



---

# Transitioning to a two-node switchless cluster

If you have an existing two-node cluster that uses cluster interconnect switches and is running ONTAP 9.1 or later, you can replace the switches with direct, back-to-back connections between the nodes.

## About this task

The switchless cluster interconnect feature cannot be used with more than two nodes.

Transitioning to a two-node switchless cluster configuration is a nondisruptive operation. The procedure that you use depends on whether you have two dedicated cluster interconnect ports on each node (as required on most systems), four dedicated cluster interconnect ports on each node, or a single cluster port on each node (a supported option on FAS22xx nodes).

This document describes transitioning a cluster with a switched cluster network to one where two nodes are directly connected for ONTAP versions 9.1 and later.

- [Transitioning to a two-node switchless cluster on systems with two cluster interconnect connections](#)
- [Transitioning to a two-node switchless cluster on systems with one cluster interconnect connections](#)
- [Transitioning to a two-node switchless cluster on systems with four cluster interconnect connections](#)

## Choices

- [Transitioning to a two-node switchless cluster on systems with two cluster interconnect connections](#) on page 1
- [Transitioning to a two-node switchless cluster running ONTAP 9.1 or later \(FAS22xx systems with a single cluster interconnect connection\)](#) on page 10
- [Transitioning to a two-node switchless cluster \(systems with four redundant cluster interconnect connections\)](#) on page 14

## Transitioning to a two-node switchless cluster on systems with two cluster interconnect connections


If you have a two-node cluster that uses cluster interconnect switches, replace the switches with direct, back-to-back connections between the nodes. This is a nondisruptive operation.

### Before you begin

- The cluster must be healthy and consist of two nodes connected by cluster switches, and the nodes must be running the same ONTAP release.
  - Each node must have two dedicated cluster ports that provide redundant cluster interconnect connections.

### 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.



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 platform.

### Steps

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

```
set -privilege advanced
.
```

The advanced prompt `*>` appears.

2. ONTAP releases 9.2 and later support automatic detection of switchless clusters, which is enabled by default.

### Example

If running 9.2 or later, verify that detection of switchless clusters is enabled by running the advanced privilege command:

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

The following example 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 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.

### Example

The following command suppresses automatic case creation for two hours:

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

4. Identify the cluster ports and verify link status and health: `network port show -ipSpace Cluster`.

The following example shows the nodes that are using e0a and e0b as cluster ports. Your nodes might be using different cluster ports because they vary by platform.

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

## Example

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

Node: node1
Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health
Admin/Oper Status
-----
e0a      Cluster      Cluster          up  9000     auto/10000 healthy
e0b      Cluster      Cluster          up  9000     auto/10000 healthy

Node: node2
Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health
Admin/Oper Status
-----
e0a      Cluster      Cluster          up  9000     auto/10000 healthy
e0b      Cluster      Cluster          up  9000     auto/10000 healthy
4 entries were displayed.
```

5. Make sure that all of 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
```

## 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, use the `network interface revert -vserver Cluster -lif <lif name>` to move those LIFs to their home ports.

6. Verify that all cluster LIFs have auto-revert enabled:

```
network interface show -vserver - fields auto-revert
```

## Example

```
Cluster -fields auto-revert
cluster::> network interface show -vserver Cluster -fields auto-revert
vserver lif      auto-revert
-----
Cluster node1_clus1 true
Cluster node1_clus2 true
Cluster node2_clus1 true
Cluster node2_clus2 true
4 entries were displayed.
```

7. Enable auto-revert on any cluster LIF where auto-revert is disabled:

```
network interface modify -vserver Cluster -lif <cluster-lif-name> -auto-revert true
```

### Example

```
cluster::> network interface modify -vserver Cluster -lif node1_clus1 -auto-revert true
```

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

```
network device-discovery show -node <node> -port <cluster port> -protocol cdp
```

### Example

The “Discovered Device” column should be the name of the cluster switch that the port is connected to. The following example shows that cluster ports e0a and e0b are correctly connected to cluster switches cs1 and cs2.

```
cluster::> network device-discovery show -node node1 -port e0a -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface      Platform
-----
Node1/cdp  e0a    cs1                       Ethernet1/25   N1K-C9999UP

cluster::> network device-discovery show -node node1 -port e0b -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface      Platform
-----
Node1/cdp  e0b    cs2                       Ethernet1/25   N1K-C9999UP
```

9. For ONTAP release 9.1, enable the switchless-cluster network option:

```
network option switchless-cluster modify -enabled true
```

### Example

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

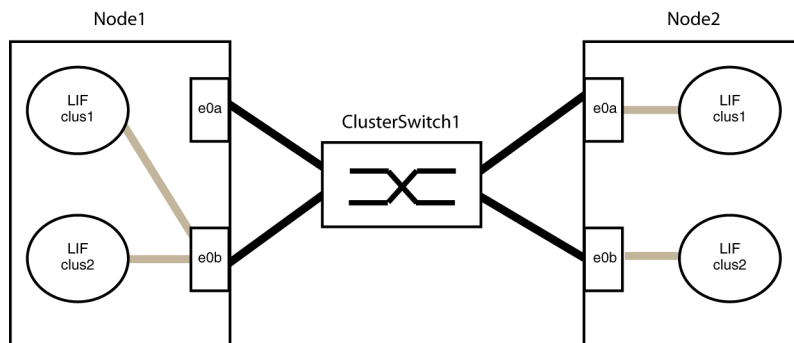
10. Migrate the LIFs on the first cluster port of node1 to the second cluster port on node1:

