



Replace switches

Cluster and storage switches

NetApp

December 12, 2024

Table of Contents

- Replace switches 1
 - Requirements for replacing Cisco Nexus 3132Q-V cluster switches 1
 - Replace Cisco Nexus 3132Q-V cluster switches 4
 - Replace Cisco Nexus 3132Q-V cluster switches with switchless connections 30

Replace switches

Requirements for replacing Cisco Nexus 3132Q-V cluster switches

Make sure you understand the configuration requirements, port connections, and cabling requirements when you replace cluster switches.

Cisco Nexus 3132Q-V requirements

- The Cisco Nexus 3132Q-V cluster switch is supported.
- The number of 10 GbE and 40 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.
- The cluster switches use the Inter-Switch Link (ISL) ports e1/31-32.
- The [Hardware Universe](#) contains information about supported cabling to Nexus 3132Q-V switches:
 - The nodes with 10 GbE cluster connections require QSFP optical modules with breakout fiber cables or QSFP to SFP+ copper break-out cables.
 - The nodes with 40 GbE cluster connections require supported QSFP/QSFP28 optical modules with fiber cables or QSFP/QSFP28 copper direct-attach cables.
 - The cluster switches use the appropriate ISL cabling: 2x QSFP28 fiber or copper direct-attach cables.
- On Nexus 3132Q-V, you can operate QSFP ports as either 40 Gb Ethernet or 4x10 Gb Ethernet modes.

By default, there are 32 ports in the 40 Gb Ethernet mode. These 40 Gb Ethernet ports are numbered in a 2-tuple naming convention. For example, the second 40 Gb Ethernet port is numbered as 1/2. The process of changing the configuration from 40 Gb Ethernet to 10 Gb Ethernet is called *breakout* and the process of changing the configuration from 10 Gb Ethernet to 40 Gb Ethernet is called *breakin*. When you break out a 40 Gb Ethernet port into 10 Gb Ethernet ports, the resulting ports are numbered using a 3-tuple naming convention. For example, the breakout ports of the second 40 Gb Ethernet port are numbered as 1/2/1, 1/2/2, 1/2/3, and 1/2/4.

- On the left side of Nexus 3132Q-V is a set of four SFP+ ports multiplexed to the first QSFP port.

By default, the RCF is structured to use the first QSFP port.

You can make four SFP+ ports active instead of a QSFP port for Nexus 3132Q-V by using the hardware profile `front portmode sfp-plus` command. Similarly, you can reset Nexus 3132Q-V to use a QSFP port instead of four SFP+ ports by using the hardware profile `front portmode qsfp` command.

- You must have configured some of the ports on Nexus 3132Q-V to run at 10 GbE or 40 GbE.

You can break-out the first six ports into 4x10 GbE mode by using the `interface breakout module 1 port 1-6 map 10g-4x` command. Similarly, you can regroup the first six QSFP+ ports from breakout configuration by using the `no interface breakout module 1 port 1-6 map 10g-4x` command.

- You must have done the planning, migration, and read the required documentation on 10 GbE and 40 GbE connectivity from nodes to Nexus 3132Q-V cluster switches.

The [Cisco Ethernet Switches](#) page has information about the ONTAP and NX-OS versions supported in

this procedure.

Cisco Nexus 5596 requirements

- The following cluster switches are supported:
 - Nexus 5596
 - Nexus 3132Q-V
- The number of 10 GbE and 40 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.
- The cluster switches use the following ports for connections to nodes:
 - Ports e1/1-40 (10 GbE): Nexus 5596
 - Ports e1/1-30 (40 GbE): Nexus 3132Q-V
- The cluster switches use the following Inter-Switch Link (ISL) ports:
 - Ports e1/41-48 (10 GbE): Nexus 5596
 - Ports e1/31-32 (40 GbE): Nexus 3132Q-V
- The [Hardware Universe](#) contains information about supported cabling to Nexus 3132Q-V switches:
 - Nodes with 10 GbE cluster connections require QSFP to SFP+ optical fiber breakout cables or QSFP to SFP+ copper breakout cables.
 - Nodes with 40 GbE cluster connections require supported QSFP/QSFP28 optical modules with fiber cables or QSFP/QSFP28 copper direct-attach cables.
- The cluster switches use the appropriate ISL cabling:
 - Beginning: Nexus 5596 to Nexus 5596 (SFP+ to SFP+)
 - 8x SFP+ fiber or copper direct-attach cables
 - Interim: Nexus 5596 to Nexus 3132Q-V (QSFP to 4xSFP+ break-out)
 - 1x QSFP to SFP+ fiber break-out or copper break-out cables
 - Final: Nexus 3132Q-V to Nexus 3132Q-V (QSFP28 to QSFP28)
 - 2x QSFP28 fiber or copper direct-attach cables
- On Nexus 3132Q-V switches, you can operate QSFP/QSFP28 ports as either 40 Gigabit Ethernet or 4 x10 Gigabit Ethernet modes.

