



Cisco® Nexus Switches

Migrating to a Two-Node Switched Cluster

For Cisco Nexus 3132Q-V switches

September 2020 | 215-11556_CO
doccomments@netapp.com

Contents

- Migrating to a two-node switched cluster with Cisco Nexus 3132Q-V cluster switches..... 3**
 - How to migrate a two-node switched cluster with Cisco Nexus 3132Q-V cluster switches.....4

- Copyright and trademark..... 14**
 - Copyright..... 14
 - Trademark..... 14

- Index..... 15**

Migrating to a two-node switched cluster with Cisco Nexus 3132Q-V cluster switches

You must be aware of certain configuration information, port connections and cabling requirements when you migrate to a two-node switched cluster with Cisco Nexus 3132Q-V cluster switches.

- The number of 10 GbE and 40/100 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.
- The 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 40/100 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 40/100 Gb Ethernet or 4 x10 Gb Ethernet modes.

By default, there are 32 ports in the 40/100 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 40/100 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 40/100 Gb Ethernet port are numbered as 1/2/1, 1/2/2, 1/2/3, 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 40/100 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 40/100 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 Ethernet Switches](#)

How to migrate a two-node switched cluster with Cisco Nexus 3132Q-V cluster switches

If you have a two-node switchless cluster, you can migrate nondisruptively to a two-node switched cluster that includes Cisco Nexus 3132Q-V cluster network switches.

Before you begin

- The configurations must be properly set up and functioning.
- The nodes must be running ONTAP 9.4 or later.
- All cluster ports must be in the **up** state.
- The Cisco Nexus 3132Q-V cluster switch must be supported.
- The existing cluster network configuration must have:
 - The Nexus 3132 cluster infrastructure that is redundant and fully functional on both switches.
The latest RCF and NX-OS versions on your switches.
 - Management connectivity on both switches.
 - Console access to both switches.
 - All cluster logical interfaces (LIFs) in the **up** state without being migrated.
 - Initial customization of the switch.
 - All the ISL ports enabled and cabled.

About this task

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

- Nexus 3132Q-V cluster switches, C1 and C2.
- The nodes are n1 and n2.

Note: The examples in this procedure use two nodes, each utilizing two 40/100 GbE cluster interconnect ports e4a and e4e. The *Hardware Universe* has details about the cluster ports on your platforms.

- n1_clus1 is the first cluster logical interface (LIF) to be connected to cluster switch C1 for node n1.
- n1_clus2 is the first cluster LIF to be connected to cluster switch C2 for node n1.
- n2_clus1 is the first cluster LIF to be connected to cluster switch C1 for node n2.
- n2_clus2 is the second cluster LIF to be connected to cluster switch C2 for node n2.
- The number of 10 GbE and 40/100 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.

Note: The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

- The cluster starts with two nodes connected and functioning in a two-node switchless cluster setting.
- The first cluster port moved to C1 (steps 1 to 20).
- The second cluster port moved to C2 (steps 21 to 32).
- Disable the two-node switchless cluster option (steps 33 to 35).

Steps

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.

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

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

a. Display the network port attributes:

network port show

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

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

```
Node: n2
```

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

4 entries were displayed.

b. Display information about the logical interfaces:

network interface show

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

| Vserver | Logical Interface | Status Admin/Oper | Network Address/Mask | Current Node | Current Port | Is Home |
|---------|-------------------|----------------------|-------------------------|-----------------|-----------------|------------|
| Cluster | n1_clus1 | up/up | 10.10.0.1/24 | n1 | e4a | true |
| | n1_clus2 | up/up | 10.10.0.2/24 | n1 | e4e | true |
| | n2_clus1 | up/up | 10.10.0.3/24 | n2 | e4a | true |
| | n2_clus2 | up/up | 10.10.0.4/24 | n2 | e4e | true |

4 entries were displayed.

3. Verify that the appropriate RCFs and image are installed on the new 3132Q-V switches as necessary for your requirements, and make any essential site customizations, such as users and passwords, network addresses, and so on.

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

- a. Go to the *Cisco Ethernet Switches* page on the NetApp Support Site.
[Cisco Ethernet Switches](#)
- b. Note your switch and the required software versions in the table on that page.
- c. Download the appropriate version of 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.

4. Click **CONTINUE** on the **Description** page, accept the license agreement, and then follow the instructions on the **Download** page to download the RCF.

Migrating to a two-node switched cluster with Cisco Nexus 3132Q-V cluster switches

- On Nexus 3132Q-V switches C1 and C2, disable all node-facing ports C1 and C2, but do not disable the ISL ports.

