Migrating to a two-node switched cluster with NetApp CN1610 cluster switches

If you have a two-node switchless cluster, you can migrate, non-disruptively, to a two-node switched cluster that includes NetApp CN1610 cluster-network switches. The procedure you use depends on whether you have two dedicated cluster-network ports on each controller or a single cluster port on each controller.

About this task

Most systems require two dedicated cluster-network ports on each controller.

FAS22xx nodes allow a single cluster port on each controller.

There are two migration options available:

- *Migrating from a switchless cluster to a switched NetApp CN1610 cluster environment* on page 1
- *Migrating from a switchless cluster (FAS22xx systems with a single cluster-network connection)* on page 12

Migrating from a switchless cluster environment to a switched NetApp CN1610 cluster environment

If you have an existing two-node switchless cluster environment, you can migrate to a two-node switched cluster environment using CN1610 cluster network switches that enables you to scale beyond two nodes.

Before you begin

Two-node switchless configuration:

- The two-node switchless configuration must be properly set up and functioning.
- The nodes must be running ONTAP 8.2 or later.
- All cluster ports must be in the **up** state.
- All cluster logical interfaces (LIFs) must be in the **up** state and on their home ports.

CN1610 cluster switch configuration:

- The CN1610 cluster switch infrastructure must be and fully functional on both switches.
- Both switches must have management network connectivity.
- There must be console access to the cluster switches.
- CN1610 node-to-node switch and switch-to-switch connections must use twinax or fiber cables.
  The *NetApp Hardware Universe* contains more information about cabling.
  Hardware Universe - Switches
- Inter-Switch Link (ISL) cables must be connected to ports 13 through 16 on both CN1610 switches.
- Initial customization of both the CN1610 switches must be completed.

Any previous site customization, such as SMTP, SNMP, and SSH should be copied to the new switches.

About this task

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

- The names of the CN1610 switches are cs1 and cs2.
- The names of the LIFs are clus1 and clus2.
- The names of the SVMs are node1 and node2.
- The `cluster::*>` prompt indicates the name of the cluster.
- The cluster ports used in this procedure are e1a and e2a.
  The *Hardware Universe* contains the latest information about the actual cluster ports for your platforms.

Steps

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

   The advanced prompt (*>) appears.

2. 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=`***x*`h`**

   x is the duration of the maintenance window in hours.

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

   #### Example

   The following command suppresses automatic case creation for two hours:

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

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

   You must not disable the ISL ports.

   #### Example

   The following example shows that node-facing ports 1 through 12 are disabled on switch cs1:

   ```
   (cs1)>
   enable

   (cs1)# configure
   (cs1)(Config)# interface 0/1-0/12
   (cs1)(Interface 0/1-0/12)# shutdown
   (cs1)(Interface 0/1-0/12)# exit
   (cs1)(Config)# exit
   ```

   The following example shows that node-facing ports 1 through 12 are disabled on switch cs2:

   ```
   (c2)>
   enable
   ```

```
(cs2)# configure
(cs2)(Config)# interface 0/1-0/12
(cs2)(Interface 0/1-0/12)# shutdown
(cs2)(Interface 0/1-0/12)# exit
(cs2)(Config)# exit
```

4. Verify that the ISL and the physical ports on the ISL between the two CN1610 cluster switches cs1 and cs2 are **up**: **show port-channel**

Example

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

```
(cs1)# show port-channel 3/1
Local Interface................................ 3/1
Channel Name................................... ISL-LAG
Link State..................................... Up
Admin Mode..................................... Enabled
Type........................................... Static
Load Balance Option............................ 7
(Enhanced hashing mode)

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

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

```
(cs2)# show port-channel 3/1
Local Interface................................ 3/1
Channel Name................................... ISL-LAG
Link State..................................... Up
Admin Mode..................................... Enabled
Type........................................... Static
Load Balance Option............................ 7
(Enhanced hashing mode)