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

### Example

```
network interface migrate -vserver Cluster -lif node1_clus1 -destination-node node1 -destination-port e0b
```

The LIF node1\_clus1 is now associated with port e0b on node1, as indicated by the gray line from LIF node1\_clus1 to port e0b in the following illustration:



11. Migrate the LIFs on the first cluster port of node2 to the second cluster port on node2:

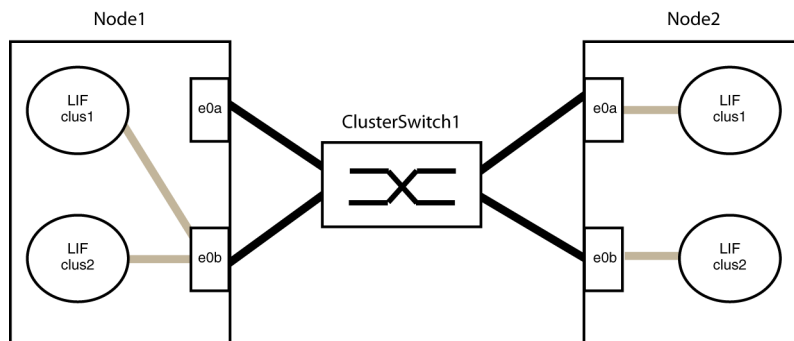
```
network interface migrate -vserver Cluster -lif lif_name -destination-node destination-node-name -destination-port port-name
```

### Example

```
network interface migrate -vserver Cluster -lif node2_clus1 -destination-node node2 -destination-port e0b
```

The LIF node2\_clus1 is now associated with port e0b on node2, as indicated by the gray line from LIF node2\_clus1 to port e0b in the following illustration:

- 12.



Confirm the LIFs successfully migrated:

```
network interface show -vserver Cluster -fields curr-port,is-home
```

### Example


The LIF migration is successful if the LIFs node1\_clus1 and node2\_clus1 have e0b as the value for “Current Port” and false for the “Is Home” value.

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

13. Disconnect the cables that link the switches from the cluster ports that no longer host cluster LIFs.

### Example


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:



14. Cable the ports that were formerly connected to the switch directly to each other, creating a direct connection between the nodes.

### Example

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



15. For ONTAP releases 9.2 and later, 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
```

### Example

The following example shows that the switchless cluster is enabled:

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

16. After the ports have been cabled and are able to process network traffic, the cluster LIFs automatically revert back to their home ports. Confirm that the LIFs were reverted:

```
network interface show -vserver Cluster -fields curr-port,is-home
```

### Example

The LIFs are reverted if the “Is Home” column is true for node1\_clus1 and node2\_clus1:

```
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.
```

17. Migrate the LIFs on the remaining cluster port connected to the switch on node1 to the cluster port directly connected to node2:


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

### Example

The following example migrates LIF node1\_clus2 from port e0b on node1 to port e0a:

```
cluster::> network interface migrate -vserver Cluster -lif node1_clus2 -destination-node node1 -destination-port e0a
```

The LIF is now associated with port e0a on node1, as indicated by the gray line from LIF clus2 to port e0a in the following illustration:



18. Migrate the LIFs on the remaining cluster port connected to the switch on node2 to the cluster port directly connected to node1:


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

### Example

The following example shows that LIF node2\_clus2 is migrated from port e0b on node2 to port e0a.

```
cluster:: network interface migrate -vserver Cluster -lif node2_clus2 -destination-node node2 -destination-port e0a
```

The LIF is now associated with port e0a on node2, as indicated by the gray line from LIF clus2 to port e0a in the following illustration:



19. Confirm that node1\_clus2 and node2\_clus2 were successfully migrated to port e0a on their home nodes. The LIF's current port should be e0a and "is-home" should be false.

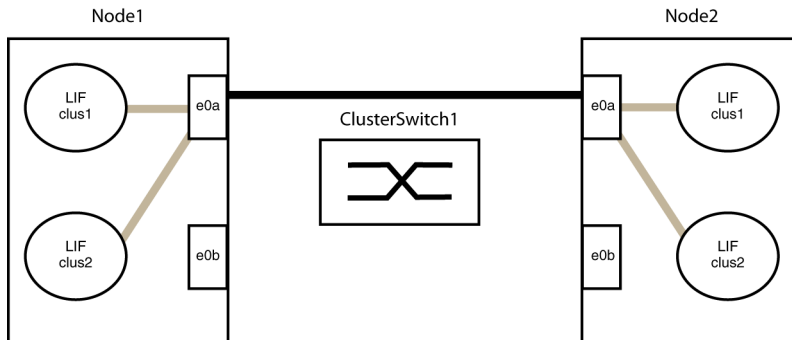
**Example**

```
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       e0a      false
Cluster node2_clus1       e0a      true
Cluster node2_clus2       e0a      false
4 entries were displayed.
```

20. Disconnect the cables linking the switches to the ports that no longer have traffic.

**Example**

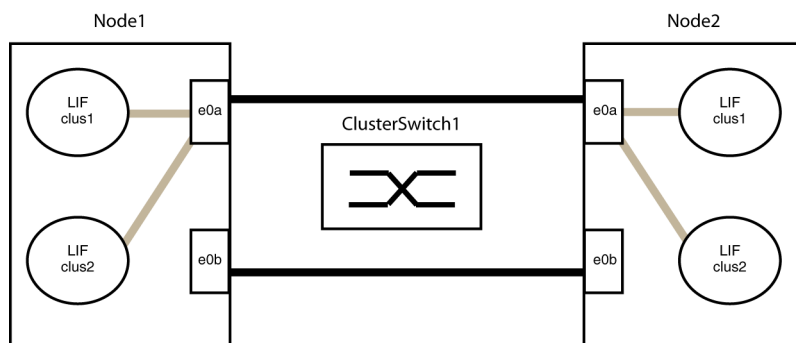
In the following illustration, the cables to the switch are removed, and cluster traffic continues through the direct connection between the e0a ports:



21. Cable the ports that were formerly connected to the switch to each other, creating a direct connection between the nodes.

**Example**





22. After the ports have been cabled and are able to process network traffic, the cluster LIFs automatically revert back to their home ports. Confirm that the LIFs have been reverted:

```
network interface show -vserver Cluster -fields curr-port,is-home
```

### Example

The LIFs have been reverted if the “Is Home” column is true for node1\_clus2 and node2\_clus2:

```
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.
```

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

```
network device-discovery show
```

### Example

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

```
cluster::> network device-discovery show -node node1 -port e0a -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
node1/cd   e0a    node2                      e0a        FAS2552

cluster::> network device-discovery show -node node1 -port e0b -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
node1/cdp  e0b    node2                      e0b        FAS2552

cluster::> network device-discovery show -node node2 -port e0a -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
node1/cdp  e0a    node1                      e0a        FAS2552
cluster::> network device-discovery show -node node2 -port e0b -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
node1/cdp  e0b    node1                      e0b        FAS2552
```

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

```
cluster 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.
```

25. Confirm connectivity between the cluster ports:

```
cluster ping-cluster local
```

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

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

27. 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](#)

## Transitioning to a two-node switchless cluster running ONTAP 9.1 or later (FAS22xx systems with a single cluster interconnect connection)

If you have FAS22xx systems in an existing two-node cluster, you can replace the switch with a single, direct, back-to-back connection between the nodes. This is a nondisruptive operation. This configuration does not provide a redundant connection if there is a cable failure or port failure.

### Before you begin

- The cluster must be healthy and consist of two nodes that are connected by a cluster switch, and the nodes must be running ONTAP 9.1 or later.


**Note:** The switchless cluster interconnect feature cannot be used with more than two nodes.

- Each node must use a single, dedicated 10 GbE cluster port that provides the cluster interconnect connection.

**Note:** This configuration is only supported for FAS22xx systems only.

### About this task

The following procedure removes the cluster switch in a two-node cluster and replaces the connection to the switch with a direct 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 from 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:

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

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=<number of hours>h
```

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

The following command suppresses automatic case creation for two hours:

### Example

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

- For ONTAP release 9.1 only, enable the switchless-cluster network option:

```
network option switchless-cluster modify -enabled true
```

### Example

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

- Designate one of the nodes as the target node. Take over the target node, entering *y* when prompted to continue:

```
storage failover takeover -ofnode target_node_name
```

As part of the takeover, the target node automatically reboots and displays the `waiting for giveback..` message. The node that is not waiting for giveback is now serving data for the partner (target) node that was taken over.

- Wait for two minutes after takeover of the impaired node to confirm that the takeover was completed successfully.
- After the target node shows the `waiting for giveback..` message, shut down the target node.

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

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

- On each controller module, disconnect the cable that connects the 10 GbE cluster port to the cluster switch.
- Connect the cluster port of each controller module to the cluster port on the other controller module, so that the controller modules have a back-to-back network connection.
- Restart the target node.

If the SP ...	Then...
Is configured	Log in to the SP of the target node, and then turn the power on from the console: <code>system power off</code> .
Is not configured	At the prompt of the target node, boot the node to ONTAP: <code>boot_ontap</code> .

The target node boots and displays the `waiting for giveback...` message.

- Give back storage to the target node:  
`storage failover giveback -ofnode target_node_name`
- 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 that takeover is possible:

```
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.
```

15. Verify that the cluster port LIFs are operating correctly:

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

### Example

The following example shows that the LIFs are operationally 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        e0a         true
node2        clus1      up/up       192.168.177.123/24  node2        e0a         true

2 entries were displayed.
```

16. Check the cluster status of the nodes from 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:

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

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=ENDcluster::*> 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](#)

# Transitioning to a two-node switchless cluster (systems with four redundant cluster interconnect connections)


If you have an existing two-node cluster with four-port cluster interconnects that are connected to cluster interconnect switches and are running ONTAP 9.1 or later, you can replace the switches with direct, back-to-back connections between the nodes. This is a nondisruptive operation.

## Before you begin

- The cluster must be healthy and consist of two nodes that are connected by cluster switches, and the nodes must be running the same ONTAP release of 9.1 or later.  
The switchless cluster interconnect feature cannot be used with more than two nodes.
- Each node must have four dedicated cluster ports per controller module, which provide redundant cluster interconnect connections  
This is the supported configuration for most nodes that support four cluster interconnects.

## 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:



## 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 from 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:

```

Node Health Eligibility
-----
node1 true true
node2 true true

2 entries were displayed.

```

- If any node shows false for health or eligibility, correct the problem.
- ONTAP releases 9.2 and later support automatic detection of switchless clusters which is enabled by default. For this release verify that detection of switchless clusters is enabled by running the advanced privilege command:  
**network options detect-switchless-cluster show**

### Example

The following example shows if the option is enabled:

```

cluster::*> net options detect-switchless-cluster show
(network options detect-switchless-cluster show)

Enable Switchless Cluster Detection: true

```

If “Enable Switchless Cluster Detection” is false contact support.

- 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 AutoSupport message notifies technical support of this maintenance task so that they can suppress automatic case creation 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

```

- Identify the cluster ports and verify link status and health:

**network port show -ipspace Cluster**

The following example shows the nodes that are using e0a and e0b as cluster ports. Your nodes might be using different cluster ports as they vary by platform.