By default, there are 32 ports in the 40 Gigabit Ethernet mode. These 40 Gigabit Ethernet ports are numbered in a 2-tuple naming convention. For example, the second 40 Gigabit Ethernet port is numbered as 1/2. The process of changing the configuration from 40 Gigabit Ethernet to 10 Gigabit Ethernet is called *breakout* and the process of changing the configuration from 10 Gigabit Ethernet to 40 Gigabit Ethernet is called *breakin*. When you break out a 40 Gigabit Ethernet port into 10 Gigabit Ethernet ports, the resulting ports are numbered using a 3-tuple naming convention. For example, the break-out ports of the second 40 Gigabit Ethernet port are numbered as 1/2/1, 1/2/2, 1/2/3, and 1/2/4.

- On the left side of Nexus 3132Q-V switches is a set of 4 SFP+ ports multiplexed to that QSFP28 port.

By default, the RCF is structured to use the QSFP28 port.



You can make 4x SFP+ ports active instead of a QSFP port for Nexus 3132Q-V switches by using the `hardware profile front portmode sfp-plus` command. Similarly, you can reset Nexus 3132Q-V switches to use a QSFP port instead of 4x SFP+ ports by using the `hardware profile front portmode qsfp` command.

- You have configured some of the ports on Nexus 3132Q-V switches to run at 10 GbE or 40 GbE.



You can break out the first six ports into 4x10 GbE mode by using the `interface breakout module 1 port 1-6 map 10g-4x` command. Similarly, you can regroup the first six QSFP+ ports from breakout configuration by using the `no interface breakout module 1 port 1-6 map 10g-4x` command.

- You have done the planning, migration, and read the required documentation on 10 GbE and 40 GbE connectivity from nodes to Nexus 3132Q-V cluster switches.
- The ONTAP and NX-OS versions supported in this procedure are on the [Cisco Ethernet Switches](#) page.

NetApp CN1610 requirements

- The following cluster switches are supported:
 - NetApp CN1610
 - Cisco Nexus 3132Q-V
- The cluster switches support the following node connections:
 - NetApp CN1610: ports 0/1 through 0/12 (10 GbE)
 - Cisco Nexus 3132Q-V: ports e1/1-30 (40 GbE)
- The cluster switches use the following inter-switch link (ISL) ports:
 - NetApp CN1610: ports 0/13 through 0/16 (10 GbE)
 - Cisco Nexus 3132Q-V: ports e1/31-32 (40 GbE)
- The [Hardware Universe](#) contains information about supported cabling to Nexus 3132Q-V switches:
 - Nodes with 10 GbE cluster connections require QSFP to SFP+ optical fiber breakout cables or QSFP to SFP+ copper breakout cables
 - Nodes with 40 GbE cluster connections require supported QSFP/QSFP28 optical modules with optical fiber cables or QSFP/QSFP28 copper direct-attach cables
- The appropriate ISL cabling is as follows:
 - Beginning: For CN1610 to CN1610 (SFP+ to SFP+), four SFP+ optical fiber or copper direct-attach cables
 - Interim: For CN1610 to Nexus 3132Q-V (QSFP to four SFP+ breakout), one QSFP to SFP+ optical fiber or copper breakout cable
 - Final: For Nexus 3132Q-V to Nexus 3132Q-V (QSFP28 to QSFP28), two QSFP28 optical fiber or copper direct-attach cables
- NetApp twinax cables are not compatible with Cisco Nexus 3132Q-V switches.

If your current CN1610 configuration uses NetApp twinax cables for cluster-node-to-switch connections or ISL connections and you want to continue using twinax in your environment, you need to procure Cisco twinax cables. Alternatively, you can use optical fiber cables for both the ISL connections and the cluster-node-to-switch connections.

- On Nexus 3132Q-V switches, you can operate QSFP/QSFP28 ports as either 40 Gb Ethernet or 4x 10 Gb Ethernet modes.

By default, there are 32 ports in the 40 Gb Ethernet mode. These 40 Gb Ethernet ports are numbered in a 2-tuple naming convention. For example, the second 40 Gb Ethernet port is numbered as 1/2. The process of changing the configuration from 40 Gb Ethernet to 10 Gb Ethernet is called *breakout* and the process of changing the configuration from 10 Gb Ethernet to 40 Gb Ethernet is called *breakin*. When you break out a 40 Gb Ethernet port into 10 Gb Ethernet ports, the resulting ports are numbered using a 3-tuple naming convention. For example, the breakout ports of the second 40 Gb Ethernet port are numbered as 1/2/1, 1/2/2, 1/2/3, and 1/2/4.