The following example shows ports 1 through 30 being disabled on Nexus 3132Q-V cluster switches C1 and C2 using a configuration supported in RCF

NX3132_RCF_v1.1_24p10g_26p40g.txt:

```
C1# copy running-config startup-config
[#####] 100%
Copy complete.
C1# configure
C1(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C1(config-if-range)# shutdown
C1(config-if-range)# exit
C1(config)# exit

C2# copy running-config startup-config
[#####] 100%
Copy complete.
C2# configure
C2(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C2(config-if-range)# shutdown
C2(config-if-range)# exit
C2(config)# exit
```

- Connect ports 1/31 and 1/32 on C1 to the same ports on C2 using supported cabling.
- Verify that the ISL ports are operational on C1 and C2:

show port-channel summary

```
C1# 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     Pol(SU)        Eth    LACP      Eth1/31(P)  Eth1/32(P)
```

```
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     Pol(SU)        Eth    LACP      Eth1/31(P)  Eth1/32(P)
```

- Display the list of neighboring devices on the switch:

show cdp neighbors

```
C1# show cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID          Local Intrfce  Hldtme Capability  Platform      Port ID
C2                  Eth1/31        174    R S I s      N3K-C3132Q-V  Eth1/31
```

Migrating to a two-node switched cluster with Cisco Nexus 3132Q-V cluster switches

```
C2          Eth1/32          174    R S I s    N3K-C3132Q-V  Eth1/32

Total entries displayed: 2

C2# show cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID      Local Intrfce  Hldtme  Capability  Platform      Port ID
C1              Eth1/31       178     R S I s    N3K-C3132Q-V  Eth1/31
C1              Eth1/32       178     R S I s    N3K-C3132Q-V  Eth1/32

Total entries displayed: 2
```

9. Display the cluster port connectivity on each node:

network device-discovery show

The following example shows a two-node switchless cluster configuration.

```
cluster::*> network device-discovery show
Node      Local Port   Discovered Device      Interface      Platform
-----
n1        /cdp
          e4a      n2          e4a          FAS9000
          e4e      n2          e4e          FAS9000
n2        /cdp
          e4a      n1          e4a          FAS9000
          e4e      n1          e4e          FAS9000
```

10. Migrate the clus1 interface to the physical port hosting clus2:

network interface migrate

Execute this command from each local node.

```
cluster::*> network interface migrate -vserver Cluster -lif n1_clus1 -source-node n1
             -destination-node n1 -destination-port e4e
cluster::*> network interface migrate -vserver Cluster -lif n2_clus1 -source-node n2
             -destination-node n2 -destination-port e4e
```

11. Verify the cluster interfaces migration:

network interface show

```
cluster::*> network interface show -role cluster
(network interface show)
Vserver   Logical Interface  Status Admin/Oper  Network Address/Mask  Current Node  Current Port  Is Home
-----
Cluster
          n1_clus1  up/up    10.10.0.1/24  n1          e4e          false
          n1_clus2  up/up    10.10.0.2/24  n1          e4e          true
          n2_clus1  up/up    10.10.0.3/24  n2          e4e          false
          n2_clus2  up/up    10.10.0.4/24  n2          e4e          true
4 entries were displayed.
```

12. Shut down cluster ports clus1 LIF on both nodes:

network port modify

```
cluster::*> network port modify -node n1 -port e4a -up-admin false
cluster::*> network port modify -node n2 -port e4a -up-admin false
```

13. Ping the remote cluster interfaces and perform an RPC server check:

cluster ping-cluster

```

cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1      e4a    10.10.0.1
Cluster n1_clus2 n1      e4e    10.10.0.2
Cluster n2_clus1 n2      e4a    10.10.0.3
Cluster n2_clus2 n2      e4e    10.10.0.4

Local = 10.10.0.1 10.10.0.2
Remote = 10.10.0.3 10.10.0.4
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 32 path(s):
    Local 10.10.0.1 to Remote 10.10.0.3
    Local 10.10.0.1 to Remote 10.10.0.4
    Local 10.10.0.2 to Remote 10.10.0.3
    Local 10.10.0.2 to Remote 10.10.0.4
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
1 paths up, 0 paths down (tcp check)
1 paths up, 0 paths down (ucp check)

```

14. Disconnect the cable from e4a on node n1.

You can refer to the running configuration and connect the first 40 GbE port on the switch C1 (port 1/7 in this example) to e4a on n1 using supported cabling on Nexus 3132Q-V.



Attention: When reconnecting any cables to a new Cisco cluster switch, the cables used must be either fiber or cabling supported by Cisco.

15. Disconnect the cable from e4a on node n2.

You can refer to the running configuration and connect e4a to the next available 40 GbE port on C1, port 1/8, using supported cabling.