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

```

cluster::> network port show -ipspace Cluster

Node: node1

Port      IPspace      Broadcast Domain  Link  MTU  Speed(Mbps) Health
-----
e0a       Cluster     Cluster           up    9000 auto/10000 healthy
e0b       Cluster     Cluster           up    9000 auto/10000 healthy
e0c       Cluster     Cluster           up    9000 auto/10000 healthy
e0d       Cluster     Cluster           up    9000 auto/10000 healthy

Node: node2

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

```

```

e0b      Cluster      Cluster      up    9000  auto/10000 healthy
e0c      Cluster      Cluster      up    9000  auto/10000 healthy
e0d      Cluster      Cluster      up    9000  auto/10000 healthy
entries were displayed

```

- Verify that all of the cluster LIFs are on their home ports.

Confirm that the “Is Home column is true for each of the cluster LIFs:

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

The following command shows if the cluster LIFs are home:

```

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 node1_clus3  true
Cluster node1_clus4  true
Cluster node2_clus1  true
Cluster node2_clus2  true
Cluster node2_clus3  true
Cluster node2_clus4  true
4 entries were displayed.

```

- Verify that all cluster LIFs have auto-revert enabled:

```
network interface show -vserver Cluster -fields auto-revert
```

#### Example

```

cluster::> network interface show -vserver Cluster -fields auto-revert
vserver lif          auto-revert
-----
Cluster node1_clus1  true
Cluster node1_clus2  true
Cluster node1_clus3  true
Cluster node1_clus4  true
Cluster node2_clus1  true
Cluster node2_clus2  true
Cluster node2_clus3  true
Cluster node2_clus4  true
8 entries were displayed.

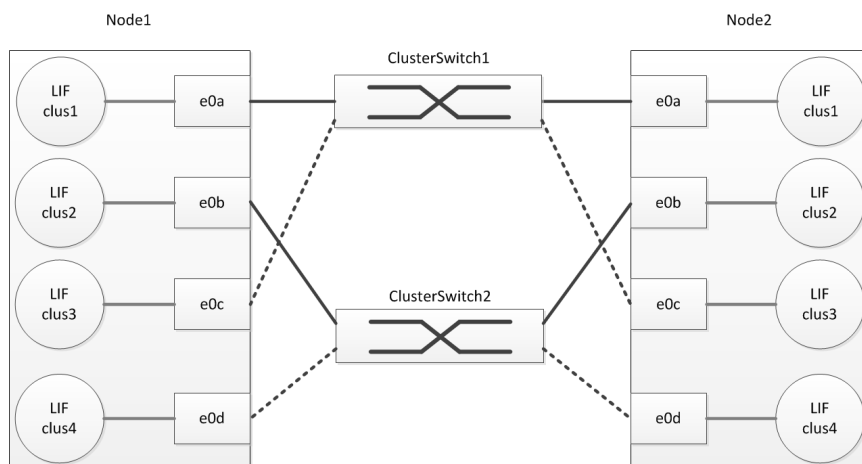
```

- Review how the cluster network is cabled.

The following illustration shows a simplified view of the two nodes, node1 and node2:

- Each node uses ports e0a, e0b, e0c, and e0d for the cluster connections.
- Each cluster port has one cluster LIF associated with it.
- The following illustration shows the connections to both *ClusterSwitch1* and *ClusterSwitch2*.





10. For ONTAP release 9.1 only, enable the switchless-cluster network option:

**network option switchless-cluster modify -enabled true**

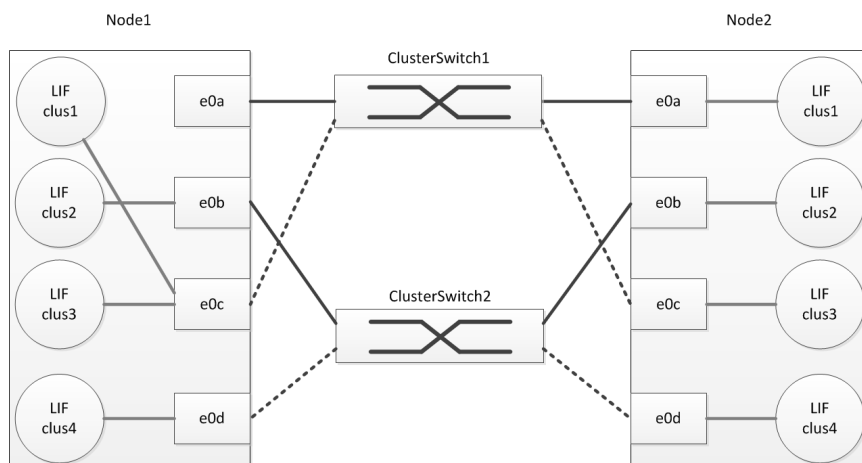
```
cluster::*> network options switchless-cluster modify -enabled true
```

11. Migrate the LIFs on the first cluster port of node1 to the third cluster port on node1:

**network interface migrate -vserver Cluster -lif lif\_name -destination-node destination-node-name -destination-port port-name**

```
network interface migrate -vserver Cluster -lif node1_clus1 -destination-node node1 -
destination-port e0c
```