- On the left side of Nexus 3132Q-V switches is a set of four SFP+ ports multiplexed to the first QSFP port.

By default, the reference configuration file (RCF) is structured to use the first QSFP port.

You can make four SFP+ ports active instead of a QSFP port for Nexus 3132Q-V switches by using the `hardware profile front portmode sfp-plus` command. Similarly, you can reset Nexus 3132Q-V switches to use a QSFP port instead of four SFP+ ports by using the `hardware profile front portmode qsfp` command.



When you use the first four SFP+ ports, it will disable the first 40GbE QSFP port.

- You must have configured some of the ports on Nexus 3132Q-V switches to run at 10 GbE or 40 GbE.

You can break out the first six ports into 4x10 GbE mode by using the `interface breakout module 1 port 1-6 map 10g-4x` command. Similarly, you can regroup the first six QSFP+ ports from *breakout* configuration by using the `no interface breakout module 1 port 1-6 map 10g-4x` command.

- You must have done the planning, migration, and read the required documentation on 10 GbE and 40 GbE connectivity from nodes to Nexus 3132Q-V cluster switches.
- The ONTAP and NX-OS versions that are supported in this procedure are listed on the [Cisco Ethernet Switches](#) page.
- The ONTAP and FASTPATH versions that are supported in this procedure are listed on the [NetApp CN1601 and CN1610 Switches](#) page.

Replace Cisco Nexus 3132Q-V cluster switches

Follow this procedure to replace a defective Cisco Nexus 3132Q-V switch in a cluster network. The replacement procedure is a nondisruptive procedure (NDO).

Review requirements

Switch requirements

Review the [Requirements for replacing Cisco Nexus 3132Q-V cluster switches](#).

Before you begin

- The existing cluster and network configuration has:
 - The Nexus 3132Q-V cluster infrastructure is redundant and fully functional on both switches.

The [Cisco Ethernet Switch](#) page has the latest RCF and NX-OS versions on your switches.

- All cluster ports are in the `up` state.
- Management connectivity exists on both switches.
- All cluster logical interfaces (LIFs) are in the `up` state and have been migrated.
- For the Nexus 3132Q-V replacement switch, make sure that:
 - Management network connectivity on the replacement switch is functional.
 - Console access to the replacement switch is in place.
 - The desired RCF and NX-OS operating system image switch is loaded onto the switch.
 - Initial customization of the switch is complete.
- [Hardware Universe](#)

Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when replacing your switch:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport before and after maintenance to disable case creation for the duration of the maintenance. See this Knowledge Base article [SU92: How to suppress automatic case creation during scheduled maintenance windows](#) for further details.
- Enable session logging for any CLI sessions. For instructions on how to enable session logging, review the "Logging Session Output" section in this Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

Replace the switch

This procedure replaces the second Nexus 3132Q-V cluster switch CL2 with new 3132Q-V switch C2.

About the examples

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

- `n1_clus1` is the first cluster logical interface (LIF) connected to cluster switch C1 for node n1.
- `n1_clus2` is the first cluster LIF connected to cluster switch CL2 or C2, for node n1.
- `n1_clus3` is the second LIF connected to cluster switch C2, for node n1.
- `n1_clus4` is the second LIF connected to cluster switch CL1, for node n1.
- The number of 10 GbE and 40 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.
- The nodes are n1, n2, n3, and n4. - The examples in this procedure use four nodes: Two nodes use four 10 GB cluster interconnect ports: e0a, e0b, e0c, and e0d. The other two nodes use two 40 GB cluster interconnect ports: e4a and e4e. See the [Hardware Universe](#) for the actual cluster ports on your platforms.

About this task

This procedure covers the following scenario:

- The cluster starts with four nodes connected to two Nexus 3132Q-V cluster switches, CL1 and CL2.
- Cluster switch CL2 is to be replaced by C2

- On each node, cluster LIFs connected to CL2 are migrated onto cluster ports connected to CL1.
- Disconnect cabling from all ports on CL2 and reconnect cabling to the same ports on the replacement switch C2.
- On each node, its migrated cluster LIFs are reverted.

Step 1: Prepare for replacement

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

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

x is the duration of the maintenance window in hours.