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

5. Display the list of neighboring devices: **show isdp neighbors**

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

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                   0/13         11        S            CN1610    0/13
cs2                   0/14         11        S            CN1610    0/14
cs2                   0/15         11        S            CN1610    0/15
cs2                   0/16         11        S            CN1610    0/16
```

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
--------------------- ------------ --------- ------------ --------- ------------
cs1                   0/13         11        S            CN1610    0/13
cs1                   0/14         11        S            CN1610    0/14
cs1                   0/15         11        S            CN1610    0/15
cs1                   0/16         11        S            CN1610    0/16
```

6. Verify that each cluster port is connected to the corresponding port on its partner cluster node: **`run * cdpd show-neighbors`**

Example

The following example shows that cluster ports e1a and e2a are connected to the same port on their cluster partner node:

```
cluster::*> run * cdpd show-neighbors
2 entries were acted on.

Node: node1
Local  Remote          Remote                 Remote           Hold  Remote
Port   Device          Interface              Platform         Time  Capability
------ --------------- ---------------------- ---------------- ----- ----------
e1a    node2           e1a                    FAS3270           137   H
e2a    node2           e2a                    FAS3270           137   H


Node: node2

Local  Remote          Remote                 Remote           Hold  Remote
Port   Device          Interface              Platform         Time  Capability
------ --------------- ---------------------- ---------------- ----- ----------
e1a    node1           e1a                    FAS3270           161   H
e2a    node1           e2a                    FAS3270           161   H
```

7. Verify that all of the cluster LIFs are **`up`** and operational: **`network interface show -role cluster`**

   Each cluster LIF should display `true` in the “Is Home” column.

Example

```
cluster::*> network interface show -role cluster
            Logical    Status     Network       Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask  Node          Port    Home
----------- ---------- ---------- ------------- ------------- ------- ----
node1
            clus1      up/up      10.10.10.1/16 node1         e1a     true
            clus2      up/up      10.10.10.2/16 node1         e2a     true
node2
            clus1      up/up      10.10.11.1/16 node2         e1a     true
            clus2      up/up      10.10.11.2/16 node2         e2a     true

4 entries were displayed.
```

Attention: The following modification and migration commands in steps *9* on page 5 through *12* on page 6 must be done from the local node.

8. Verify that all cluster ports are **up**: **`network port show -role cluster`**

Example

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

                                       Auto-Negot  Duplex     Speed (Mbps)
Node   Port   Role         Link  MTU   Admin/Oper  Admin/Oper Admin/Oper
------ ------ ------------ ----- ----- ----------- ---------- ------------
node1
       e1a    clus1        up    9000  true/true  full/full   auto/10000
       e2a    clus2        up    9000  true/true  full/full   auto/10000
node2
       e1a    clus1        up    9000  true/true  full/full   auto/10000
       e2a    clus2        up    9000  true/true  full/full   auto/10000

4 entries were displayed.
```

9. Set the `-auto-revert` parameter to **`false`** on cluster LIFs clus1 and clus2 on both nodes: **`network interface modify`**

Example

```
cluster::*> network interface modify -vserver node1 -lif clus1 -auto-revert false
cluster::*> network interface modify -vserver node1 -lif clus2 -auto-revert false
cluster::*> network interface modify -vserver node2 -lif clus1 -auto-revert false
cluster::*> network interface modify -vserver node2 -lif clus2 -auto-revert false
```

10. Ping the cluster ports to verify the cluster connectivity: **`cluster ping-cluster local`**

The command output shows connectivity between all of the cluster ports.

11. Migrate clus1 to port e2a on the console of each node: **`network interface migrate`**

Example

The following example shows the process for migrating clus1 to port e2a on node1 and node2:

```
cluster::*> network interface migrate -vserver node1 -lif clus1 -source-node node1 -dest-
node node1 -dest-port e2a
cluster::*> network interface migrate -vserver node2 -lif clus1 -source-node node2 -dest-
node node2 -dest-port e2a
```

12. Verify that the migration took place: **network interface show**

    Example

    The following example verifies that clus1 is migrated to port e2a on node1 and node2:

```
cluster::*> network interface show -role cluster
            Logical    Status     Network       Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask  Node          Port    Home
----------- ---------- ---------- ------------- ------------- ------- ----
node1
            clus1      up/up    10.10.10.1/16   node1         e2a     false
            clus2      up/up    10.10.10.2/16   node1         e2a     true
node2
            clus1      up/up    10.10.11.1/16   node2         e2a     false
            clus2      up/up    10.10.11.2/16   node2         e2a     true

4 entries were displayed.
```

13. Shut down cluster port e1a on both nodes: **network port modify**

    Example

    The following example shows how to shut down the port e1a on node1 and node2:

```
cluster::*> network port modify -node node1 -port e1a -up-admin false
cluster::*> network port modify -node node2 -port e1a -up-admin false
```

14. Verify the port status: **network port show**

    Example

    The following example shows that port e1a is **down** on node1 and node2:

```
cluster::*> network port show -role cluster
                                  Auto-Negot  Duplex     Speed (Mbps)
Node   Port   Role         Link   MTU Admin/Oper  Admin/Oper Admin/Oper
------ ------ ------------ ---- ----- ----------- ---------- ------------
node1
       e1a    clus1        down  9000  true/true  full/full   auto/10000
       e2a    clus2        up    9000  true/true  full/full   auto/10000
node2
       e1a    clus1        down  9000  true/true  full/full   auto/10000
       e2a    clus2        up    9000  true/true  full/full   auto/10000

4 entries were displayed.
```

15. Disconnect the cable from cluster port e1a on node1, and then connect e1a to port 1 on cluster switch cs1, using the appropriate cabling supported by the CN1610 switches.

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

    Hardware Universe - Switches

16. Disconnect the cable from cluster port e1a on node2, and then connect e1a to port 2 on cluster switch cs1, using the appropriate cabling supported by the CN1610 switches.

17. Enable all of the node-facing ports on cluster switch cs1.

Example

The following example shows that ports 1 through 12 are enabled on switch cs1:

```
(cs1)# configure
(cs1)(Config)# interface 0/1-0/12
(cs1)(Interface 0/1-0/12)# no shutdown
(cs1)(Interface 0/1-0/12)# exit
(cs1)(Config)# exit
```

18. Enable the first cluster port e1a on each node: **`network port modify`**

Example

The following example shows how to enable the port e1a on node1 and node2:

```
cluster::*> network port modify -node node1 -port e1a -up-admin true
cluster::*> network port modify -node node2 -port e1a -up-admin true
```

19. Verify that all of the cluster ports are **`up`**: **`network port show`**

Example

The following example shows that all of the cluster ports are **`up`** on node1 and node2:

```
cluster::*> network port show -role cluster
                                  Auto-Negot  Duplex     Speed (Mbps)
Node   Port   Role         Link   MTU Admin/Oper  Admin/Oper Admin/Oper
------ ------ ------------ ---- ----- ----------- ---------- ------------
node1
       e1a    clus1        up    9000  true/true  full/full   auto/10000
       e2a    clus2        up    9000  true/true  full/full   auto/10000
node2
       e1a    clus1        up    9000  true/true  full/full   auto/10000
       e2a    clus2        up    9000  true/true  full/full   auto/10000

4 entries were displayed.
```

20. Revert clus1 (which was previously migrated) to e1a on both nodes: **`network interface revert`**

Example

The following example shows how to revert clus1 to the port e1a on node1 and node2:

```
cluster::*> network interface revert -vserver node1 -lif clus1
cluster::*> network interface revert -vserver node2 -lif clus1
```

21. Verify that all of the cluster LIFs are **`up`**, operational, and display as `true` in the “Is Home” column: **`network interface show`**

Example

The following example shows that all of the LIFs are **`up`** on node1 and node2 and that the “Is Home” column results are **`true`**:

```
cluster::*> network interface show -role cluster
            Logical    Status     Network       Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask  Node          Port    Home
----------- ---------- ---------- ------------- ------------- ------- ----
node1
            clus1      up/up    10.10.10.1/16   node1         e1a     true
            clus2      up/up    10.10.10.2/16   node1         e2a     true
node2
            clus1      up/up    10.10.11.1/16   node2         e1a     true
            clus2      up/up    10.10.11.2/16   node2         e2a     true

4 entries were displayed.
```

22. Display information about the status of the nodes in the cluster: **`cluster show`**

Example

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

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

23. Migrate clus2 to port e1a on the console of each node: **`network interface migrate`**

Example

The following example shows the process for migrating clus2 to port e1a on node1 and node2:

```
cluster::*> network interface migrate -vserver node1 -lif clus2 -source-node node1 -dest-
node node1 -dest-port e1a
cluster::*> network interface migrate -vserver node2 -lif clus2 -source-node node2 -dest-
node node2 -dest-port e1a
```

24. Verify that the migration took place: **`network interface show`**

Example

The following example verifies that clus2 is migrated to port e1a on node1 and node2:

```
cluster::*> network interface show -role cluster
            Logical    Status     Network       Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask  Node          Port    Home
----------- ---------- ---------- ------------- ------------- ------- ----
node1
            clus1      up/up    10.10.10.1/16   node1         e1a     true
            clus2      up/up    10.10.10.2/16   node1         e1a     false
node2
            clus1      up/up    10.10.11.1/16   node2         e1a     true
            clus2      up/up    10.10.11.2/16   node2         e1a     false

4 entries were displayed.
```

25. Shut down cluster port e2a on both nodes: **`network port modify`**

Example

The following example shows how to shut down the port e2a on node1 and node2:

```
cluster::*> network port modify -node node1 -port e2a -up-admin false
cluster::*> network port modify -node node2 -port e2a -up-admin false
```

26. Verify the port status: **network port show**

Example

The following example shows that port e2a is **down** on node1 and node2:

```
cluster::*> network port show -role cluster
                                        Auto-Negot  Duplex     Speed (Mbps)
Node   Port   Role         Link   MTU Admin/Oper  Admin/Oper Admin/Oper
------ ------ ------------ ---- ----- ----------- ---------- ------------
node1
       e1a    clus1        up    9000  true/true  full/full   auto/10000
       e2a    clus2        down  9000  true/true  full/full   auto/10000
node2
       e1a    clus1        up    9000  true/true  full/full   auto/10000
       e2a    clus2        down  9000  true/true  full/full   auto/10000

4 entries were displayed.
```

27. Disconnect the cable from cluster port e2a on node1, and then connect e2a to port 1 on cluster switch cs2, using the appropriate cabling supported by the CN1610 switches.

28. Disconnect the cable from cluster port e2a on node2, and then connect e2a to port 2 on cluster switch cs2, using the appropriate cabling supported by the CN1610 switches.

29. Enable all of the node-facing ports on cluster switch cs2.

Example

The following example shows that ports 1 through 12 are enabled on switch cs2:

```
(cs2)# configure
(cs2)(Config)# interface 0/1-0/12
(cs2)(Interface 0/1-0/12)# no shutdown
(cs2)(Interface 0/1-0/12)# exit
(cs2)(Config)# exit
```

30. Enable the second cluster port e2a on each node:

Example

The following example shows how to enable the port e2a on node1 and node2:

```
cluster::*> network port modify -node node1 -port e2a -up-admin true
cluster::*> network port modify -node node2 -port e2a -up-admin true
```

31. Verify that all of the cluster ports are **up**: **network port show -role cluster**

Example

The following example shows that all of the cluster ports are **up** on node1 and node2:

```
cluster::*> network port show -role cluster
                                     Auto-Negot  Duplex     Speed (Mbps)
Node   Port   Role         Link   MTU Admin/Oper  Admin/Oper Admin/Oper
------ ------ ------------ ---- ----- ----------- ---------- ------------
node1
       e1a    clus1        up    9000  true/true  full/full   auto/10000
       e2a    clus2        up    9000  true/true  full/full   auto/10000
node2
       e1a    clus1        up    9000  true/true  full/full   auto/10000
       e2a    clus2        up    9000  true/true  full/full   auto/10000

4 entries were displayed.
```

32. Revert clus2 (which was previously migrated) to e2a on both nodes: **`network interface revert`**

Example

The following example shows how to revert clus2 to the port e2a on node1 and node2:

```
cluster::*> network interface revert -vserver node1 -lif clus2
cluster::*> network interface revert -vserver node2 -lif clus2
```

33. Verify that all of the interfaces display **`true`** in the "Is Home" column: **`network interface show`**.

Example

The following example shows that all of the LIFs are **`up`** on node1 and node2 and that the "Is Home" column results are **`true`**:

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

             Logical    Status     Network            Current     Current Is
Vserver      Interface  Admin/Oper Address/Mask       Node        Port    Home
-----------  ---------- ---------- ------------------ ----------- ------- ----
node1
             clus1      up/up      10.10.10.1/16      node1       e1a     true
             clus2      up/up      10.10.10.2/16      node1       e2a     true
node2
             clus1      up/up      10.10.11.1/16      node2       e1a     true
             clus2      up/up      10.10.11.2/16      node2       e2a     true
```

34. Ping the cluster ports to verify the cluster connectivity: **`cluster ping-cluster local`**

The command output shows connectivity between all of the cluster ports.

35. Verify that both nodes have two connections to each switch: **`show isdp neighbors`**

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          132       H            FAS3270   e1a
node2                 0/2          163       H            FAS3270   e1a
cs2                   0/13         11        S            CN1610    0/13
```

```
cs2                    0/14         11        S            CN1610    0/14
cs2                    0/15         11        S            CN1610    0/15
cs2                    0/16         11        S            CN1610    0/16

(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          132       H            FAS3270   e2a
node2                  0/2          163       H            FAS3270   e2a
cs1                    0/13         11        S            CN1610    0/13
cs1                    0/14         11        S            CN1610    0/14
cs1                    0/15         11        S            CN1610    0/15
cs1                    0/16         11        S            CN1610    0/16
```

36. Display information about the devices in your configuration: **network device discovery show**

37. Disable the two-node switchless configuration settings on both nodes: **network options switchless-cluster modify**

Example

The following example shows how to disable the switchless configuration settings:

```
cluster::*> network options switchless-cluster modify -enabled false
```

38. Verify that the settings are disabled: **network options switchless-cluster show**

Example

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

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

39. Configure clusters clus1 and clus2 to auto revert on each node and confirm:

Example

```
cluster::*> network interface modify -vserver node1 -lif clus1 -auto-revert true
cluster::*> network interface modify -vserver node1 -lif clus2 -auto-revert true
cluster::*> network interface modify -vserver node2 -lif clus1 -auto-revert true
cluster::*> network interface modify -vserver node2 -lif clus2 -auto-revert true
```

40. Verify the status of the node members in the cluster: **cluster show**

Example

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

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

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

    Example

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

42. Change the privilege level back to admin: **`set -privilege admin`**

Related information

Hardware Universe
NetApp CN1601 and CN1610 description page
CN1601 and CN1610 Switch Setup and Configuration Guide
NetApp KB Article 1010449: How to suppress automatic case creation during scheduled maintenance windows

Migrating to a two-node switched cluster in FAS22xx systems with a single cluster-network connection

If you have FAS22xx systems in an existing two-node switchless cluster in which each controller module has a single, back-to-back 10 GbE connection for cluster connectivity, you can use the switchless cluster networking option and replace the direct back-to-back connectivity with switch connections.

Before you begin

- Two cluster connections are required to migrate from a switchless configuration to a switched configuration.
- The cluster must be healthy and consist of two nodes connected with back-to-back connectivity.
- The nodes must be running ONTAP 8.2 or later.
- The switchless cluster feature cannot be used with more than two nodes.
- All cluster ports must be in the **`up`** state.

About this task

This procedure is a nondisruptive procedure that removes the direct cluster connectivity in a switchless environment and replaces each connection to the switch with a connection to the partner node.

Steps

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

   The advanced prompt (`*>`) appears.

2. Check the cluster status of the nodes at the system console of either node: **`cluster show`**

   Example

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

   ```
   cluster::*> cluster show
   Node                 Health  Eligibility   Epsilon
   -------------------- ------- ------------  ------------
   ```

