Lenovo

Cisco Nexus 3232C cluster switches

First edition (August 2021)

Contents

Contents i

Chapter 1. Installing a Cisco Nexus 3232C cluster switch in a rack 1

Chapter 2. Installing NX-OS software and RCFs on Cisco Nexus 3232C cluster switches 5

Installing the NX-OS software 8

Installing the Reference Configuration File (RCF) . . 11

Chapter 3. Migrating from a two-node switchless cluster to a cluster with Cisco Nexus 3232C cluster switches 19

How to migrate from a two-node switchless cluster to a cluster with Cisco Nexus 3232C cluster switches 19

Chapter 4. Replacing a Cisco Nexus 3232C cluster switch 29

How to replace a Cisco Nexus 3232C cluster switch 29

Chapter 5. Replacing a Cisco Nexus 3232C storage switch 41

Steps to replace a Cisco Nexus 3232C storage switch 41

Appendix A. Contacting Support . . . 45

Appendix B. Notices 47

Trademarks 48

Chapter 1. Installing a Cisco Nexus 3232C cluster switch in a rack

You can install the Cisco® Nexus 3232C switch in a rack with the standard brackets that are included with the switch.

You must have reviewed the initial preparation requirements, kit contents, and safety precautions.

Cisco Nexus 3000 Series Hardware Installation Guide

- For each switch, you must supply the eight 10-32 or 12-24 screws and clip nuts to mount the brackets and slider rails to the front and rear rack posts.
- You must use the Cisco standard rail kit to install the switch in a rack.

Note: The jumper cords should be included with your switches. If they were not shipped with the switches, you can order them from Lenovo (part number X1558A-R6).

Step 1. Install the rack-mount brackets on the Nexus 3232C switch chassis.

a. Position a front rack-mount bracket on one side of the switch chassis so that the mounting ear is aligned with the chassis faceplate (on the PSU or fan side), and then use four M4 screws to attach the bracket to the chassis.

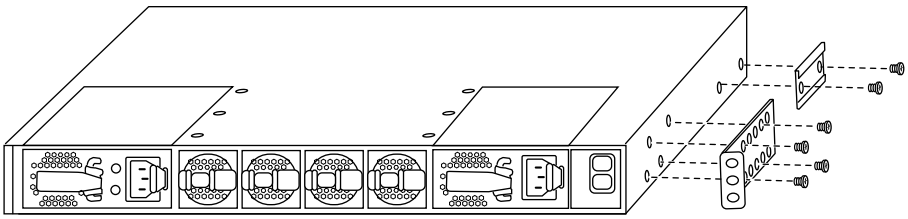

b. Repeat step 2a with the other front rack-mount bracket on the other side of the switch.

c. Install the rear rack-mount bracket on the switch chassis.

d. Repeat step 2c with the other rear rack-mount bracket on the other side of the switch.

Step 2. Install the clip nuts in the square hole locations for all four IEA posts.

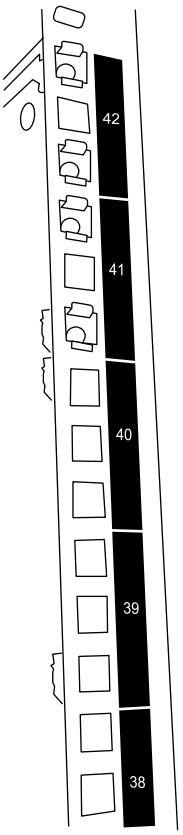

The two 3232C switches will always be mounted in the top 2U of the rack RU41 and 42.

Step 3. Install the slider rails in the rack.

a. Position the first slider rail at the RU42 mark on the back side of the rear left post, insert screws with the matching thread type, and then tighten the screws with your fingers.

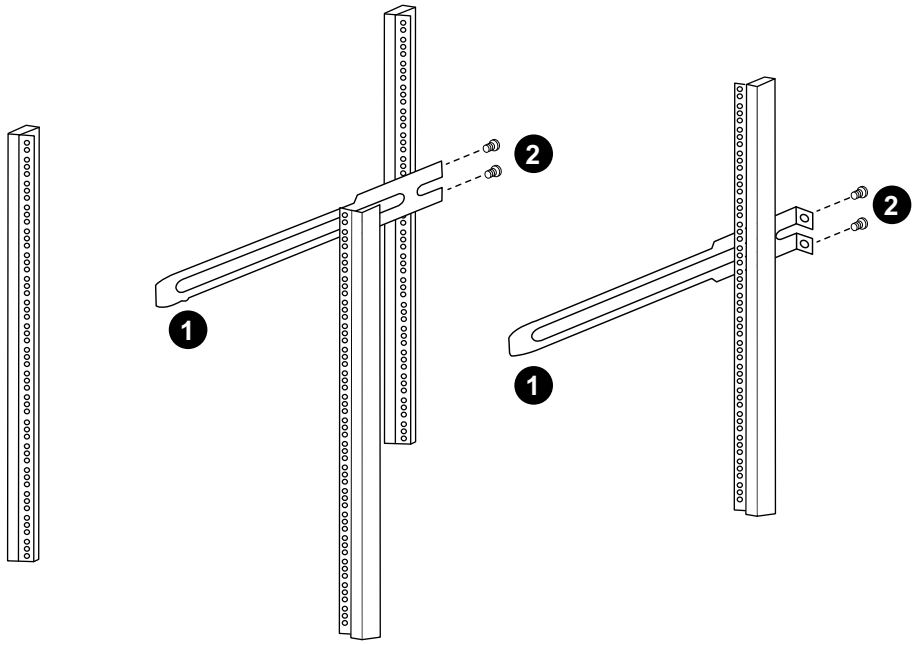

1. As you gently slide the slider rail, align it to the screw holes in the rack.

2. Tighten the screws of the slider rails to the rack posts.

b. Repeat step 4a for the right side rear post.

c. Repeat steps 4a and 4b at the RU41 locations on the rack.

Step 4. Install the switch in the rack.

Attention: This step requires two people: one person to support the switch from the front and another to guide the switch into the rear slider rails.

a. Position the back of the switch at RU41.