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

2. Display information about the devices in your configuration:

```
network device-discovery show
```

Show example

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

Node	Local Port	Discovered Device	Interface	Platform
n1	/cdp			
	e0a	CL1	Ethernet1/1/1	N3K-C3132Q-V
	e0b	CL2	Ethernet1/1/1	N3K-C3132Q-V
	e0c	CL2	Ethernet1/1/2	N3K-C3132Q-V
	e0d	CL1	Ethernet1/1/2	N3K-C3132Q-V
n2	/cdp			
	e0a	CL1	Ethernet1/1/3	N3K-C3132Q-V
	e0b	CL2	Ethernet1/1/3	N3K-C3132Q-V
	e0c	CL2	Ethernet1/1/4	N3K-C3132Q-V
	e0d	CL1	Ethernet1/1/4	N3K-C3132Q-V
n3	/cdp			
	e4a	CL1	Ethernet1/7	N3K-C3132Q-V
	e4e	CL2	Ethernet1/7	N3K-C3132Q-V
n4	/cdp			
	e4a	CL1	Ethernet1/8	N3K-C3132Q-V
	e4e	CL2	Ethernet1/8	N3K-C3132Q-V

12 entries were displayed

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

a. Display the network port attributes:

```
network port show
```

Show example

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

Node: n1

Ignore

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

Node: n2

Ignore

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

Node: n3

Ignore

Health      Health      Speed (Mbps)
```

```
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper
Status    Status
-----
-----
e4a      Cluster      Cluster      up   9000 auto/40000 -
-
e4e      Cluster      Cluster      up   9000 auto/40000 -
-

Node: n4

Ignore

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

Speed (Mbps)

12 entries were displayed.
```

b. Display information about the logical interfaces:

```
network interface show
```

Show example

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

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

Cluster				
e0a	n1_clus1	up/up	10.10.0.1/24	n1
e0b	n1_clus2	up/up	10.10.0.2/24	n1
e0c	n1_clus3	up/up	10.10.0.3/24	n1
e0d	n1_clus4	up/up	10.10.0.4/24	n1
e0a	n2_clus1	up/up	10.10.0.5/24	n2
e0b	n2_clus2	up/up	10.10.0.6/24	n2
e0c	n2_clus3	up/up	10.10.0.7/24	n2
e0d	n2_clus4	up/up	10.10.0.8/24	n2
e0a	n3_clus1	up/up	10.10.0.9/24	n3
e0e	n3_clus2	up/up	10.10.0.10/24	n3
e0a	n4_clus1	up/up	10.10.0.11/24	n4
e0e	n4_clus2	up/up	10.10.0.12/24	n4

12 entries were displayed.

c. Display the information on the discovered cluster switches:

```
system cluster-switch show
```

Show example

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

Switch                               Type                               Address
Model
-----
CL1                                  cluster-network                    10.10.1.101
NX3132V
  Serial Number: FOX000001
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                               7.0(3)I4(1)
  Version Source: CDP

CL2                                  cluster-network                    10.10.1.102
NX3132V
  Serial Number: FOX000002
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                               7.0(3)I4(1)
  Version Source: CDP

2 entries were displayed.
```

4. Verify that the appropriate RCF and image are installed on the new Nexus 3132Q-V switch as necessary for your requirements, and make any essential site customizations.

You must prepare the replacement switch at this time. If you need to upgrade the RCF and image, you must follow these steps:

- a. On the NetApp Support Site, go to the [Cisco Ethernet Switch](#) page.
 - b. Note your switch and the required software versions in the table on that page.
 - c. Download the appropriate version of the RCF.
 - d. Click **CONTINUE** on the **Description** page, accept the license agreement, and then follow the instructions on the **Download** page to download the RCF.
 - e. Download the appropriate version of the image software.
5. Migrate the LIFs associated to the cluster ports connected to switch C2:

```
network interface migrate
```

Show example

This example shows that the LIF migration is done on all the nodes:

```
cluster::*> network interface migrate -vserver Cluster -lif n1_clus2
-source-node n1 -destination-node n1 -destination-port e0a
cluster::*> network interface migrate -vserver Cluster -lif n1_clus3
-source-node n1 -destination-node n1 -destination-port e0d
cluster::*> network interface migrate -vserver Cluster -lif n2_clus2
-source-node n2 -destination-node n2 -destination-port e0a
cluster::*> network interface migrate -vserver Cluster -lif n2_clus3
-source-node n2 -destination-node n2 -destination-port e0d
cluster::*> network interface migrate -vserver Cluster -lif n3_clus2
-source-node n3 -destination-node n3 -destination-port e4a
cluster::*> network interface migrate -vserver Cluster -lif n4_clus2
-source-node n4 -destination-node n4 -destination-port e4a
```

