



Cisco® Nexus Switches

Replacing a Cisco Nexus 3232C Cluster Switch

September 2020 | 215-13128_BO
doccomments@netapp.com

Contents

Replacing a Cisco Nexus 3232C cluster switch.....	3
How to replace a Cisco Nexus 3232C cluster switch.....	3
Copyright and trademark.....	14
Copyright.....	14
Trademark.....	14
Index.....	15

Replacing a Cisco Nexus 3232C cluster switch

You must be aware of certain configuration information, port connections and cabling requirements when you replace Cisco Nexus 3232C cluster switches.

You must verify the following conditions exist before installing the NX-OS software and RCFs on a Cisco Nexus cluster switch:

- Your system can support Cisco Nexus 3232C switches.
- The cluster must be fully functioning.
- You must have consulted the switch compatibility table on the Cisco Ethernet Switch page for the supported ONTAP, NX-OS, and RCF versions.

Note: You should be aware there can be dependencies between command syntax in the RCF and NX-OS versions.

- You must have referred to the appropriate software and upgrade guides available on the Cisco web site for complete documentation on the Cisco switch upgrade and downgrade procedures.
- You must have downloaded the applicable RCFs.

How to replace a Cisco Nexus 3232C cluster switch

You can nondisruptively replace a defective Cisco Nexus 3232C switch in a cluster by performing a specific sequence of tasks.

Before you begin

The existing cluster and network configuration must have the following characteristics:

- The Nexus 3232C cluster infrastructure must be redundant and fully functional on both switches.
The Cisco Ethernet Switches page has the latest RCF and NX-OS versions on your switches.
- All cluster ports must be in the up state.
- Management connectivity must exist on both switches.
- All cluster logical interfaces (LIFs) must be in the up state and must not have been migrated.

The replacement Cisco Nexus 3232C switch must have the following characteristics:

- Management network connectivity must be functional.
- Console access to the replacement switch must be in place.
- The appropriate RCF and NX-OS operating system image must be loaded onto the switch.
- Initial customization of the switch must be complete.

About this task

Procedure summary

- **Display and migrate the cluster ports to switch C2** ([Steps 1-7](#))
- **Reconnect ISL cables from switch CL2 to switch C2, then migrate ISLs to switch CL1 and C2** ([Steps 8-14](#))
- **Revert all LIFs to originally assigned ports** ([Steps 15-18](#))
- **Verify all ports and LIF are correctly migrated** ([Steps 19-21](#))

This procedure replaces the second Nexus 3232C cluster switch CL2 with the new 3232C switch C2. The examples in this procedure use the following switch and node nomenclature:

- The four nodes are n1, n2, n3, and n4.
- 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 40/100 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.

The examples in this procedure use four nodes. Two of the 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 to verify the correct cluster ports for your platform.

[Hardware Universe](#)

This procedure describes the following scenario:

- The cluster initially has four nodes connected to two Nexus 3232C cluster switches, CL1 and CL2.
- You plan to replace cluster switch CL2 with C2 (steps 1 to 21):
 - On each node, you migrate the cluster LIFs connected to cluster switch CL2 to cluster ports connected to cluster switch CL1.
 - You disconnect the cabling from all ports on cluster switch CL2 and reconnect the cabling to the same ports on the replacement cluster switch C2.
 - You revert the migrated cluster LIFs on each node.

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. Display information about the devices in your configuration:

```
network device-discovery show
```

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

Node	Local Port	Discovered Device	Interface	Platform
n1	/cdp			
	e0a	CL1	Ethernet1/1/1	N3K-C3232C
	e0b	CL2	Ethernet1/1/1	N3K-C3232C
	e0c	CL2	Ethernet1/1/2	N3K-C3232C
n2	e0d	CL1	Ethernet1/1/2	N3K-C3232C
	/cdp			
	e0a	CL1	Ethernet1/1/3	N3K-C3232C
	e0b	CL2	Ethernet1/1/3	N3K-C3232C
n3	e0c	CL2	Ethernet1/1/4	N3K-C3232C
	e0d	CL1	Ethernet1/1/4	N3K-C3232C
	/cdp			
	e4a	CL1	Ethernet1/7	N3K-C3232C
n4	e4e	CL2	Ethernet1/7	N3K-C3232C
	/cdp			
n4	e4a	CL1	Ethernet1/8	N3K-C3232C
	e4e	CL2	Ethernet1/8	N3K-C3232C

12 entries were displayed

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

a. Display the network port attributes:

network port show -role cluster

```
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
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
```

Port	IPspace	Broadcast Domain	Link	MTU	Speed(Mbps) Admin/Oper	Health Status	Ignore Health 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
```

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: n4
```

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	-	

12 entries were displayed.

b. Display information about the logical interfaces (LIFs):

network interface show -role cluster

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

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

12 entries were displayed.

c. Display the discovered cluster switches:

system cluster-switch show

The following output example displays the cluster switches:

```
cluster::> system cluster-switch show
Switch                Type                Address            Model
-----
CL1                   cluster-network    10.10.1.101       NX3232C
Serial Number: FOX000001
Is Monitored: true
Reason:
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version
                  7.0(3)I6(1)
Version Source: CDP

CL2                   cluster-network    10.10.1.102       NX3232C
Serial Number: FOX000002
Is Monitored: true
Reason:
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version
                  7.0(3)I6(1)
Version Source: CDP

2 entries were displayed.
```