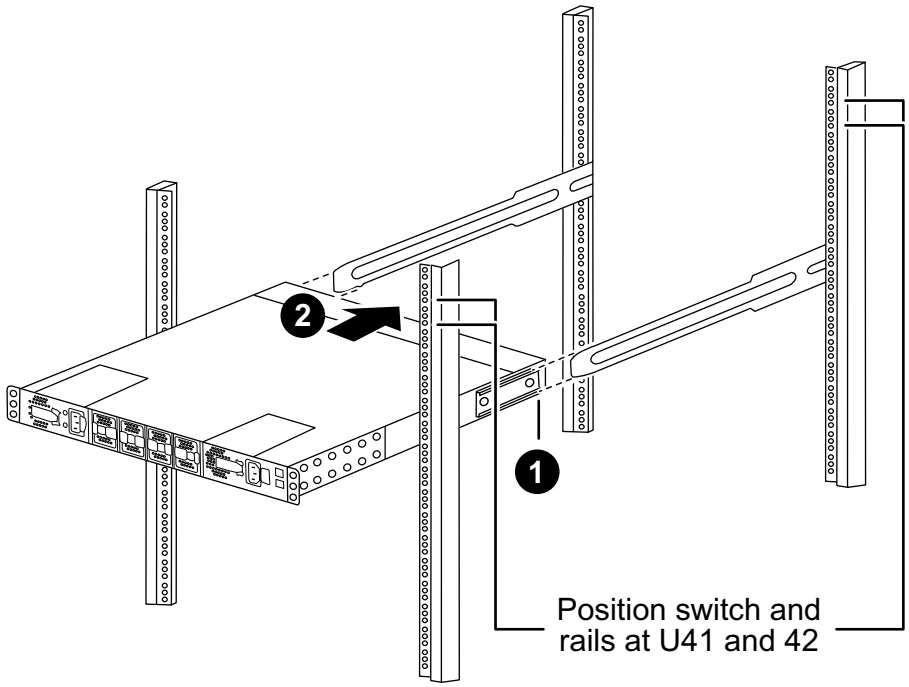

1. 1. As the chassis is pushed toward the rear posts, align the two rear rack-mount guides with the slider rails.
2. 2. Gently slide the switch until the front rack-mount brackets are flush with the front posts.

b. Attach the switch to the rack.

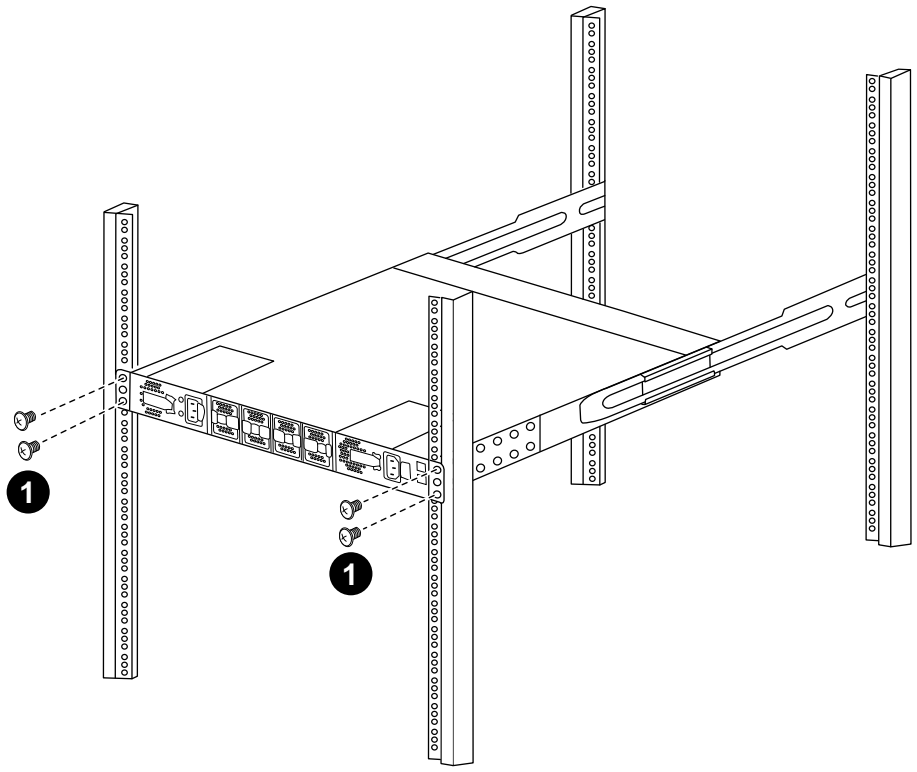

With one person holding the front of the chassis level, the other person should fully tighten the four rear screws to the rack posts.

c. With the chassis now supported without assistance, fully tighten the front screws to the posts.

d. Repeat steps 5a through 5c for the second switch at the RU42 location.

Note: By using the fully installed switch as a support, it is not necessary to hold the front of the second switch during the installation process.

Step 5. When the switches are installed, connect the jumper cords to the switch power inlets.

Step 6. Connect the male plugs of both jumper cords to the closest available PDU outlets.

Note: To maintain redundancy, the two cords must be connected to different PDUs.

Step 7. Connect the management port on each 3232C switch to either of the management switches (if ordered) or connect them directly to your management network. The management port is the upper-right port located on the PSU side of the switch. The CAT6 cable for each switch needs to be routed through the pass-through panel after the switches are installed to connect to the management switches or management network.

Chapter 2. Installing NX-OS software and RCFs on Cisco Nexus 3232C cluster switches

The Cisco® NX-OS software and reference configuration files (RCFs) must be installed on Cisco Nexus 3232C cluster switches.

The following conditions must exist before you install the NX-OS software and Reference Configurations Files (RCFs) on the cluster switch:

- The cluster must be fully functioning (there should be no errors in the logs or similar issues).
- You must have checked or set your desired boot configuration in the RCF to reflect the desired boot images if you are installing only NX-OS and keeping your current RCF version.
- If you need to change the boot configuration to reflect the current boot images, you must do so before reapplying the RCF so that the correct version is instantiated on future reboots.
- You must have a console connection to the switch, required when installing the RCF.
- Consult the compatible ONTAP and NX-OS versions in the change list.
- There can be command dependencies between the command syntax in the RCF and that found in versions of NX-OS.
- 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 on *Cisco Nexus 3000 Series Switches*.

  Cisco Nexus 3000 Series Switches
- You must have the current RCF.

The examples in this procedure use two nodes. These nodes use two 10GbE cluster interconnect ports *e0a* and *e0b*.

See the *Lenovo Press* to verify the correct cluster ports on your platforms.

Lenovo Press

Note: The command outputs might vary depending on different releases of ONTAP.

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

- The names of the two Cisco switches are *cs1* and *cs2*.
- The node names are *cluster1-01* and *cluster1-02*.
- The cluster LIF names are *cluster1-01_clus1* and *cluster1-01_clus2* for cluster1-01 and *cluster1-02_clus1* and *cluster1-02_clus2* for cluster1-02.
- The `cluster1::*>` prompt indicates the name of the cluster.

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

Step 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=x h`

In `MAINT=x h`, *x* is the duration of the maintenance window in hours.

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

Step 2. Change the privilege level to advanced, entering **y** when prompted to continue: `set -privilege advanced`
The advanced prompt (*>) appears.