16. Enable all node-facing ports on C1.

The following example shows ports 1 through 30 being enabled on Nexus 3132Q-V cluster switches C1 and C2 using the configuration supported in RCF

NX3132_RCF_v1.1_24p10g_26p40g.txt:

```

C1# configure
C1(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C1(config-if-range)# no shutdown
C1(config-if-range)# exit
C1(config)# exit

```

17. Enable the first cluster port, e4a, on each node:

network port modify

```

cluster::*> network port modify -node n1 -port e4a -up-admin true
cluster::*> network port modify -node n2 -port e4a -up-admin true

```

18. Verify that the clusters are up on both nodes:

network port show

```

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

```

Ignore

Migrating to a two-node switched cluster with Cisco Nexus 3132Q-V cluster switches

```

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

Node: n2

Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health Health
-----
e4a      Cluster     Cluster         up  9000    auto/40000 -      -
e4e      Cluster     Cluster         up  9000    auto/40000 -      -
4 entries were displayed.
    
```

19. For each node, revert all of the migrated cluster interconnect LIFs:

network interface revert

The following example shows the migrated LIFs being reverted to their home ports.

```

cluster::*> network interface revert -vserver Cluster -lif n1_clus1
cluster::*> network interface revert -vserver Cluster -lif n2_clus1
    
```

20. Verify that all of the cluster interconnect ports are now reverted to their home ports:

network interface show

The Is Home column should display a value of **true** for all of the ports listed in the Current Port column. If the displayed value is **false**, the port has not been reverted.

```

cluster::*> network interface show -role cluster
(network interface show)
Vserver      Logical      Status      Network      Current      Current      Is
Interface    Admin/Oper   Address/Mask Node          Port         Home
-----
Cluster
n1_clus1     up/up       10.10.0.1/24 n1           e4a         true
n1_clus2     up/up       10.10.0.2/24 n1           e4e         true
n2_clus1     up/up       10.10.0.3/24 n2           e4a         true
n2_clus2     up/up       10.10.0.4/24 n2           e4e         true
4 entries were displayed.
    
```

21. Display the cluster port connectivity on each node:

network device-discovery show

```

cluster::*> network device-discovery show
Local      Discovered
Node       Port      Device      Interface      Platform
-----
n1         /cdp
e4a       C1         Ethernet1/7   N3K-C3132Q-V
e4e       n2         e4e           FAS9000
n2         /cdp
e4a       C1         Ethernet1/8   N3K-C3132Q-V
e4e       n1         e4e           FAS9000
    
```

22. On the console of each node, migrate clus2 to port e4a:

network interface migrate

```

cluster::*> network interface migrate -vserver Cluster -lif n1_clus2 -source-node n1
-destination-node n1 -destination-port e4a
cluster::*> network interface migrate -vserver Cluster -lif n2_clus2 -source-node n2
-destination-node n2 -destination-port e4a
    
```

23. Shut down cluster ports clus2 LIF on both nodes:

network port modify

The following example shows the specified ports being shut down on both nodes:

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

24. Verify the cluster LIF status:

network interface show

```
cluster::*> network interface show -role cluster
(network interface show)
Vserver      Logical      Status      Network      Current      Current      Is
-----      -
Cluster
n1_clus1     up/up       10.10.0.1/24  n1           e4a          true
n1_clus2     up/up       10.10.0.2/24  n1           e4a          false
n2_clus1     up/up       10.10.0.3/24  n2           e4a          true
n2_clus2     up/up       10.10.0.4/24  n2           e4a          false
4 entries were displayed.
```

25. Disconnect the cable from e4e on node n1.

You can refer to the running configuration and connect the first 40 GbE port on the switch C2 (port 1/7 in this example) to e4e on n1 using supported cabling on Nexus 3132Q-V.

26. Disconnect the cable from e4e on node n2.

You can refer to the running configuration and connect e4e to the next available 40 GbE port on C2, port 1/8, using supported cabling.