The LIF node1\_clus1 is now associated with port e0c on node1, as indicated by the gray line from LIF node1\_clus1 to port e0c in the following illustration:

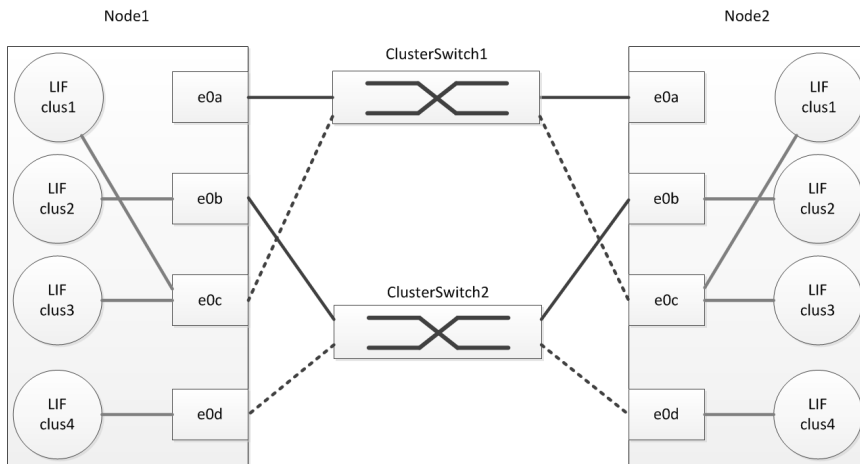


12. Migrate the LIFs on the first cluster port of node2 to the third cluster port on node2:

**network interface migrate -vserver Cluster -lif lif\_name -destination-node destination-node-name -destination-port port-name**

```
network interface migrate -vserver Cluster -lif node2_clus1 -destination-node node2 -
destination-port e0c
```

The LIF node2\_clus1 is now associated with port e0c on node2, as indicated by the gray line from LIF node2\_clus1 to port e0c in the following illustration:



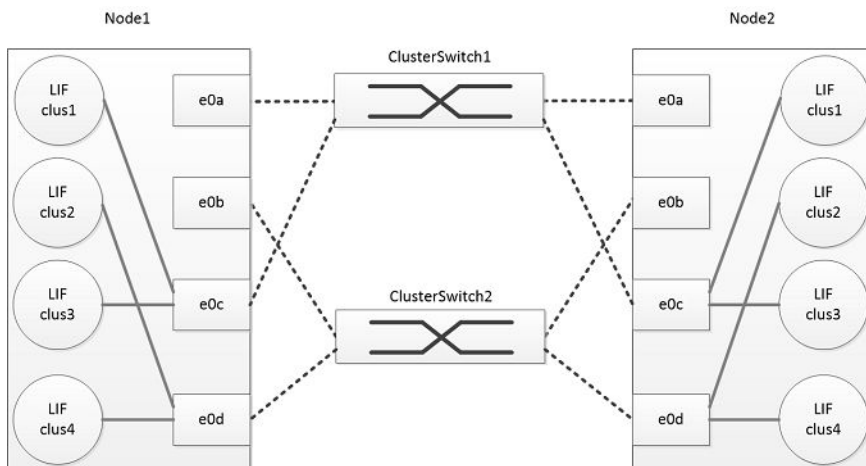
13. Repeat step 11 on page 17 through step 12 on page 17 to migrate the LIFs from second cluster port to the fourth cluster port on node 1 and node 2.
14. Confirm that the LIFs are migrated:

**network interface show -role cluster -fields curr-port,is-home**

### Example

In the following example, node1\_clus1 and node2\_clus1 have been migrated to port e0c on their respective nodes. Node1\_clus2 and node2\_clus2 have been migrated to port e0d on their respective nodes:


```
cluster::> network interface show -vserver Cluster -fields curr-port,is-home
vserver lif                curr-port is-home
-----
Cluster node1_clus1       e0c      false
Cluster node1_clus2       e0d      false
Cluster node1_clus3       e0c      true
Cluster node1_clus4       e0d      true
Cluster node2_clus1       e0c      false
Cluster node2_clus2       e0d      false
Cluster node2_clus3       e0c      true
Cluster node2_clus4       e0d      true
8 entries were displayed.
```



- Disconnect the cables connecting the switches from the cluster ports that no longer host cluster LIFs.

**Example**

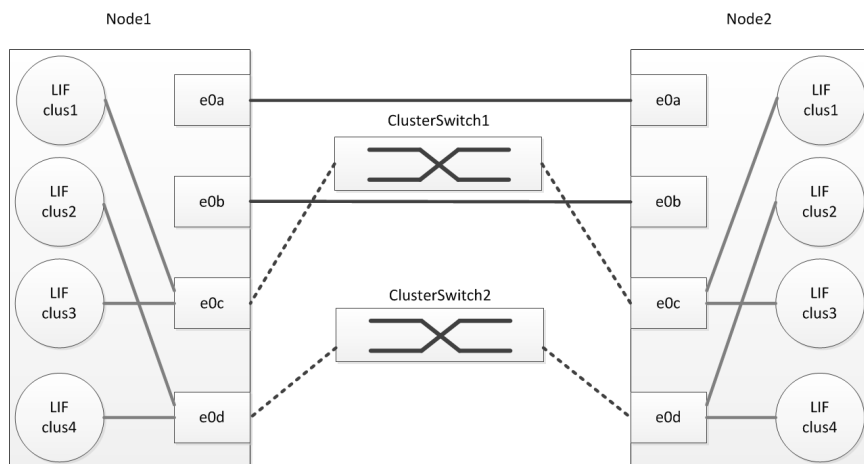
In the following example, the cables are disconnected from port e0a and e0b on each node, and cluster traffic continues through the switch and ports e0c and e0d on each node:



- Cable the ports that were formerly connected to the switch to each other, creating a direct connection between the controllers.