Step 3. Display how many cluster interconnect interfaces are configured in each node for each cluster interconnect switch: `network device-discovery show -protocol cdp`

```
cluster1::*> network device-discovery show -protocol cdp

Node/       Local  Discovered
Protocol    Port   Device (LLDP: ChassisID)  Interface          Platform
----------- ------ ------------------------- ----------------- --------
cluster1-02/cdp
            e0a    cs1                       Eth1/2             N3K-C3232C
            e0b    cs2                       Eth1/2             N3K-C3232C
cluster1-01/cdp
            e0a    cs1                       Eth1/1             N3K-C3232C
            e0b    cs2                       Eth1/1             N3K-C3232C

4 entries were displayed.
```

Step 4. Check the administrative or operational status of each cluster interface.

a. Display the network port attributes: `network port show –ipspace Cluster`

```
cluster1::*> network port show -ipspace Cluster

Node: cluster1-02
                                                            Speed(Mbps) Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
--------- ------------ ---------------- ---- ---- ----------- ------
e0a       Cluster      Cluster          up   9000  auto/10000 healthy
e0b       Cluster      Cluster          up   9000  auto/10000 healthy

Node: cluster1-01
                                                            Speed(Mbps) Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
--------- ------------ ---------------- ---- ---- ----------- ------
e0a       Cluster      Cluster          up   9000  auto/10000 healthy
e0b       Cluster      Cluster          up   9000  auto/10000 healthy

4 entries were displayed.
```

b. Display information about the LIFs: `network interface show -vserver Cluster`

```
cluster1::*> network interface show -vserver Cluster

            Logical             Status     Network            Current       Current Is
Vserver     Interface           Admin/Oper Address/Mask       Node          Port    Home
----------- ------------------- ---------- ------------------ ------------- ------- ----
Cluster
            cluster1-01_clus1  up/up      169.254.209.69/16  cluster1-01   e0a     true
            cluster1-01_clus2  up/up      169.254.49.125/16  cluster1-01   e0b     true
            cluster1-02_clus1  up/up      169.254.47.194/16  cluster1-02   e0a     true
            cluster1-02_clus2  up/up      169.254.19.183/16  cluster1-02   e0b     true

4 entries were displayed.
```

Step 5. Ping the remote cluster LIFs: `cluster ping-cluster -node` *node-name*

```
cluster1::*> cluster ping-cluster -node cluster1-02
Host is cluster1-02
Getting addresses from network interface table...
Cluster cluster1-01_clus1 169.254.209.69 cluster1-01     e0a
Cluster cluster1-01_clus2 169.254.49.125 cluster1-01     e0b
Cluster cluster1-02_clus1 169.254.47.194 cluster1-02     e0a
Cluster cluster1-02_clus2 169.254.19.183 cluster1-02     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 169.254.19.183 to Remote 169.254.209.69
    Local 169.254.19.183 to Remote 169.254.49.125
    Local 169.254.47.194 to Remote 169.254.209.69
    Local 169.254.47.194 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)
```

Step 6. Verify that the `auto-revert` command is enabled on all cluster LIFs: `network interface show -vserver Cluster -fields auto-revert`

```
cluster1::*> network interface show -vserver Cluster -fields auto-revert

          Logical
Vserver   Interface            Auto-revert
--------- -------------------- ------------
Cluster
          cluster1-01_clus1    true
          cluster1-01_clus2    true
          cluster1-02_clus1    true
          cluster1-02_clus2    true
4 entries were displayed.
```

Step 7. For ONTAP 9.8 and later, enable the Ethernet switch health monitor log collection feature for collecting switch-related log files, using the commands: `system switch ethernet log setup-password` `system switch ethernet log enable-collection`

```
cluster1::*> system switch ethernet log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
cs1
cs2

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs1
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>

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs2
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>

cluster1::*> system  switch ethernet 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.

cluster1::*>
```

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

Installing the NX-OS software

You can use this procedure to install the NX-OS software on the Nexus 3232C cluster switch.

Step 1. Connect the cluster switch to the management network.

Step 2. Use the `ping` command to verify connectivity to the server hosting the NX-OS software and the RCF.This example verifies that the switch can reach the server at IP address 172.19.2.1:

```
cs2# ping 172.19.2.1
Pinging 172.19.2.1 with 0 bytes of data:

Reply From 172.19.2.1: icmp_seq = 0. time= 5910 usec.
```

Step 3. Copy the NX-OS software and EPLD images to the Nexus 3232C switch.

```
cs2# copy sftp: bootflash: vrf management
Enter source filename: /code/nxos.9.3.4.bin
Enter hostname for the sftp server: 172.19.2.1
Enter username: user1

Outbound-ReKey for 172.19.2.1:22
Inbound-ReKey for 172.19.2.1:22
user1@172.19.2.1's password:
sftp> progress
Progress meter enabled
sftp> get   /code/nxos.9.3.4.bin  /bootflash/nxos.9.3.4.bin
/code/nxos.9.3.4.bin  100% 1261MB   9.3MB/s   02:15
sftp> exit
Copy complete, now saving to disk (please wait)...
Copy complete.


cs2# copy sftp: bootflash: vrf management
Enter source filename: /code/n9000-epld.9.3.4.img
Enter hostname for the sftp server: 172.19.2.1
Enter username: user1

Outbound-ReKey for 172.19.2.1:22
Inbound-ReKey for 172.19.2.1:22
user1@172.19.2.1's password:
sftp> progress
Progress meter enabled
sftp> get   /code/n9000-epld.9.3.4.img  /bootflash/n9000-epld.9.3.4.img
/code/n9000-epld.9.3.4.img  100%  161MB   9.5MB/s   00:16
sftp> exit
Copy complete, now saving to disk (please wait)...
Copy complete.
```

Step 4. Verify the running version of the NX-OS software:

```
cs2# show version
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2019, Cisco and/or its affiliates.
All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under their own
licenses, such as open source.  This software is provided "as is," and unless
otherwise stated, there is no warranty, express or implied, including but not
limited to warranties of merchantability and fitness for a particular purpose.
Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or
GNU General Public License (GPL) version 3.0  or the GNU
Lesser General Public License (LGPL) Version 2.1 or
Lesser General Public License (LGPL) Version 2.0.
A copy of each such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://opensource.org/licenses/gpl-3.0.html and
http://www.opensource.org/licenses/lgpl-2.1.php and
```

```
http://www.gnu.org/licenses/old-licenses/library.txt.

Software
  BIOS: version 08.37
  NXOS: version 9.3(3)
  BIOS compile time:  01/28/2020
  NXOS image file is: bootflash:///nxos.9.3.3.bin
 NXOS compile time:  12/22/2019 2:00:00 [12/22/2019 14:00:37]


Hardware
  cisco Nexus3000 C3232C Chassis (Nexus 9000 Series)
  Intel(R) Xeon(R) CPU E5-2403 v2 @ 1.80GHz with 8154432 kB of memory.
  Processor Board ID FO??????GD

  Device name: cs2
  bootflash:   53298520 kB
Kernel uptime is 0 day(s), 0 hour(s), 3 minute(s), 36 second(s)

Last reset at 74117 usecs after Tue Nov 24 06:24:23 2020
  Reason: Reset Requested by CLI command reload
  System version: 9.3(3)
  Service:

plugin
  Core Plugin, Ethernet Plugin

Active Package(s):

cs2#
```

Step 5. Install the NX-OS image.