```
node1                true    true          false
node2                true    true          false

2 entries were displayed.
```

3. Check the status of the HA pair at the system console of either node: **storage failover show**

Example

The following example shows the status of node1 and node2:

```
Node           Partner        Possible State Description
-------------- -------------- -------- -------------------------------------
node1          node2          true      Connected to node2
node2          node1          true      Connected to node1

2 entries were displayed.
```

4. 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=*x*h**

   x is the duration of the maintenance window in hours.

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

Example

The following command suppresses automatic case creation for two hours:

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

5. Verify that the current state of the switchless cluster is **true**, and then disable the switchless cluster mode: **network options switchless-cluster modify -enabled false**

6. Take over the target node: **storage failover takeover -ofnode *target_node_name***

   It does not matter which node is the target node. When it is taken over, the target node automatically reboots and displays the `Waiting for giveback...` message.

   The active node is now serving data for the partner (target) node that was taken over.

7. Wait for two minutes after takeover of the impaired node to confirm that the takeover was completed successfully.

8. With the target node showing the `Waiting for giveback...` message, shut it down.

   The method you use to shut down the node depends on whether you use remote management through the node Service Processor (SP).

If SP	Then...
Is configured	Log in to the impaired node SP, and then power off the system: **system power off**
Is not configured	At the impaired node prompt, press **Ctrl-C**, and then respond **y** to halt the node.

9. On each controller module, disconnect the cable that connects the 10 GbE cluster port to the switchless cluster.

10. Connect the 10 GbE cluster port to the switch on both controller modules.

11. Verify that the 10 GbE cluster ports connected on the switch are configured to be part of the same VLAN.

If you plan to connect the cluster ports on each controller module to different switches, then you must verify that the ports on which the cluster ports are connected on each switch are configured for the same VLAN and that trunking is properly configured on both switches.

12. Give back storage to the target node: **`storage failover giveback -ofnode node2`**

13. Monitor the progress of the giveback operation: **`storage failover show-giveback`**

14. After the giveback operation is complete, confirm that the HA pair is healthy and takeover is possible: **`storage failover show`**

    Example

    The output should be similar to the following:

```
Node           Partner        Possible State Description
-------------- -------------- -------- -------------------------------------
node1          node2          true      Connected to node2
node2          node1          true      Connected to node1

2 entries were displayed.
```

15. Verify that the cluster port LIFs are operating correctly: **`network interface show -role cluster`**

    Example

    The following example shows that the LIFs are **`up`** on node1 and node2 and that the “Is Home” column results are **`true`**:

```
cluster::*> network interface show -role cluster
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
node1
            clus1        up/up    192.168.177.121/24  node1        e1a     true
node2
            clus1        up/up    192.168.177.123/24  node2        e1a     true

2 entries were displayed.
```

16. Check the cluster status of the nodes at the system console of either node: **`cluster show`**

    Example

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

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

2 entries were displayed.
```

17. Ping the cluster ports to verify the cluster connectivity: **`cluster ping-cluster local`**

    The command output should show connectivity between all of the cluster ports.

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

Example

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

19. Change the privilege level back to admin: **`set -privilege admin`**

Related information

NetApp KB Article 1010449: How to suppress automatic case creation during scheduled maintenance windows

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

doccomments@netapp.com

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Trademark

http://www.netapp.com/us/legal/netapptmlist.aspx