**Example**

In the following example, e0a and e0b on node1 are connected to their corresponding ports, e0a and e0b on node2:



- For ONTAP releases 9.2 and later, 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
```

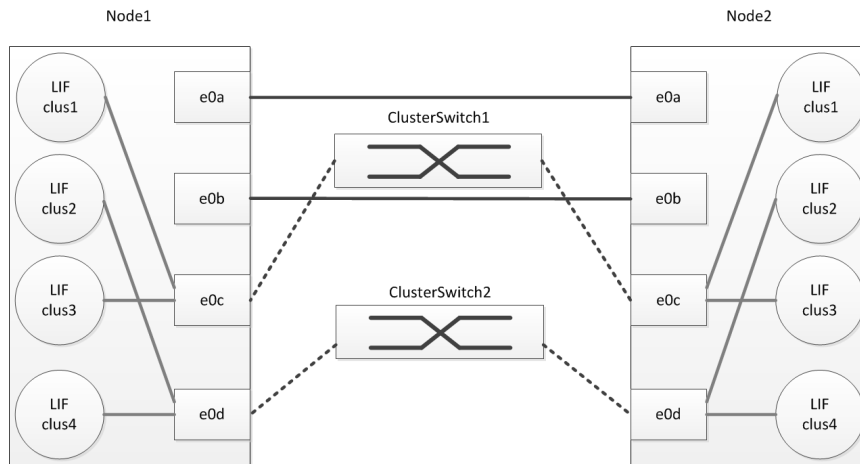
18. Once the ports have been cabled and are able to process network traffic, the cluster LIFs automatically revert back to their home ports. Confirm that the LIFs are reverted:

```
network interface show -vserver Cluster -fields curr-port,is-home
```

The LIFs are reverted if the “Is Home” column is true for node1\_clus1, node1\_clus2, node2\_clus1 and node2\_clus2:

```
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 node1_clus3       e0c      true
Cluster node1_clus4       e0d      true
Cluster node2_clus1       e0a      true
Cluster node2_clus2       e0b      true
Cluster node2_clus3       e0c      true
Cluster node2_clus4       e0d      true
      8 entries were displayed.
```

The cluster LIFs are now hosted on their home ports and the directly connected cluster ports are processing the network packets.



19. Verify that there are no communication issues between the nodes in the cluster and that all cluster port have a value of up for the “Link” column:

```
network port show -ipSPACE Cluster
```

#### Example

```
Cluster*::> network port show -ipSPACE
Cluster
```

Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Oper
-----						
node1						
	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
node2						
	e0a	Cluster	Cluster	up	9000	auto/10000
	e0b	Cluster	Cluster	up	9000	auto/10000

```

      e0c      Cluster      Cluster      up      9000      auto/10000
      e0d      Cluster      Cluster      up      9000      auto/10000

8 entries were displayed.

```

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

**cluster show**

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

```

Node Health Eligibility
-----
node1 true true
node2 true true

2 entries were displayed.

```

21. If any node shows false for health or eligibility, correct the problem.

22. Display information about the devices in your configuration:

**network device-discovery show -port <cluster-port name> -protocol cdp**

Verify that the ports on both nodes are correctly connected:

**network device-discovery show**

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

```

cluster::> network device-discovery show -node node1 -port e0a -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface      Platform
-----
node1/cdp
      e0a   node2                e0a            FAS2552

cluster::> network device-discovery show -node node1 -port e0b -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface      Platform
-----
node1/cdp
      e0b   node2                e0b            FAS2552

cluster::> network device-discovery show -node node2 -port e0a -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface      Platform
-----
node1/cdp
      e0a   node1                e0a            FAS2552

cluster::> network device-discovery show -node node2 -port e0b -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface      Platform
-----
node1/cdp
      e0b   node1                e0b            FAS2552

```

23. Migrate the LIFs on the third cluster port of node1 to the first cluster port on node1:

**network interface migrate -vserver Cluster -lif lif\_name -destination-node destination-node-name -destination-port port-name**

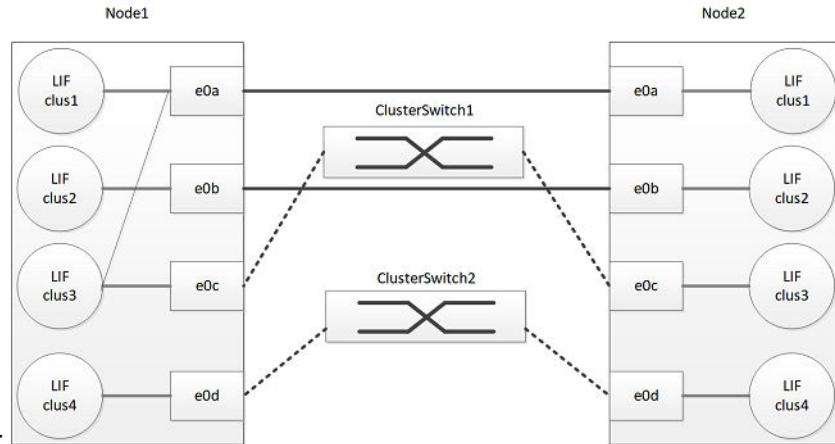
```