Installing the image file causes it to be loaded every time the switch is rebooted.

```
cs2# install all nxos bootflash:nxos.9.3.4.bin
Installer will perform compatibility check first. Please wait.
Installer is forced disruptive

Verifying image bootflash:/nxos.9.3.4.bin for boot variable "nxos".
[####################] 100% -- SUCCESS

Verifying image type.
[####################] 100% -- SUCCESS

Preparing "nxos" version info using image bootflash:/nxos.9.3.4.bin.
[####################] 100% -- SUCCESS

Preparing "bios" version info using image bootflash:/nxos.9.3.4.bin.
[####################] 100% -- SUCCESS

Performing module support checks.
[####################] 100% -- SUCCESS

Notifying services about system upgrade.
[####################] 100% -- SUCCESS


Compatibility check is done:
Module  bootable          Impact              Install-type  Reason
------- ----------------- ------------------- ------------- ------
     1     yes             disruptive          reset         default upgrade is not hitless


Images will be upgraded according to following table:
Module       Image       Running-Version(pri:alt)                New-Version          Upg-Required
------------ ----------- --------------------------------------- -------------------- ------------
     1       nxos        9.3(3)                                    9.3(4)               yes
```

```
     1       bios         v08.37(01/28/2020):v08.32(10/18/2016)   v08.37(01/28/2020)   no


Switch will be reloaded for disruptive upgrade.
Do you want to continue with the installation (y/n)?  [n] y

Install is in progress, please wait.

Performing runtime checks.
[####################] 100% -- SUCCESS

Setting boot variables.
[####################] 100% -- SUCCESS

Performing configuration copy.
[####################] 100% -- SUCCESS

Module 1: Refreshing compact flash and upgrading bios/loader/bootrom.
Warning: please do not remove or power off the module at this time.
[####################] 100% -- SUCCESS


Finishing the upgrade, switch will reboot in 10 seconds.
cs2#
```

Step 6. Verify the new version of NX-OS software after the switch has rebooted: `show version`

```
cs2# show version
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2020, Cisco and/or its affiliates.
All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under their own
licenses, such as open source.  This software is provided "as is," and unless
otherwise stated, there is no warranty, express or implied, including but not
limited to warranties of merchantability and fitness for a particular purpose.
Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or
GNU General Public License (GPL) version 3.0  or the GNU
Lesser General Public License (LGPL) Version 2.1 or
Lesser General Public License (LGPL) Version 2.0.
A copy of each such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://opensource.org/licenses/gpl-3.0.html and
http://www.opensource.org/licenses/lgpl-2.1.php and
http://www.gnu.org/licenses/old-licenses/library.txt.

Software
  BIOS: version 08.37
  NXOS: version 9.3(4)
  BIOS compile time:  01/28/2020
  NXOS image file is: bootflash:///nxos.9.3.4.bin
  NXOS compile time:  4/28/2020 21:00:00 [04/29/2020 06:28:31]


Hardware
 cisco Nexus3000 C3232C Chassis (Nexus 9000 Series)
  Intel(R) Xeon(R) CPU E5-2403 v2 @ 1.80GHz with 8154432 kB of memory.
  Processor Board ID FO??????GD

  Device name: rtpnpi-mcc01-8200-ms-A1
  bootflash:   53298520 kB
Kernel uptime is 0 day(s), 0 hour(s), 3 minute(s), 14 second(s)

Last reset at 196755 usecs after Tue Nov 24 06:37:36 2020
  Reason: Reset due to upgrade
  System version: 9.3(3)
  Service:
```

```
plugin
  Core Plugin, Ethernet Plugin

Active Package(s):

cs2#
```

Step 7. Upgrade the EPLD image and reboot the switch.

```
cs2# show version module 1 epld

EPLD Device                       Version
---------------------------------------
MI   FPGA                         0x12
IO   FPGA                         0x11

cs2# install epld bootflash:n9000-epld.9.3.4.img module 1
Compatibility check:
Module         Type          Upgradable     Impact      Reason
------  -----------------  ----------    ----------   ------
     1          SUP           Yes           disruptive   Module Upgradable

Retrieving EPLD versions.... Please wait.
Images will be upgraded according to following table:
Module  Type   EPLD              Running-Version   New-Version  Upg-Required
------  ----  -------------      ---------------   -----------  ------------
     1   SUP   MI FPGA                   0x12        0x12            No
     1   SUP   IO FPGA                   0x11        0x12            Yes
The above modules require upgrade.
The switch will be reloaded at the end of the upgrade
Do you want to continue (y/n) ?  [n] y

Proceeding to upgrade Modules.

Starting Module 1 EPLD Upgrade

Module 1 : IO FPGA [Programming] : 100.00% (     64 of      64 sectors)
Module 1 EPLD upgrade is successful.
Module         Type  Upgrade-Result
------  ------------------  --------------
     1          SUP         Success

Module 1 EPLD upgrade is successful.
```

Step 8. After the switch reboot, log in again and verify that the new version of EPLD loaded successfully.

```
cs2# show version module 1 epld

EPLD Device                       Version
---------------------------------------
MI   FPGA                         0x12
IO   FPGA                         0x12
```

Installing the Reference Configuration File (RCF)

You can install the RCF after setting up the Nexus 3232C switch for the first time.You can also use this procedure to upgrade your RCF version.

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

- The names of the two Cisco switches are *cs1* and *cs2*.
- The node names are *cluster1-01*, *cluster1-02*, *cluster1-03*, and *cluster1-04*.
- The cluster LIF names are *cluster1-01_clus1*, *cluster1-01_clus2*, *cluster1-02_clus1*, *cluster1-02_clus2* , *cluster1-03_clus1*, *cluster1-03_clus2*, *cluster1-04_clus1* , and *cluster1-04_clus2*.
- The `cluster1::*>` prompt indicates the name of the cluster.