6. Verify cluster's health:

```
network interface show
```

Show example

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

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

Cluster				
e0a	n1_clus1	up/up	10.10.0.1/24	n1
e0a	true			
e0a	n1_clus2	up/up	10.10.0.2/24	n1
e0a	false			
e0d	n1_clus3	up/up	10.10.0.3/24	n1
e0d	false			
e0d	n1_clus4	up/up	10.10.0.4/24	n1
e0d	true			
e0a	n2_clus1	up/up	10.10.0.5/24	n2
e0a	true			
e0a	n2_clus2	up/up	10.10.0.6/24	n2
e0a	false			
e0d	n2_clus3	up/up	10.10.0.7/24	n2
e0d	false			
e0d	n2_clus4	up/up	10.10.0.8/24	n2
e0d	true			
e4a	n3_clus1	up/up	10.10.0.9/24	n3
e4a	true			
e4a	n3_clus2	up/up	10.10.0.10/24	n3
e4a	false			
e4a	n4_clus1	up/up	10.10.0.11/24	n4
e4a	true			
e4a	n4_clus2	up/up	10.10.0.12/24	n4
e4a	false			

12 entries were displayed.

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

```
network port modify
```

Show example

This example shows the specified ports being shut down on all nodes:

```
cluster::*> network port modify -node n1 -port e0b -up-admin false
cluster::*> network port modify -node n1 -port e0c -up-admin false
cluster::*> network port modify -node n2 -port e0b -up-admin false
cluster::*> network port modify -node n2 -port e0c -up-admin false
cluster::*> network port modify -node n3 -port e4e -up-admin false
cluster::*> network port modify -node n4 -port e4e -up-admin false
```

8. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Node	Date		Source LIF	Destination LIF	Packet Loss
n1					
	3/5/2022 19:21:18 -06:00		n1_clus2	n2_clus1	none
	3/5/2022 19:21:20 -06:00		n1_clus2	n2_clus2	none
n2					
	3/5/2022 19:21:18 -06:00		n2_clus2	n1_clus1	none
	3/5/2022 19:21:20 -06:00		n2_clus2	n1_clus2	none
n3					
...					
...					
n4					
...					
...					

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```
cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1 e0a 10.10.0.1
Cluster n1_clus2 n1 e0b 10.10.0.2
Cluster n1_clus3 n1 e0c 10.10.0.3
Cluster n1_clus4 n1 e0d 10.10.0.4
Cluster n2_clus1 n2 e0a 10.10.0.5
Cluster n2_clus2 n2 e0b 10.10.0.6
```

```
Cluster n2_clus3 n2      e0c 10.10.0.7
Cluster n2_clus4 n2      e0d 10.10.0.8
Cluster n3_clus1 n4      e0a 10.10.0.9
Cluster n3_clus2 n3      e0e 10.10.0.10
Cluster n4_clus1 n4      e0a 10.10.0.11
Cluster n4_clus2 n4      e0e 10.10.0.12

Local = 10.10.0.1 10.10.0.2 10.10.0.3 10.10.0.4
Remote = 10.10.0.5 10.10.0.6 10.10.0.7 10.10.0.8 10.10.0.9 10.10.0.10
10.10.0.11 10.10.0.12
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 32 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 32 path(s):
  Local 10.10.0.1 to Remote 10.10.0.5
  Local 10.10.0.1 to Remote 10.10.0.6
  Local 10.10.0.1 to Remote 10.10.0.7
  Local 10.10.0.1 to Remote 10.10.0.8
  Local 10.10.0.1 to Remote 10.10.0.9
  Local 10.10.0.1 to Remote 10.10.0.10
  Local 10.10.0.1 to Remote 10.10.0.11
  Local 10.10.0.1 to Remote 10.10.0.12
  Local 10.10.0.2 to Remote 10.10.0.5
  Local 10.10.0.2 to Remote 10.10.0.6
  Local 10.10.0.2 to Remote 10.10.0.7
  Local 10.10.0.2 to Remote 10.10.0.8
  Local 10.10.0.2 to Remote 10.10.0.9
  Local 10.10.0.2 to Remote 10.10.0.10
  Local 10.10.0.2 to Remote 10.10.0.11
  Local 10.10.0.2 to Remote 10.10.0.12
  Local 10.10.0.3 to Remote 10.10.0.5
  Local 10.10.0.3 to Remote 10.10.0.6
  Local 10.10.0.3 to Remote 10.10.0.7
  Local 10.10.0.3 to Remote 10.10.0.8
  Local 10.10.0.3 to Remote 10.10.0.9
  Local 10.10.0.3 to Remote 10.10.0.10
  Local 10.10.0.3 to Remote 10.10.0.11
  Local 10.10.0.3 to Remote 10.10.0.12
  Local 10.10.0.4 to Remote 10.10.0.5
  Local 10.10.0.4 to Remote 10.10.0.6
  Local 10.10.0.4 to Remote 10.10.0.7
  Local 10.10.0.4 to Remote 10.10.0.8
  Local 10.10.0.4 to Remote 10.10.0.9
```