.

network interface migrate -vserver Cluster -lif node1_clus3 -destination-node node1 -
destination-port e0a

```

The LIF node1\_clus3 is now associated with port e0a on node1, as indicated by the gray line from LIF node1\_clus3 to port e0a in the following illustration:



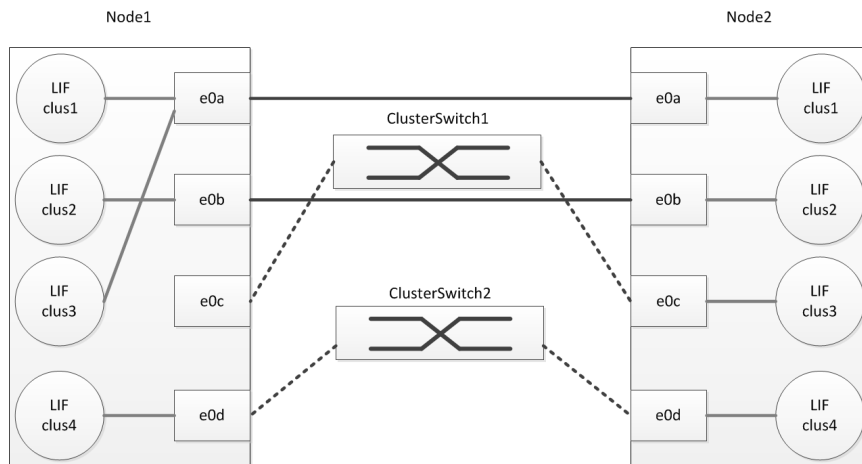
e0a in the following illustration:

- Migrate the LIFs on the third cluster port of node2 to the first cluster port on node2:

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

```
network interface migrate -vserver Cluster -lif node2_clus3 -destination-node node2 -destination-port e0a
```

The LIF node1\_clus3 is now associated with port e0a on node2, as indicated by the gray line from LIF node2\_clus3 to port e0a in the following illustration:



- Repeat step 11 on page 17 through step 12 on page 17 to migrate the LIFs from fourth cluster port to the second cluster port on node 1 and node 2.

- Confirm that the LIFs have been migrated:

```
network interface show -role cluster -fields curr-port,is-home
```

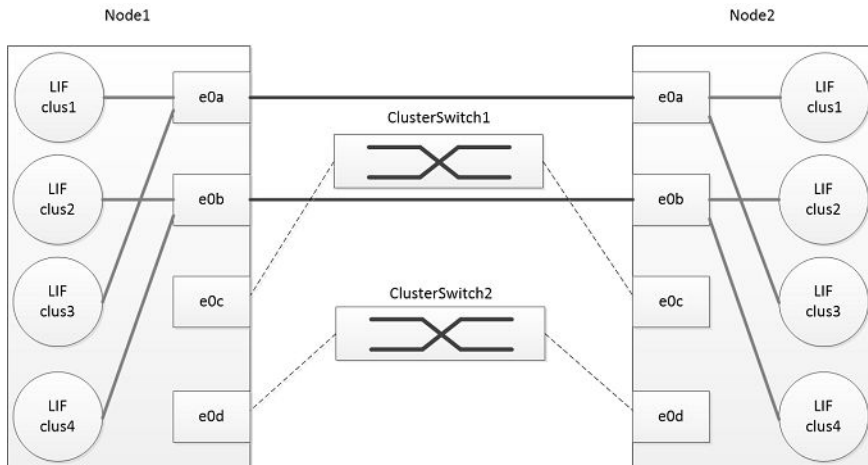
In the following example, node1\_clus3 and node2\_clus3 have been migrated to port e0a on their respective nodes. Node1\_clus4 and node2\_clus4 have been migrated to port e0b on their respective nodes:

```
cluster::> network interface show -vserver Cluster -fields curr-port,is-home
vserver lif                               curr-port is-home
-----
Cluster node1_clus1                       e0a       true
```