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

Step 1. Display the cluster ports on each node that are connected to the cluster switches: `network device-discovery show`

```
cluster1::*> network device-discovery show
Node/       Local  Discovered
Protocol    Port   Device (LLDP: ChassisID)  Interface          Platform
----------- ------ ------------------------- ----------------  --------
cluster1-01/cdp
            e0a    cs1                       Ethernet1/7       N3K-C3232C
            e0d    cs2                       Ethernet1/7       N3K-C3232C
cluster1-02/cdp
            e0a    cs1                       Ethernet1/8       N3K-C3232C
            e0d    cs2                       Ethernet1/8       N3K-C3232C
cluster1-03/cdp
            e0a    cs1                       Ethernet1/1/1     N3K-C3232C
            e0b    cs2                       Ethernet1/1/1     N3K-C3232C
cluster1-04/cdp
            e0a    cs1                       Ethernet1/1/2     N3K-C3232C
            e0b    cs2                       Ethernet1/1/2     N3K-C3232C
cluster1::*>
```

Step 2. Check the administrative and operational status of each cluster port.

a. Verify that all the cluster ports are up with a healthy status: `network port show –role cluster`

```
cluster1::*> network port show -role cluster

Node: cluster1-01
                                                                        Ignore
                                                  Speed(Mbps) Health   Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status   Status
--------- ------------ ---------------- ---- ---- ----------- -------- ------
e0a       Cluster      Cluster          up   9000  auto/100000 healthy false
e0d       Cluster      Cluster          up   9000  auto/100000 healthy false

Node: cluster1-02
                                                                        Ignore
                                                  Speed(Mbps) Health   Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status   Status
--------- ------------ ---------------- ---- ---- ----------- -------- ------
e0a       Cluster      Cluster          up   9000  auto/100000 healthy false
e0d       Cluster      Cluster          up   9000  auto/100000 healthy false
8 entries were displayed.

Node: cluster1-03

   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: cluster1-04
                                                                        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
cluster1::*>
```

b. Verify that all the cluster interfaces (LIFs) are on the home port: `network interface show -role cluster`

```
cluster1::*> network interface show -role cluster
            Logical             Status     Network            Current      Current Is
Vserver     Interface           Admin/Oper Address/Mask       Node         Port    Home
----------- ------------------- ---------- ------------------ ------------ ------- ----
```

```
Cluster
            cluster1-01_clus1  up/up     169.254.3.4/23     cluster1-01  e0a     true
            cluster1-01_clus2  up/up     169.254.3.5/23     cluster1-01  e0d     true
            cluster1-02_clus1  up/up     169.254.3.8/23     cluster1-02  e0a     true
            cluster1-02_clus2  up/up     169.254.3.9/23     cluster1-02  e0d     true
            cluster1-03_clus1  up/up     169.254.1.3/23     cluster1-03  e0a     true
            cluster1-03_clus2  up/up     169.254.1.1/23     cluster1-03  e0b     true
            cluster1-04_clus1  up/up     169.254.1.6/23     cluster1-04  e0a     true
            cluster1-04_clus2  up/up     169.254.1.7/23     cluster1-04  e0b     true
8 entries were displayed.
cluster1::*>
```

c. Verify that the cluster displays information for both cluster switches: `system cluster-switch show -is-monitoring-enabled-operational true`

```
cluster1::*> system cluster-switch show -is-monitoring-enabled-operational true
Switch                        Type               Address          Model
--------------------------- ------------------ ---------------- ---------------
cs1                           cluster-network    10.233.205.92    NX3232C
     Serial Number: FOXXXXXXXGS
      Is Monitored: true
            Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS) Software, Version
                    9.3(4)
    Version Source: CDP

cs2                           cluster-network    10.233.205.93    NX3232C
     Serial Number: FOXXXXXXXGD
      Is Monitored: true
            Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS) Software, Version
                    9.3(4)
    Version Source: CDP

2 entries were displayed.
```

Step 3. Disable auto-revert on the cluster LIFs.

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto-revert false
```

Step 4. On cluster switch cs2, shut down the ports connected to the cluster ports of the nodes.

```
cs2(config)# interface eth1/1/1-2,eth1/7-8
cs2(config-if-range)# shutdown
```

Step 5. Verify that the cluster ports have migrated to the ports hosted on cluster switch cs1. This might take a few seconds. `network interface show -role cluster`

```
cluster1::*> network interface show -role cluster
            Logical            Status     Network            Current       Current Is
Vserver     Interface          Admin/Oper Address/Mask       Node          Port    Home
----------- ------------------ ---------- ------------------ ------------- ------- ----
Cluster
            cluster1-01_clus1 up/up      169.254.3.4/23     cluster1-01   e0a     true
            cluster1-01_clus2 up/up      169.254.3.5/23     cluster1-01   e0a     false
            cluster1-02_clus1 up/up      169.254.3.8/23     cluster1-02   e0a     true
            cluster1-02_clus2 up/up      169.254.3.9/23     cluster1-02   e0a     false
            cluster1-03_clus1 up/up      169.254.1.3/23     cluster1-03   e0a     true
            cluster1-03_clus2 up/up      169.254.1.1/23     cluster1-03   e0a     false
            cluster1-04_clus1 up/up      169.254.1.6/23     cluster1-04   e0a     true
            cluster1-04_clus2 up/up      169.254.1.7/23     cluster1-04   e0a     false
8 entries were displayed.
cluster1::*>
```

Step 6. Verify that the cluster is healthy: `cluster show`

```
cluster1::*> cluster show
Node                 Health  Eligibility   Epsilon
-------------------- ------- ------------  -------
cluster1-01          true    true          false
cluster1-02          true    true          false
cluster1-03          true    true          true
cluster1-04          true    true          false
4 entries were displayed.
cluster1::*>
```

Step 7. Clean the configuration on switch cs2 and perform a basic setup.

a. Clean the configuration. This step requires a console connection to the switch.

```
cs2# write erase
Warning: This command will erase the startup-configuration.
Do you wish to proceed anyway? (y/n)  [n] y
cs2# reload
This command will reboot the system. (y/n)?  [n] y
cs2#
```

b. Perform a basic setup of the switch.

Step 8. Copy the RCF to the bootflash of switch cs2 using one of the following transfer protocols: FTP, TFTP, SFTP, or SCP. For more information on Cisco commands, see the appropriate guide in the Cisco Nexus 3000 Series NX-OS Command Reference guides.This example shows TFTP being used to copy an RCF to the bootflash on switch cs2:

```
cs2# copy tftp: bootflash: vrf management
Enter source filename: Nexus_3232C_RCF_v1.6-Cluster-HA-Breakout.txt
Enter hostname for the tftp server: 172.22.201.50
Trying to connect to tftp server.....Connection to Server Established.
TFTP get operation was successful
Copy complete, now saving to disk (please wait)...
```

Step 9. Apply the RCF previously downloaded to the bootflash.

For more information on Cisco commands, see the appropriate guide in the Cisco Nexus 3000 Series NX-OS Command Reference guides.

This example shows the RCF file `Nexus_3232C_RCF_v1.6-Cluster-HA-Breakout.txt` being installed on switch cs2:

```
cs2# copy Nexus_3232C_RCF_v1.6-Cluster-HA-Breakout.txt running-config echo-commands
```

Step 10. Examine the banner output from the `show banner motd` command. You must read and follow the instructions under **Important Notes** to ensure the proper configuration and operation of the switch.