4. Verify that the appropriate RCF and image are installed on the new Nexus 3232C switch and make any necessary site customizations.
 - a. Go to the NetApp Support Site.
mysupport.netapp.com
 - b. Go to the **Cisco Ethernet Switches** page and note the required software versions in the table.
[Cisco Ethernet Switches](#)
 - c. Download the appropriate version of the RCF.
 - d. Click **CONTINUE** on the **Description** page, accept the license agreement, and then navigate to the **Download** page.
 - e. Download the correct version of the image software from the **Cisco® Cluster and Management Network Switch Reference Configuration File Download** page.
[Cisco® Cluster and Management Network Switch Reference Configuration File Download](#)

5. Migrate the cluster LIFs to the physical node ports connected to the replacement switch C2:

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

You must migrate all the cluster LIFs individually as shown in the following example:

```
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 -destinationnode
n4 -destination-port e4a
```

6. Verify the status of the cluster ports and their home designations:

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

```
cluster::*> network interface show -role cluster
(network interface show)
Vserver      Logical   Status   Network   Current   Current Is
-----
Cluster      Interface Admin/Oper Address/Mask Node       Port     Home
-----
n1_clus1    up/up    10.10.0.1/24    n1        e0a      true
n1_clus2    up/up    10.10.0.2/24    n1        e0a      false
```

```
n1_clus3 up/up 10.10.0.3/24 n1 e0d false
n1_clus4 up/up 10.10.0.4/24 n1 e0d true
n2_clus1 up/up 10.10.0.5/24 n2 e0a true
n2_clus2 up/up 10.10.0.6/24 n2 e0a false
n2_clus3 up/up 10.10.0.7/24 n2 e0d false
n2_clus4 up/up 10.10.0.8/24 n2 e0d true
n3_clus1 up/up 10.10.0.9/24 n3 e4a true
n3_clus2 up/up 10.10.0.10/24 n3 e4a false
n4_clus1 up/up 10.10.0.11/24 n4 e4a true
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 the original switch CL2:

network port modify -node node-name -port port-name -up-admin false

The following example shows the cluster interconnect ports are 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. Ping the remote cluster interfaces and perform an RPC server check:

cluster ping-cluster -node node-name

The following example shows node n1 being pinged and the RPC status indicated afterward:

```
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 9000 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 cluster switch CL1.

For more information on Cisco commands, see the guides listed in the [Cisco Nexus 3000 Series NX-OS Command References](#).

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

10. Remove all the cables attached to the cluster switch CL2 and reconnect them to the replacement switch C2 for all the nodes.
11. Remove the inter-switch link (ISL) cables from ports e1/31 and e1/32 on cluster switch CL2 and reconnect them to the same ports on the replacement switch C2.
12. Bring up ISL ports 1/31 and 1/32 on the cluster switch CL1.

For more information on Cisco commands, see the guides listed in the [Cisco Nexus 3000 Series NX-OS Command References](#).

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

13. Verify that the ISLs are up on CL1.

For more information on Cisco commands, see the guides listed in the [Cisco Nexus 3000 Series NX-OS Command References](#).

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

```
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)
```

14. Verify that the ISLs are up on cluster switch C2.

For more information on Cisco commands, see the guides listed in the [Cisco Nexus 3000 Series NX-OS Command References](#).

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

15. On all nodes, bring up all the cluster interconnect ports connected to the replacement switch C2:

network port modify -node node-name -port port-name -up-admin true

```

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
  
```

16. Revert all the migrated cluster interconnect LIFs on all the nodes:

network interface revert -vserver cluster -lif lif-name

You must revert all the cluster interconnect LIFs individually as shown in the following 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
  
```

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

network interface show

The following example shows that all the LIFs have been successfully reverted because the ports listed under the **Current Port** column have a status of **true** in the **Is Home** column. If a port has a value of **false**, the LIF 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           e0a         true
n1_clus2     up/up        10.10.0.2/24 n1           e0b         true
n1_clus3     up/up        10.10.0.3/24 n1           e0c         true
n1_clus4     up/up        10.10.0.4/24 n1           e0d         true
n2_clus1     up/up        10.10.0.5/24 n2           e0a         true
n2_clus2     up/up        10.10.0.6/24 n2           e0b         true
n2_clus3     up/up        10.10.0.7/24 n2           e0c         true
n2_clus4     up/up        10.10.0.8/24 n2           e0d         true
n3_clus1     up/up        10.10.0.9/24 n3           e4a         true
n3_clus2     up/up        10.10.0.10/24 n3          e4e         true
n4_clus1     up/up        10.10.0.11/24 n4          e4a         true
n4_clus2     up/up        10.10.0.12/24 n4          e4e         true
12 entries were displayed.
  