```

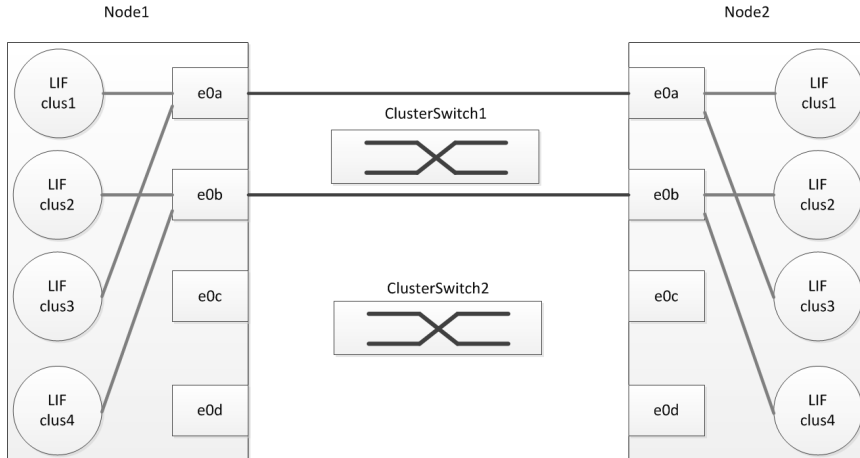
Cluster node1_clus2      e0b      true
Cluster node1_clus3      e0a      false
Cluster node1_clus4      e0b      false
Cluster node2_clus1      e0a      true
Cluster node2_clus2      e0b      true
Cluster node2_clus3      e0a      false
Cluster node2_clus4      e0b      false
8 entries were displayed.

```

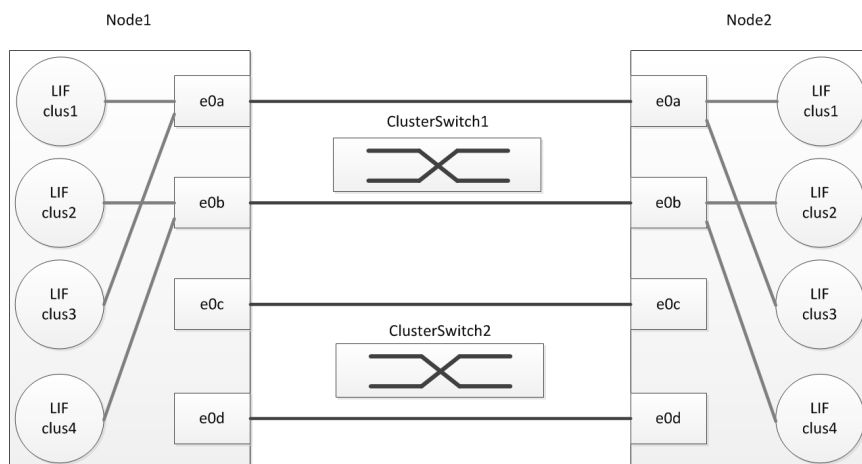


27. Disconnect the cables linking the switches to the ports that no longer host cluster LIFs.

In the following illustration, the cables to the switch are removed, and cluster traffic continues through the direct connection between the e0a and e0b ports:



28. Cable the ports that were formerly connected to the switch to each other, creating a direct connection between the nodes.



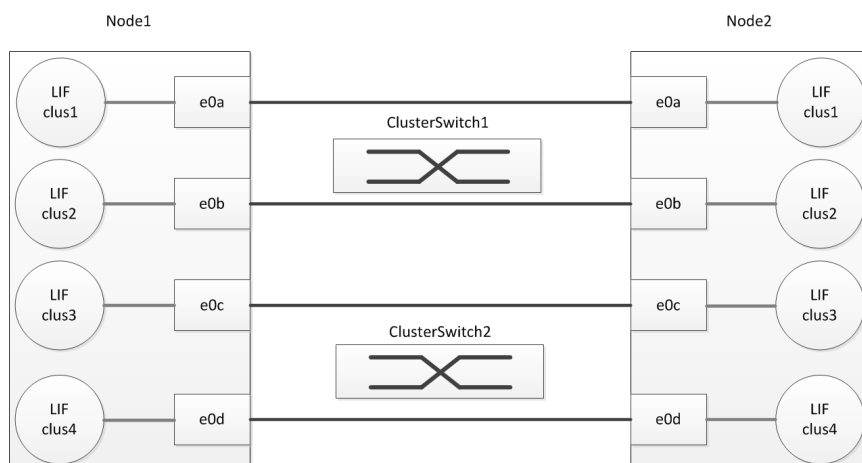
29. After the ports have been cabled and are able to process network traffic, the cluster LIFs automatically revert back to their home ports. Confirm that the LIFs were reverted:

**network interface show -vserver Cluster -fields curr-port,is-home**

The LIFs are reverted if the “Is Home” column is true for node1\_clus1 and node2\_clus1:

```
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 node1_clus3 e0c true
Cluster node1_clus4 e0d true
Cluster node2_clus1 e0a true
Cluster node2_clus2 e0b true
Cluster node2_clus3 e0c true
Cluster node2_clus4 e0d true
8 entries were displayed.
```

The cluster LIFs are now hosted on their home ports and the directly connected cluster ports are processing network packets.



30. Check that all ports have a value of up for the “Link” column, the communication between the nodes is good, and that no remote procedure call failures are displayed:

**network port show -ipSPACE Cluster**



Speed (Mbps)	Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper
	node1						
		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
	node2						
		e0a	Cluster	Cluster	up	9000	auto/10000
		e0b	Cluster	Cluster	up	9000	auto/10000
		e0c	Cluster	Cluster	up	9000	auto/10000
		e0d	Cluster	Cluster	up	9000	auto/10000

8 entries were displayed.

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

**cluster show**

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

Node	Health	Eligibility
node1	true	true
node2	true	true

2 entries were displayed.

32. If any node shows false for health or eligibility, correct the problem.

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

**network device-discovery show -protocol cdp**

The following example shows that cluster ports e0c and e0d are correctly connected to the corresponding port on the cluster partner:

```
cid="OwvSy">cluster::> network device-discovery show -node node1 -port e0c -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface      Platform
node1/cdp  e0c    node2                      e0c            FAS2552

cluster::> network device-discovery show -node node1 -port e0d -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface      Platform
node1/cdp  e0d    node2                      e0d            FAS2552

cluster::> network device-discovery show -node node2 -port e0c -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface      Platform
Node2/cdp  e0c    node1                      e0c            FAS2552

cluster::> network device-discovery show -node node2 -port e0d -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface      Platform
Node2/cdp  e0d    node1                      e0d            FAS2552
```

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

**cluster show**

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

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

2 entries were displayed.
```

35. 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.

36. If you suppressed automatic case creation, reenable 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
```

37. 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](#)

## Copyright

Copyright © 2019 NetApp, Inc. All rights reserved. Printed in the U.S.

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

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

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

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

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

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

# Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

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

## 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](mailto: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