```
**************************************************************************
* Lenovo Reference Configuration File (RCF)
*
* Switch   : Cisco Nexus 3232C
* Filename : Nexus_3232C_RCF_v1.6-Cluster-HA-Breakout.txt
* Date     : Oct-20-2020
* Version  : v1.6
*
* Port Usage : Breakout configuration
* Ports  1- 3: Breakout mode (4x10GbE) Intra-Cluster Ports, int e1/1/1-4,
* e1/2/1-4, e1/3/1-4
* Ports  4- 6: Breakout mode (4x25GbE) Intra-Cluster/HA Ports, int e1/4/1-4,
* e1/5/1-4, e1/6/1-4
* Ports  7-30: 40/100GbE Intra-Cluster/HA Ports, int e1/7-30
* Ports 31-32: Intra-Cluster ISL Ports, int e1/31-32
* Ports 33-34: 10GbE Intra-Cluster 10GbE Ports, int e1/33-34
*
* IMPORTANT NOTES
* - Load Nexus_3232C_RCF_v1.6-Cluster-HA.txt for non breakout config
*
* - This RCF utilizes QoS and requires TCAM re-configuration, requiring RCF
*   to be loaded twice with the Cluster Switch rebooted in between.
*
* - Perform the following 4 steps to ensure proper RCF installation:
*
*   (1) Apply RCF first time, expect following messages:
*       - Please save config and reload the system...
*       - Edge port type (portfast) should only be enabled on ports...
*       - TCAM region is not configured for feature QoS class IPv4 ingress...
*
*   (2) Save running-configuration and reboot Cluster Switch
*
*   (3) After reboot, apply same RCF second time and expect following messages:
*       - % Invalid command at '^' marker
```

```
*       - Syntax error while parsing...
*
*   (4) Save running-configuration again
*******************************************************************************
```

Step 11. Verify that the RCF file is the correct newer version: `show running-config`

When you check the output to verify you have the correct RCF, make sure that the following information is correct:

- The RCF banner
- The node and port settings
- Customizations

The output varies according to your site configuration. Check the port settings and refer to the release notes for any changes specific to the RCF that you have installed.

Step 12. After you verify the RCF versions and switch settings are correct, copy the running-config file to the startup-config file.For more information on Cisco commands, see the appropriate guide in the Cisco Nexus 3000 Series NX-OS Command Reference guides.

```
cs2# copy running-config startup-config [########################################] 100% Copy complete
```

Step 13. Reboot switch cs2. You can ignore the “cluster ports down” events reported on the nodes while the switch reboots.

```
cs2# reload
This command will reboot the system. (y/n)?  [n] y
```

Step 14. Apply the same RCF and save the running configuration for a second time.

```
cs2# copy Nexus_3232C_RCF_v1.6-Cluster-HA-Breakout.txt running-config echo-commands
cs2# copy running-config startup-config [########################################] 100% Copy complete
```

Step 15. Verify the health of cluster ports on the cluster.

a. Verify that e0d ports are up and healthy across all nodes in the cluster: `network port show -role cluster`

```
cluster1::*> network port show -role cluster

Node: cluster1-01
                                                                        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: cluster1-02
                                                                        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: cluster1-03
                                                                        Ignore
                                                  Speed(Mbps) Health   Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status   Status
--------- ------------ ---------------- ---- ---- ----------- -------- ------
e0a       Cluster      Cluster          up   9000  auto/100000 healthy false
e0d       Cluster      Cluster          up   9000  auto/100000 healthy false

Node: cluster1-04
                                                                        Ignore
                                                  Speed(Mbps) Health   Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status   Status
--------- ------------ ---------------- ---- ---- ----------- -------- ------
e0a       Cluster      Cluster          up   9000  auto/100000 healthy false
```

```
e0d       Cluster      Cluster          up   9000  auto/100000 healthy false
8 entries were displayed.
```

b. Verify the switch health from the cluster (this might not show switch cs2, since LIFs are not homed on e0d).

```
cluster1::*> network device-discovery show -protocol cdp
Node/       Local  Discovered
Protocol    Port   Device (LLDP: ChassisID)  Interface         Platform
----------- ------ ------------------------- ----------------- --------
cluster1-01/cdp
            e0a    cs1                       Ethernet1/7       N3K-C3232C
            e0d    cs2                       Ethernet1/7       N3K-C3232C
cluster01-2/cdp
            e0a    cs1                       Ethernet1/8       N3K-C3232C
            e0d    cs2                       Ethernet1/8       N3K-C3232C
cluster01-3/cdp
            e0a    cs1                       Ethernet1/1/1     N3K-C3232C
            e0b    cs2                       Ethernet1/1/1     N3K-C3232C
cluster1-04/cdp
            e0a    cs1                       Ethernet1/1/2     N3K-C3232C
            e0b    cs2                       Ethernet1/1/2     N3K-C3232C

cluster1::*> system cluster-switch show -is-monitoring-enabled-operational true
Switch                      Type               Address          Model
--------------------------- ------------------ ---------------- -----
cs1                         cluster-network    10.233.205.90    N3K-C3232C
     Serial Number: FOXXXXXXXGD
      Is Monitored: true
            Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS) Software, Version
                    9.3(4)
    Version Source: CDP

cs2                         cluster-network    10.233.205.91    N3K-C3232C
     Serial Number: FOXXXXXXXGS
      Is Monitored: true
            Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS) Software, Version
                    9.3(4)
    Version Source: CDP

2 entries were displayed.
```

Note: You might observe the following output on the cs1 switch console depending on the RCF version previously loaded on the switch:

```
2020 Nov 17 16:07:18 cs1 %$ VDC-1 %$ %STP-2-UNBLOCK_CONSIST_PORT: Unblocking port port-channel1 on VLAN0092.
 Port consistency restored.
2020 Nov 17 16:07:23 cs1 %$ VDC-1 %$ %STP-2-BLOCK_PVID_PEER: Blocking port-channel1 on VLAN0001. Inconsistent peer vlan.
2020 Nov 17 16:07:23 cs1 %$ VDC-1 %$ %STP-2-BLOCK_PVID_LOCAL: Blocking port-channel1 on VLAN0092. Inconsistent local vlan.
```

Step 16. On cluster switch cs1, shut down the ports connected to the cluster ports of the nodes.The following example uses the interface example output from step 1:

```
cs1(config)# interface eth1/1/1-2,eth1/7-8
cs1(config-if-range)# shutdown
```

Step 17. Verify that the cluster LIFs have migrated to the ports hosted on switch cs2. This might take a few seconds. `network interface show -role cluster`

```
cluster1::*> network interface show -role cluster
            Logical            Status     Network            Current             Current Is
Vserver     Interface          Admin/Oper Address/Mask       Node                Port    Home
----------- ------------------ ---------- ------------------ ------------------- ------- ----
Cluster
            cluster1-01_clus1  up/up      169.254.3.4/23     cluster1-01         e0d     false
            cluster1-01_clus2  up/up      169.254.3.5/23     cluster1-01         e0d     true
            cluster1-02_clus1  up/up      169.254.3.8/23     cluster1-02         e0d     false
            cluster1-02_clus2  up/up      169.254.3.9/23     cluster1-02         e0d     true
            cluster1-03_clus1  up/up      169.254.1.3/23     cluster1-03         e0b     false
```