```

18. Verify that the cluster ports are connected:

network port show -role cluster

```

cluster::*> network port show -role cluster
(network port show)
Node: n1
-----
Port      IPspace      Broadcast Domain  Link  MTU      Speed(Mbps)  Health  Ignore
Admin/Oper  Status       Status          Health  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

Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health      Ignore
Admin/Oper Status      Health
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

Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health      Ignore
Admin/Oper Status      Health
Status
-----
e4a      Cluster      Cluster      up  9000 auto/40000 -    -
e4e      Cluster      Cluster      up  9000 auto/40000 -    -
Node: n4

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

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

cluster ping-cluster -node node-name

The following example shows node n1 being pinged and the RPC status indicated afterward:

```

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 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)

```

20. Display the information about the devices in your configuration by entering the following commands:

You can execute the following commands in any order:

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

```

cluster::> network device-discovery show
Node      Local   Discovered
Port      Device
-----
n1        /cdp
e0a      C1      Ethernet1/1/1   N3K-C3232C
e0b      C2      Ethernet1/1/1   N3K-C3232C
e0c      C2      Ethernet1/1/2   N3K-C3232C
e0d      C1      Ethernet1/1/2   N3K-C3232C
n2        /cdp
e0a      C1      Ethernet1/1/3   N3K-C3232C
e0b      C2      Ethernet1/1/3   N3K-C3232C
e0c      C2      Ethernet1/1/4   N3K-C3232C
e0d      C1      Ethernet1/1/4   N3K-C3232C
n3        /cdp
e4a      C1      Ethernet1/7     N3K-C3232C
e4e      C2      Ethernet1/7     N3K-C3232C
n4        /cdp
e4a      C1      Ethernet1/8     N3K-C3232C
e4e      C2      Ethernet1/8     N3K-C3232C

```

12 entries were displayed.

```

cluster::*> network port show -role cluster
(network port show)
Node: n1
Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health Ignore
Admin/Oper  Status      Health
-----
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
Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health Ignore
Admin/Oper  Status      Health
-----
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
Port      IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health Ignore
Admin/Oper  Status      Health
-----
e4a      Cluster      Cluster      up  9000  auto/40000  -

```

```
e4e      Cluster      Cluster      up  9000 auto/40000 -    -
Node: n4

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

12 entries were displayed.

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

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
Cluster	nml_clus1	up/up	10.10.0.1/24	n1	e0a	true
	n1_clus2	up/up	10.10.0.2/24	n1	e0b	true
	n1_clus3	up/up	10.10.0.3/24	n1	e0c	true
	n1_clus4	up/up	10.10.0.4/24	n1	e0d	true
	n2_clus1	up/up	10.10.0.5/24	n2	e0a	true
	n2_clus2	up/up	10.10.0.6/24	n2	e0b	true
	n2_clus3	up/up	10.10.0.7/24	n2	e0c	true
	n2_clus4	up/up	10.10.0.8/24	n2	e0d	true
	n3_clus1	up/up	10.10.0.9/24	n3	e4a	true
	n3_clus2	up/up	10.10.0.10/24	n3	e4e	true
	n4_clus1	up/up	10.10.0.11/24	n4	e4a	true
	n4_clus2	up/up	10.10.0.12/24	n4	e4e	true

12 entries were displayed.

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

Switch	Type	Address	Model
CL1	cluster-network	10.10.1.101	NX3232C
Serial Number: FOX000001 Is Monitored: true Reason: Software Version: Cisco Nexus Operating System (NX-OS) Software, Version 7.0(3)I6(1) Version Source: CDP			
CL2	cluster-network	10.10.1.102	NX3232C
Serial Number: FOX000002 Is Monitored: true Reason: Software Version: Cisco Nexus Operating System (NX-OS) Software, Version 7.0(3)I6(1) Version Source: CDP			
C2	cluster-network	10.10.1.103	NX3232C
Serial Number: FOX000003 Is Monitored: true Reason: Software Version: Cisco Nexus Operating System (NX-OS) Software, Version 7.0(3)I6(1) Version Source: CDP			

21. Delete the replaced cluster switch CL2 if it has not been removed automatically:

```
system cluster-switch delete -device cluster-switch-name
```

22. Verify that the proper cluster switches are monitored:

```
system cluster-switch show
```

The following example shows the cluster switches are monitored because the **Is Monitored** state is **true**.

Switch	Type	Address	Model
CL1	cluster-network	10.10.1.101	NX3232C
Serial Number: FOX000001 Is Monitored: true Reason: Software Version: Cisco Nexus Operating System (NX-OS) Software, Version 7.0(3)I6(1) Version Source: CDP			
C2	cluster-network	10.10.1.103	NX3232C
Serial Number: FOX000002 Is Monitored: true Reason:			

```
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version 7.0(3)I6(1)
Version Source: CDP
2 entries were displayed.
```

23. 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:
CL1
C2

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

Enter the switch name: CL1
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.

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

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

Related information

[Cisco Ethernet Switch description page](#)
[Hardware Universe](#)

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

cluster switches

replacing a Cisco Nexus 3232C non-disruptively [3](#)

requirements for replacing Cisco Nexus 3232C [3](#)

S

switches

replacing a Cisco Nexus 3232C non-disruptively in a cluster [3](#)