)# exit
   (cs1)(Config)# exit
   ```

2. Verify that the ISL and the physical ports on the ISL between the two BES-53248 switches cs1 and cs2 are up:

   ```
   show port-channel
   ```

Show example

The following example shows that the ISL ports are up on switch cs1:

```
(cs1)# show port-channel 1/1
Local Interface................................ 1/1
Channel Name................................... Cluster-ISL
Link State..................................... Up
Admin Mode..................................... Enabled
Type........................................... Dynamic
Port channel Min-links......................... 1
Load Balance Option............................ 7
(Enhanced hashing mode)

Mbr     Device/         Port        Port
Ports   Timeout         Speed       Active
------  -------------   ---------   -------
0/55    actor/long      100G Full   True
        partner/long
0/56    actor/long      100G Full   True
        partner/long
(cs1) #
```

The following example shows that the ISL ports are up on switch cs2:

```
(cs2)# show port-channel 1/1
Local Interface................................ 1/1
Channel Name................................... Cluster-ISL
Link State..................................... Up
Admin Mode..................................... Enabled
Type........................................... Dynamic
Port channel Min-links......................... 1
Load Balance Option............................ 7
(Enhanced hashing mode)

Mbr     Device/         Port        Port
Ports   Timeout         Speed       Active
------  -------------   ---------   -------
0/55    actor/long      100G Full   True
        partner/long
0/56    actor/long      100G Full   True
        partner/long
```

3. Display the list of neighboring devices:

```
show isdp neighbors
```

This command provides information about the devices that are connected to the system.

Show example

The following example lists the neighboring devices on switch cs1:

```
(cs1)# show isdp neighbors

Capability Codes: R - Router, T - Trans Bridge, B - Source Route
Bridge,
              S - Switch, H - Host, I - IGMP, r - Repeater
Device ID      Intf      Holdtime  Capability    Platform     Port ID
-------------- --------- --------- ------------- ------------ ---------
cs2            0/55      176       R             BES-53248    0/55
cs2            0/56      176       R             BES-53248    0/56
```

The following example lists the neighboring devices on switch cs2:

```
(cs2)# show isdp neighbors

Capability Codes: R - Router, T - Trans Bridge, B - Source Route
Bridge,
              S - Switch, H - Host, I - IGMP, r - Repeater
Device ID      Intf      Holdtime  Capability    Platform     Port ID
-------------- --------- --------- ------------- ------------ ---------
cs2            0/55      176       R             BES-53248    0/55
cs2            0/56      176       R             BES-53248    0/56
```

4. Verify that all cluster ports are up:

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

Show example

```
cluster1::*> network port show -ipspace Cluster

Node: node1

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

Node: node2

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

5. Verify that all cluster LIFs are up and operational:

```
network interface show -vserver Cluster
```

Show example

```
cluster1::*> network interface show -vserver Cluster

            Logical      Status     Network                 Current
Current Is
Vserver       Interface    Admin/Oper Address/Mask          Node
Port     Home
----------- ------------ ---------- ------------------ -------------
------- -----
Cluster
            node1_clus1  up/up      169.254.209.69/16  node1
e0a      true
            node1_clus2  up/up      169.254.49.125/16  node1
e0b      true
            node2_clus1  up/up      169.254.47.194/16  node2
e0a      true
            node2_clus2  up/up      169.254.19.183/16  node2
e0b      true
```

6. Disable auto-revert on the cluster LIFs.

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto
-revert false
```

7. Disconnect the cable from cluster port e0a on node1, and then connect e0a to port 1 on cluster switch cs1, using the appropriate cabling supported by the BES-53248 switches.

   The *NetApp Hardware Universe* contains more information about cabling.

8. Disconnect the cable from cluster port e0a on node2, and then connect e0a to port 2 on cluster switch cs1, using the appropriate cabling supported by the BES-53248 switches.

9. Enable all node-facing ports on cluster switch cs1.

   The following example shows that ports 1 through 16 are enabled on switch cs1:

```
(cs1)# configure
(cs1)(Config)# interface 0/1-0/16
(cs1)(Interface 0/1-0/16)# no shutdown
(cs1)(Interface 0/1-0/16)# exit
(cs1)(Config)# exit
```

10. Verify that all cluster ports are up:

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

Show example

```
cluster1::*> 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
```

11. Verify that all cluster LIFs are up and operational:

```
network interface show -vserver Cluster
```

Show example

```
cluster1::*> network interface show -vserver Cluster

         Logical      Status     Network                Current
Current Is
Vserver  Interface    Admin/Oper Address/Mask           Node         Port
Home
-------- ------------ ---------- ------------------ -----------
------- ----
Cluster
         node1_clus1  up/up      169.254.209.69/16  node1        e0a
false
         node1_clus2  up/up      169.254.49.125/16  node1        e0b
true
         node2_clus1  up/up      169.254.47.194/16  node2        e0a
false
         node2_clus2  up/up      169.254.19.183/16  node2        e0b
true
```

12. Display information about the status of the nodes in the cluster:

cluster show

Show example

The following example displays information about the health and eligibility of the nodes in the cluster:

```
cluster1::*> cluster show

Node                    Health  Eligibility   Epsilon
-------------------- ------- ------------   ------------
node1                   true    true          false
node2                   true    true          false
```

13. Disconnect the cable from cluster port e0b on node1, and then connect e0b to port 1 on cluster switch cs2, using the appropriate cabling supported by the BES-53248 switches.

14. Disconnect the cable from cluster port e0b on node2, and then connect e0b to port 2 on cluster switch cs2, using the appropriate cabling supported by the BES-53248 switches.

15. Enable all node-facing ports on cluster switch cs2.

The following example shows that ports 1 through 16 are enabled on switch cs2:

```
(cs2)# configure
(cs2)(Config)# interface 0/1-0/16
(cs2)(Interface 0/1-0/16)# no shutdown
(cs2)(Interface 0/1-0/16)# exit
(cs2)(Config)# exit
```

16. Verify that all cluster ports are up:

`network port show -ipspace Cluster`

Show example

```
cluster1::*> 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
```

Step 3: Verify the configuration

1. Enable auto-revert on the cluster LIFs.

   ```
   cluster1::*> network interface modify -vserver Cluster -lif * -auto
   -revert true
   ```

2. Verify that the cluster LIFs have reverted to their home ports (this might take a minute):

   `network interface show -vserver Cluster`

   If the cluster LIFs have not reverted to their home port, manually revert them:

   `network interface revert -vserver Cluster -lif *`

3. Verify that all interfaces display `true` for `Is Home`:

   `network interface show -vserver Cluster`

   > **i**   This might take several minutes to complete.

   Show example

   ```
   cluster1::*> network interface show -vserver Cluster

             Logical      Status     Network                Current
   Current Is
   Vserver   Interface    Admin/Oper Address/Mask           Node          Port
   Home
   --------- ------------ ---------- ------------------- ----------
   ------- ----
   Cluster
             node1_clus1  up/up      169.254.209.69/16   node1         e0a
   true
             node1_clus2  up/up      169.254.49.125/16   node1         e0b
   true
             node2_clus1  up/up      169.254.47.194/16   node2         e0a
   true
             node2_clus2  up/up      169.254.19.183/16   node2         e0b
   true
   ```

4. Verify that both nodes each have one connection to each switch:

   `show isdp neighbors`

Show example

The following example shows the appropriate results for both switches:

```
(cs1)# show isdp neighbors

Capability Codes: R - Router, T - Trans Bridge, B - Source Route
Bridge,
                S - Switch, H - Host, I - IGMP, r - Repeater
Device ID      Intf          Holdtime  Capability   Platform -- Port
ID
-------------- ------------ --------- ------------ -----------
----------
node1          0/1            175       H            FAS2750      e0a
node2          0/2            157       H            FAS2750      e0a
cs2            0/55           178       R            BES-53248    0/55
cs2            0/56           178       R            BES-53248    0/56


(cs2)# show isdp neighbors

Capability Codes: R - Router, T - Trans Bridge, B - Source Route
Bridge,
                S - Switch, H - Host, I - IGMP, r - Repeater
Device ID      Intf          Holdtime  Capability   Platform     Port
ID
-------------- ------------ --------- ------------ -----------
------------
node1          0/1            137       H            FAS2750      e0b
node2          0/2            179       H            FAS2750      e0b
cs1            0/55           175       R            BES-53248    0/55
cs1            0/56           175       R            BES-53248    0/56
```

5. Display information about the discovered network devices in your cluster:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/       Local  Discovered
Protocol    Port   Device (LLDP: ChassisID)  Interface
Platform
----------- ------ ------------------------- ----------------
----------------
node2       /cdp
            e0a    cs1                        0/2                BES-
53248
            e0b    cs2                        0/2                BES-
53248
node1       /cdp
            e0a    cs1                        0/1                BES-
53248
            e0b    cs2                        0/1                BES-
53248
```

6. Verify that the settings are disabled:

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

(i) It might take several minutes for the command to complete. Wait for the '3 minute lifetime to expire' announcement.

The `false` output in the following example shows that the configuration settings are disabled:

```
cluster1::*> network options switchless-cluster show
Enable Switchless Cluster: false
```

7. Verify the status of the node members in the cluster:

```
cluster show
```

Show example

The following example shows information about the health and eligibility of the nodes in the cluster:

```
cluster1::*> cluster show

Node                    Health  Eligibility   Epsilon
--------------------    ------- ------------  --------
node1                   true    true          false
node2                   true    true          false
```

8. Verify that the cluster network has full connectivity using the command:

`cluster ping-cluster -node node-name`

Show example

```
cluster1::*> cluster ping-cluster -node local

Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 192.168.168.26 node1 e0a
Cluster node1_clus2 192.168.168.27 node1 e0b
Cluster node2_clus1 192.168.168.28 node2 e0a
Cluster node2_clus2 192.168.168.29 node2 e0b
Local = 192.168.168.28 192.168.168.29
Remote = 192.168.168.26 192.168.168.27
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
...............
Detected 1500 byte MTU on 4 path(s):
    Local 192.168.168.28 to Remote 192.168.168.26
    Local 192.168.168.28 to Remote 192.168.168.27
    Local 192.168.168.29 to Remote 192.168.168.26
    Local 192.168.168.29 to Remote 192.168.168.27
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

9. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

Show example

```
cluster1::*> system node autosupport invoke -node \* -type all
-message MAINT=END
```

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

What's next?

After your migration completes, you might need to install the required configuration file to support the Ethernet Switch Health Monitor (CSHM) for BES-53248 cluster switches. See Enable log collection.