```
Local 10.10.0.4 to Remote 10.10.0.10
Local 10.10.0.4 to Remote 10.10.0.11
Local 10.10.0.4 to Remote 10.10.0.12
```

Larger than PMTU communication succeeds on 32 path(s)

RPC status:

8 paths up, 0 paths down (tcp check)

8 paths up, 0 paths down (udp check)

9. Shut down the ports 1/31 and 1/32 on CL1, and the active Nexus 3132Q-V switch:

```
shutdown
```

Show example

This example shows the ISL ports 1/31 and 1/32 being shut down on switch CL1:

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

Step 2: Configure ports

1. Remove all the cables attached to the Nexus 3132Q-V switch CL2 and reconnect them to the replacement switch C2 on all nodes.
2. Remove the ISL cables from ports e1/31 and e1/32 on CL2 and reconnect them to the same ports on the replacement switch C2.
3. Bring up ISLs ports 1/31 and 1/32 on the Nexus 3132Q-V switch CL1:

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

4. Verify that the ISLs are up on CL1:

```
show port-channel
```

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

Show example

```
CL1# show port-channel summary
Flags: D - Down          P - Up in port-channel (members)
      I - Individual    H - Hot-standby (LACP only)
      s - Suspended     r - Module-removed
      S - Switched     R - Routed
      U - Up (port-channel)
      M - Not in use. Min-links not met

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

5. Verify that the ISLs are up on C2:

```
show port-channel summary
```

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

Show example

```
C2# show port-channel summary
Flags: D - Down          P - Up in port-channel (members)
      I - Individual    H - Hot-standby (LACP only)
      s - Suspended     r - Module-removed
      S - Switched     R - Routed
      U - Up (port-channel)
      M - Not in use. Min-links not met

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

6. On all nodes, bring up all the cluster interconnect ports connected to the Nexus 3132Q-V switch C2:

```
network port modify
```

Show example

```
cluster::*> network port modify -node n1 -port e0b -up-admin true
cluster::*> network port modify -node n1 -port e0c -up-admin true
cluster::*> network port modify -node n2 -port e0b -up-admin true
cluster::*> network port modify -node n2 -port e0c -up-admin true
cluster::*> network port modify -node n3 -port e4e -up-admin true
cluster::*> network port modify -node n4 -port e4e -up-admin true
```

7. For all nodes, revert all of the migrated cluster interconnect LIFs:

```
network interface revert
```

Show example

```
cluster::*> network interface revert -vserver Cluster -lif n1_clus2
cluster::*> network interface revert -vserver Cluster -lif n1_clus3
cluster::*> network interface revert -vserver Cluster -lif n2_clus2
cluster::*> network interface revert -vserver Cluster -lif n2_clus3
Cluster::*> network interface revert -vserver Cluster -lif n3_clus2
Cluster::*> network interface revert -vserver Cluster -lif n4_clus2
```

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

```
network interface show
```

Show example

This example shows that all the LIFs are successfully reverted because the ports listed under the Current Port column have a status of true in the Is Home column. If the Is Home column value is false, the LIF has not been reverted.

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

```

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

Cluster					
e0a	true	n1_clus1	up/up	10.10.0.1/24	n1
e0b	true	n1_clus2	up/up	10.10.0.2/24	n1
e0c	true	n1_clus3	up/up	10.10.0.3/24	n1
e0d	true	n1_clus4	up/up	10.10.0.4/24	n1
e0a	true	n2_clus1	up/up	10.10.0.5/24	n2
e0b	true	n2_clus2	up/up	10.10.0.6/24	n2
e0c	true	n2_clus3	up/up	10.10.0.7/24	n2
e0d	true	n2_clus4	up/up	10.10.0.8/24	n2
e4a	true	n3_clus1	up/up	10.10.0.9/24	n3
e4e	true	n3_clus2	up/up	10.10.0.10/24	n3
e4a	true	n4_clus1	up/up	10.10.0.11/24	n4
e4e	true	n4_clus2	up/up	10.10.0.12/24	n4

12 entries were displayed.