```
            cluster1-03_clus2  up/up      169.254.1.1/23     cluster1-03         e0b     true
            cluster1-04_clus1  up/up      169.254.1.6/23     cluster1-04         e0b     false
            cluster1-04_clus2  up/up      169.254.1.7/23     cluster1-04         e0b     true
8 entries were displayed.
cluster1::*>
```

Step 18. Verify that the cluster is healthy: `cluster show`

```
cluster1::*> cluster show
Node                 Health   Eligibility   Epsilon
-------------------- -------- ------------- -------
cluster1-01          true     true          false
cluster1-02          true     true          false
cluster1-03          true     true          true
cluster1-04          true     true          false
4 entries were displayed.
cluster1::*>
```

Step 19. Repeat Steps 7 to 14 on switch cs1.

Step 20. Enable auto-revert on the cluster LIFs.

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto-revert True
```

Step 21. Reboot switch cs1.You do this to trigger the cluster LIFs to revert to their home ports. You can ignore the “cluster ports down” events reported on the nodes while the switch reboots.

```
cs1# reload
This command will reboot the system. (y/n)?  [n] y
```

Step 22. Verify that the switch ports connected to the cluster ports are up.

```
cs1# show interface brief | grep up
.
.
Eth1/1/1      1       eth  access up      none                    10G(D) --
Eth1/1/2      1       eth  access up      none                    10G(D) --
Eth1/7        1       eth  trunk  up      none                   100G(D) --
Eth1/8        1       eth  trunk  up      none                   100G(D) --
.
.
```

Step 23. Verify that the ISL between cs1 and cs2 is functional: `show port-channel summary`

```
cs1# 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
        b - BFD Session Wait
        S - Switched    R - Routed
        U - Up (port-channel)
        p - Up in delay-lacp mode (member)
        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)
cs1#
```

Step 24. Verify that the cluster LIFs have reverted to their home port: `network interface show -role cluster`

```
cluster1::*> network interface show -role cluster
             Logical            Status     Network            Current             Current Is
Vserver      Interface          Admin/Oper Address/Mask       Node                Port    Home
------------ ------------------ ---------- ------------------ ------------------- ------- ----
Cluster
             cluster1-01_clus1  up/up      169.254.3.4/23     cluster1-01         e0d     true
             cluster1-01_clus2  up/up      169.254.3.5/23     cluster1-01         e0d     true
             cluster1-02_clus1  up/up      169.254.3.8/23     cluster1-02         e0d     true
             cluster1-02_clus2  up/up      169.254.3.9/23     cluster1-02         e0d     true
             cluster1-03_clus1  up/up      169.254.1.3/23     cluster1-03         e0b     true
             cluster1-03_clus2  up/up      169.254.1.1/23     cluster1-03         e0b     true
             cluster1-04_clus1  up/up      169.254.1.6/23     cluster1-04         e0b     true
             cluster1-04_clus2  up/up      169.254.1.7/23     cluster1-04         e0b     true
8 entries were displayed.
```

```
cluster1::*>
```

Step 25. Verify that the cluster is healthy: `cluster show`

```
cluster1::*> cluster show
Node                 Health  Eligibility   Epsilon
-------------------- ------- ------------- -------
cluster1-01          true    true          false
cluster1-02          true    true          false
cluster1-03          true    true          true
cluster1-04          true    true          false
4 entries were displayed.
cluster1::*>
```

Step 26. Ping the remote cluster interfaces to verify connectivity: `cluster ping-cluster -node local`

```
cluster1::*> cluster ping-cluster -node local
Host is cluster1-03
Getting addresses from network interface table...
Cluster cluster1-03_clus1 169.254.1.3 cluster1-03 e0a
Cluster cluster1-03_clus2 169.254.1.1 cluster1-03 e0b
Cluster cluster1-04_clus1 169.254.1.6 cluster1-04 e0a
Cluster cluster1-04_clus2 169.254.1.7 cluster1-04 e0b
Cluster cluster1-01_clus1 169.254.3.4 cluster1-01 e0a
Cluster cluster1-01_clus2 169.254.3.5 cluster1-01 e0d
Cluster cluster1-02_clus1 169.254.3.8 cluster1-02 e0a
Cluster cluster1-02_clus2 169.254.3.9 cluster1-02 e0d
Local = 169.254.1.3 169.254.1.1
Remote = 169.254.1.6 169.254.1.7 169.254.3.4 169.254.3.5 169.254.3.8 169.254.3.9
Cluster Vserver Id = 4294967293
Ping status:
............
Basic connectivity succeeds on 12 path(s)
Basic connectivity fails on 0 path(s)
................................................
Detected 9000 byte MTU on 12 path(s):
    Local 169.254.1.3 to Remote 169.254.1.6
    Local 169.254.1.3 to Remote 169.254.1.7
    Local 169.254.1.3 to Remote 169.254.3.4
    Local 169.254.1.3 to Remote 169.254.3.5
    Local 169.254.1.3 to Remote 169.254.3.8
    Local 169.254.1.3 to Remote 169.254.3.9
    Local 169.254.1.1 to Remote 169.254.1.6
    Local 169.254.1.1 to Remote 169.254.1.7
    Local 169.254.1.1 to Remote 169.254.3.4
    Local 169.254.1.1 to Remote 169.254.3.5
    Local 169.254.1.1 to Remote 169.254.3.8
    Local 169.254.1.1 to Remote 169.254.3.9
Larger than PMTU communication succeeds on 12 path(s)
RPC status:
6 paths up, 0 paths down (tcp check)
6 paths up, 0 paths down (udp check)
```

Chapter 3. Migrating from a two-node switchless cluster to a cluster with Cisco Nexus 3232C cluster switches

You must be aware of certain configuration information, port connections, and cabling requirements when you migrate a two-node switchless cluster to a cluster with Cisco Nexus 3232C cluster switches.

You must have the following before you begin the migration process:

- Available ports for node connections:
  The cluster switches use the Inter-Switch Link (ISL) ports e1/31-32.
- Appropriate cables for cluster connections:
  - The nodes with 10 GbE cluster connections require QSFP optical modules with breakout fiber cables or QSFP to SFP+ copper breakout 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 require the appropriate ISL cabling: 2x QSFP28 fiber or copper direct-attach cables.

How to migrate from a two-node switchless cluster to a cluster with Cisco Nexus 3232C cluster switches

If you have a two-node switchless cluster, you can migrate nondisruptively to a two-node switched cluster that includes Cisco Nexus 3232C cluster network switches.