27. Enable all node-facing ports on C2.

The following example shows ports 1 through 30 being enabled on Nexus 3132Q-V cluster switches C1 and C2 using a configuration supported in RCF

NX3132_RCF_v1.1_24p10g_26p40g.txt:

```
C2# configure
C2(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C2(config-if-range)# no shutdown
C2(config-if-range)# exit
C2(config)# exit
```

28. Enable the second cluster port, e4e, on each node:

network port modify

The following example shows the specified ports being brought up:

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

29. For each node, revert all of the migrated cluster interconnect LIFs:

network interface revert

The following example shows the migrated LIFs being reverted to their home ports.

```
cluster::*> network interface revert -vserver Cluster -lif n1_clus2
cluster::*> network interface revert -vserver Cluster -lif n2_clus2
```

30. Verify that all of the cluster interconnect ports are now reverted to their home ports:

network interface show

The Is Home column should display a value of **true** for all of the ports listed in the Current Port column. If the displayed value is **false**, the port has not been reverted.

```
cluster::*> network interface show -role cluster
(network interface show)
Vserver      Logical      Status      Network      Current      Current      Is
-----      -
Cluster
n1_clus1     up/up       10.10.0.1/24  n1           e4a          true
n1_clus2     up/up       10.10.0.2/24  n1           e4e          true
```

Migrating to a two-node switched cluster with Cisco Nexus 3132Q-V cluster switches

```

n2_clus1 up/up 10.10.0.3/24 n2 e4a true
n2_clus2 up/up 10.10.0.4/24 n2 e4e true
4 entries were displayed.

```

- 31. Verify that all of the cluster interconnect ports are in the **up** state.

```

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

```

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

```

Node: n2

```

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

```

4 entries were displayed.

```

- 32. Display the cluster switch port numbers each cluster port is connected to on each node:

network device-discovery show

```

cluster::*> network device-discovery show

```

| Node | Local Port | Discovered Device | Interface | Platform |
|------|------------|-------------------|-------------|--------------|
| n1 | /cdp | | | |
| | e4a | C1 | Ethernet1/7 | N3K-C3132Q-V |
| n2 | e4e | C2 | Ethernet1/7 | N3K-C3132Q-V |
| | /cdp | | | |
| n1 | e4a | C1 | Ethernet1/8 | N3K-C3132Q-V |
| | e4e | C2 | Ethernet1/8 | N3K-C3132Q-V |

- 33. Display discovered and monitored cluster switches:

system cluster-switch show

```

cluster::*> system cluster-switch show

```

| Switch | Type | Address | Model |
|--|-----------------|-------------|---------|
| C1 | 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 | 10.10.1.102 | 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.

```

- 34. Disable the two-node switchless configuration settings on any node:

network options switchless-cluster

```
network options switchless-cluster modify -enabled false
```

35. Verify that the switchless-cluster option has been disabled.

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

36. Ping the remote cluster interfaces and perform an RPC server check:

```
cluster ping-cluster
```

```
cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1      e4a  10.10.0.1
Cluster n1_clus2 n1      e4e  10.10.0.2
Cluster n2_clus1 n2      e4a  10.10.0.3
Cluster n2_clus2 n2      e4e  10.10.0.4

Local = 10.10.0.1 10.10.0.2
Remote = 10.10.0.3 10.10.0.4
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 32 path(s):
  Local 10.10.0.1 to Remote 10.10.0.3
  Local 10.10.0.1 to Remote 10.10.0.4
  Local 10.10.0.2 to Remote 10.10.0.3
  Local 10.10.0.2 to Remote 10.10.0.4
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
1 paths up, 0 paths down (tcp check)
1 paths up, 0 paths down (ucp check)
```

37. Enable the cluster switch health monitor log collection feature for collecting switch-related log files:

```
system cluster-switch log setup-password
```

```
system cluster-switch log enable-collection
```

```
cluster::*> system cluster-switch log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
C1
C2

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

Enter the switch name: C1
RSA key fingerprint is e5:8b:c6:dc:e2:18:18:09:36:63:d9:63:dd:03:d9:cc
Do you want to continue? {y|n}::[n] y

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

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

Enter the switch name: C2
RSA key fingerprint is 57:49:86:a1:b9:80:6a:61:9a:86:8e:3c:e3:b7:1f:b1
Do you want to continue? {y|n}:: [n] y

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

cluster::*> system cluster-switch log enable-collection
```

```
Do you want to enable cluster log collection for all nodes in the cluster?  
{y|n}: [n] y  
  
Enabling cluster switch log collection.  
  
cluster::*>
```

Note: If any of these commands return an error, contact NetApp support.

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

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

Copyright and trademark

Copyright

Copyright © 2020 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>

Index

C

- Cisco Nexus 5010 switches
 - migrating from a switchless to a switched environment [3, 4](#)
- Cisco Nexus 5020 switches
 - migrating from a switchless to a switched environment [3, 4](#)
- Cisco Nexus 5596 switches
 - migrating from a switchless to a switched environment [3, 4](#)

T

- two-node switchless clusters
 - migrating to a two-node switched Cisco Nexus 5010 environment [3, 4](#)
 - migrating to a two-node switched Cisco Nexus 5020 environment [3, 4](#)
 - migrating to a two-node switched Cisco Nexus 5596 environment [3, 4](#)