9. Verify that the cluster ports are connected:

```
network port show
```

Show example

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

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e0a      Cluster      Cluster      up    9000  auto/10000  -
-
e0b      Cluster      Cluster      up    9000  auto/10000  -
-
e0c      Cluster      Cluster      up    9000  auto/10000  -
-
e0d      Cluster      Cluster      up    9000  auto/10000  -
-

Node: n2

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e0a      Cluster      Cluster      up    9000  auto/10000  -
-
e0b      Cluster      Cluster      up    9000  auto/10000  -
-
e0c      Cluster      Cluster      up    9000  auto/10000  -
-
e0d      Cluster      Cluster      up    9000  auto/10000  -
-

Node: n3

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
```

```

Status
-----
e4a      Cluster      Cluster      up    9000 auto/40000 -
-
e4e      Cluster      Cluster      up    9000 auto/40000 -
-

Node: n4

Ignore

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

12 entries were displayed.

```

10. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Node	Date		Source LIF	Destination LIF	Packet Loss
n1					
	3/5/2022 19:21:18 -06:00		n1_clus2	n2_clus1	none
	3/5/2022 19:21:20 -06:00		n1_clus2	n2_clus2	none
n2					
	3/5/2022 19:21:18 -06:00		n2_clus2	n1_clus1	none
	3/5/2022 19:21:20 -06:00		n2_clus2	n1_clus2	none
n3					
...					
...					
n4					
...					
...					

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```
cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1 e0a 10.10.0.1
Cluster n1_clus2 n1 e0b 10.10.0.2
Cluster n2_clus1 n2 e0a 10.10.0.5
Cluster n2_clus2 n2 e0b 10.10.0.6
Cluster n2_clus3 n2 e0c 10.10.0.7
Cluster n2_clus4 n2 e0d 10.10.0.8
```

```
Cluster n3_clus1 n3      e0a 10.10.0.9
Cluster n3_clus2 n3      e0e 10.10.0.10
Cluster n4_clus1 n4      e0a 10.10.0.11
Cluster n4_clus2 n4      e0e 10.10.0.12

Local = 10.10.0.1 10.10.0.2 10.10.0.3 10.10.0.4
Remote = 10.10.0.5 10.10.0.6 10.10.0.7 10.10.0.8 10.10.0.9 10.10.0.10
10.10.0.11 10.10.0.12
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 32 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 32 path(s):
  Local 10.10.0.1 to Remote 10.10.0.5
  Local 10.10.0.1 to Remote 10.10.0.6
  Local 10.10.0.1 to Remote 10.10.0.7
  Local 10.10.0.1 to Remote 10.10.0.8
  Local 10.10.0.1 to Remote 10.10.0.9
  Local 10.10.0.1 to Remote 10.10.0.10
  Local 10.10.0.1 to Remote 10.10.0.11
  Local 10.10.0.1 to Remote 10.10.0.12
  Local 10.10.0.2 to Remote 10.10.0.5
  Local 10.10.0.2 to Remote 10.10.0.6
  Local 10.10.0.2 to Remote 10.10.0.7
  Local 10.10.0.2 to Remote 10.10.0.8
  Local 10.10.0.2 to Remote 10.10.0.9
  Local 10.10.0.2 to Remote 10.10.0.10
  Local 10.10.0.2 to Remote 10.10.0.11
  Local 10.10.0.2 to Remote 10.10.0.12
  Local 10.10.0.3 to Remote 10.10.0.5
  Local 10.10.0.3 to Remote 10.10.0.6
  Local 10.10.0.3 to Remote 10.10.0.7
  Local 10.10.0.3 to Remote 10.10.0.8
  Local 10.10.0.3 to Remote 10.10.0.9
  Local 10.10.0.3 to Remote 10.10.0.10
  Local 10.10.0.3 to Remote 10.10.0.11
  Local 10.10.0.3 to Remote 10.10.0.12
  Local 10.10.0.4 to Remote 10.10.0.5
  Local 10.10.0.4 to Remote 10.10.0.6
  Local 10.10.0.4 to Remote 10.10.0.7
  Local 10.10.0.4 to Remote 10.10.0.8
  Local 10.10.0.4 to Remote 10.10.0.9
  Local 10.10.0.4 to Remote 10.10.0.10
  Local 10.10.0.4 to Remote 10.10.0.11
```

```
Local 10.10.0.4 to Remote 10.10.0.12
```

```
Larger than PMTU communication succeeds on 32 path(s)
```

```
RPC status:
```

```
8 paths up, 0 paths down (tcp check)
```

```
8 paths up, 0 paths down (udp check)
```

Step 3: Verify the configuration

1. Display the information about the devices in your configuration:

- network device-discovery show
- network port show -role cluster
- network interface show -role cluster
- system cluster-switch show

Show example

```
cluster::> network device-discovery show
      Local   Discovered
Node   Port     Device           Interface         Platform
-----
n1     /cdp
      e0a     C1               Ethernet1/1/1     N3K-C3132Q-V
      e0b     C2               Ethernet1/1/1     N3K-C3132Q-V
      e0c     C2               Ethernet1/1/2     N3K-C3132Q-V
      e0d     C1               Ethernet1/1/2     N3K-C3132Q-V
n2     /cdp
      e0a     C1               Ethernet1/1/3     N3K-C3132Q-V
      e0b     C2               Ethernet1/1/3     N3K-C3132Q-V
      e0c     C2               Ethernet1/1/4     N3K-C3132Q-V
      e0d     C1               Ethernet1/1/4     N3K-C3132Q-V
n3     /cdp
      e4a     C1               Ethernet1/7       N3K-C3132Q-V
      e4e     C2               Ethernet1/7       N3K-C3132Q-V
n4     /cdp
      e4a     C1               Ethernet1/8       N3K-C3132Q-V
      e4e     C2               Ethernet1/8       N3K-C3132Q-V
12 entries were displayed.
```

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

Ignore

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

Node: n2

Ignore

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

Node: n3

Ignore

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

Node: n4

Ignore

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

12 entries were displayed.

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

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

Cluster				
e0a	n1_clus1	up/up	10.10.0.1/24	n1
e0b	n1_clus2	up/up	10.10.0.2/24	n1
e0c	n1_clus3	up/up	10.10.0.3/24	n1
e0d	n1_clus4	up/up	10.10.0.4/24	n1
e0a	n2_clus1	up/up	10.10.0.5/24	n2
e0b	n2_clus2	up/up	10.10.0.6/24	n2
e0c	n2_clus3	up/up	10.10.0.7/24	n2
e0d	n2_clus4	up/up	10.10.0.8/24	n2
e4a	n3_clus1	up/up	10.10.0.9/24	n3
e4e	n3_clus2	up/up	10.10.0.10/24	n3
e4a	n4_clus1	up/up	10.10.0.11/24	n4
e4e	n4_clus2	up/up	10.10.0.12/24	n4

12 entries were displayed.

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

```
Switch                               Type                               Address
Model
-----
CL1                                   cluster-network                   10.10.1.101
NX3132V
  Serial Number: FOX000001
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        7.0(3)I4(1)
  Version Source: CDP

CL2                                   cluster-network                   10.10.1.102
NX3132V
  Serial Number: FOX000002
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        7.0(3)I4(1)
  Version Source: CDP

C2                                   cluster-network                   10.10.1.103
NX3132V
  Serial Number: FOX000003
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        7.0(3)I4(1)
  Version Source: CDP

3 entries were displayed.
```

2. Remove the replaced Nexus 3132Q-V switch, if it is not already removed automatically:

```
system cluster-switch delete
```

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

3. Verify that the proper cluster switches are monitored:

```
system cluster-switch show
```

Show example

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

Switch                                Type                                Address
Model
-----
CL1                                    cluster-network                    10.10.1.101
NX3132V
  Serial Number: FOX000001
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        7.0(3)I4(1)
  Version Source: CDP

C2                                    cluster-network                    10.10.1.103
NX3132V
  Serial Number: FOX000002
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        7.0(3)I4(1)
  Version Source: CDP

2 entries were displayed.
```

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

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

What's next?

[Configure switch health monitoring](#)

Replace Cisco Nexus 3132Q-V cluster switches with switchless connections

You can migrate from a cluster with a switched cluster network to one where two nodes

are directly connected for ONTAP 9.3 and later.

Review requirements

Guidelines

Review the following guidelines:

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


What you'll need

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

Migrate the switches

About this task

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



About the examples

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

Step 1: Prepare for migration

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

```
set -privilege advanced
```

The advanced prompt `*>` appears.

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

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

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

Show example

The following example output shows if the option is enabled.

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

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

3. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message
MAINT=<number_of_hours>h
```

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

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

Show example

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


Step 2: Configure ports and cabling

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

```
network port show -ipSpace Cluster
```

In the following example for nodes with cluster ports "e0a" and "e0b", one group is identified as "node1:e0a" and "node2:e0a" and the other group as "node1:e0b" and "node2:e0b". Your nodes might be

using different cluster ports because they vary by system.



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

Show example

```

cluster::> network port show -ipspace Cluster
Node: nodel

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

Node: node2

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

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

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

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

Show example

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

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

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

4. Disable auto-revert for the cluster LIFs:

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

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

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

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

Show example

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

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

6. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

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

```

7. Verify that the cluster is healthy:

```
cluster ring show
```

All units must be either master or secondary.


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



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


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

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



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

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



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

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

The following example shows that the switchless cluster is enabled:

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

10. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

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

```



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


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



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


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

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



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

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



Step 3: Verify the configuration

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

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

Show example

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

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

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

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

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

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

Show example

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

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

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

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

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

```
cluster show
```

Show example

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

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

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

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

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

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

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

7. Change the privilege level back to admin:

```
set -privilege admin
```

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

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