- The configurations must be properly set up and functioning.The two nodes must be connected and functioning in a two-node switchless cluster setting.
- All cluster ports must be in the up state.
- The Cisco Nexus 3232C cluster switch must be supported.
- The existing cluster network configuration must have the following:
  - A redundant and fully functional Nexus 3232C cluster infrastructure 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 having been migrated
  - Initial customization of the switch
  - All ISL ports enabled and cabled

Procedure summary

- **I. Display and migrate physical and logical ports (Steps 1-10)**
- **II. Shut down the reassigned LIFs and disconnect the cables (Steps 11-14)**
- **III. Enable the cluster ports (Steps 15-20)**
- **IV. Enable the reassigned LIFs (Steps 21-33)**

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

- Nexus 3232C cluster switches, C1 and C2.

- The nodes are n1 and n2.

Note: The examples in this procedure use two nodes, each utilizing two 100 GbE cluster interconnect ports e3a and e3b. The *Lenovo Press* 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 Lenovo Datacenter Support Download page for a DM model.

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

Step 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`

In MAINT=*x*h, *x* is the duration of the maintenance window in hours.

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

Step 2. 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
                                                                        Ignore
                                                  Speed(Mbps) Health   Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status   Status
--------- ------------ ---------------- ---- ---- ----------- -------- -----
e3a       Cluster      Cluster          up   9000 auto/100000  -
e3b       Cluster      Cluster          up   9000 auto/100000  -        -
Node: n2
                                                                        Ignore
                                                  Speed(Mbps) Health   Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status   Status
--------- ------------ ---------------- ---- ---- ----------- -------- -----
e3a       Cluster      Cluster          up   9000 auto/100000  -
e3b       Cluster      Cluster          up   9000 auto/100000  -
4 entries were displayed.
```

b. Display information about the logical interfaces and their designated home nodes: `network interface show -role cluster`

```
cluster::*> network interface show -role cluster
 (network interface show)
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ---
Cluster
            n1_clus1   up/up      10.10.0.1/24       n1            e3a     true
            n1_clus2   up/up      10.10.0.2/24       n1            e3b     true
            n2_clus1   up/up      10.10.0.3/24       n2            e3a     true
            n2_clus2   up/up      10.10.0.4/24       n2            e3b     true

4 entries were displayed.
```

c. Verify that switchless cluster detection is enabled using the advanced privilege command:`network options detect-switchless-cluster show`

The output in the following example shows that switchless cluster detection is enabled:

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

Step 3. Verify that the appropriate RCFs and image are installed on the new 3232C switches and make any necessary site customizations such as adding users, passwords, and network addresses.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 Lenovo Datacenter Support Site for a DM Series Product.

b. Download the appropriate version of RCF.

c. Refer to the Cisco website for switch image downloads.

Step 4. On Nexus 3232C switches C1 and C2, disable all node-facing ports C1 and C2, but do not disable the ISL ports e1/31-32.For more information on Cisco commands, see the guides listed in the Cisco Nexus 3000 Series NX-OS Command References.

The following example shows ports 1 through 30 being disabled on Nexus 3232C cluster switches C1 and C2:

```
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
```

```
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
```

Step 5. Connect ports 1/31 and 1/32 on C1 to the same ports on C2 using supported cabling.

Step 6. Verify that the ISL ports are operational on C1 and C2: `show port-channel summary`
For more information on Cisco commands, see the guides listed in the Cisco Nexus 3000 Series NX-OS Command References.

The following example shows the Cisco `show port-channel summary` command being used to verify the ISL ports are operational on C1 and C2:

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

Step 7. Display the list of neighboring devices on the switch.For more information on Cisco commands, see the guides listed in the Cisco Nexus 3000 Series NX-OS Command References.

The following example shows the Cisco command `show cdp neighbors` being used to display the neighboring devices on the switch:

```
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-C3232C  Eth1/31
C2                  Eth1/32        174    R S I s     N3K-C3232C  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-C3232C  Eth1/31
C1                  Eth1/32        178    R S I s     N3K-C3232C  Eth1/32
Total entries displayed: 2
```

Step 8. Display the cluster port connectivity on each node: `network device-discovery show`

The following example shows the cluster port connectivity displayed for a two-node switchless cluster configuration:

```
cluster::*> network device-discovery show
            Local  Discovered
Node        Port   Device              Interface        Platform
----------- ------ ------------------ ---------------- ----------------
n1         /cdp
            e3a    n2                  e3a              DM7100F
            e3b    n2                  e3b              DM7100F
n2         /cdp
            e3a    n1                  e3a              DM7100F
            e3b    n1                  e3b              DM7100F
```

Step 9. Migrate the n1_clus1 and n2_clus1 LIFs to the physical ports of their destination nodes: `network interface migrate -vserver cluster -lif` *lif-name* `source-node` *source-node-name* `-destination-port` *destination-port-name*

You must execute the command for each local node as shown in the following example:

```
cluster::*> network interface migrate -vserver cluster -lif n1_clus1 -source-node n1
–destination-node n1 -destination-port e3b
cluster::*> network interface migrate -vserver cluster -lif n2_clus1 -source-node n2
–destination-node n2 -destination-port e3b
```

Step 10. Verify the cluster interfaces have successfully migrated: `network interface show -role cluster`

The following example shows the “Is Home” status for the n1_clus1 and n2_clus1 LIFs has become “false” after the migration is completed:

```
cluster::*> network interface show -role cluster
 (network interface show)
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
Cluster
            n1_clus1   up/up      10.10.0.1/24       n1            e3b     false
            n1_clus2   up/up      10.10.0.2/24       n1            e3b     true
            n2_clus1   up/up      10.10.0.3/24       n2            e3b     false
            n2_clus2   up/up      10.10.0.4/24       n2            e3b     true
 4 entries were displayed.
```

Step 11. Shut down cluster ports for the n1_clus1 and n2_clus1 LIFs, which were migrated in step 10: `network port modify -node` *node-name* `-port` *port-name* `-up-admin false`

You must execute the command for each port as shown in the following example:

```
cluster::*> network port modify -node n1 -port e3a -up-admin false
cluster::*> network port modify -node n2 -port e3a -up-admin false
```

Step 12. 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        e3a    10.10.0.1
Cluster n1_clus2 n1        e3b    10.10.0.2
Cluster n2_clus1 n2        e3a    10.10.0.3
Cluster n2_clus2 n2        e3b    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 9000 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)
```

Step 13. Disconnect the cable from e3a on node n1.You can refer to the running configuration and connect the first 100 GbE port on the switch C1 (port 1/7 in this example) to e3a on n1 using cabling supported for Nexus 3232C switches.

Step 14. Disconnect the cable from e3a on node n2.You can refer to the running configuration and connect e3a to the next available 100 GbE port on C1, port 1/8, using supported cabling.

Step 15. Enable all node-facing ports on C1.For more information on Cisco commands, see the guides listed in the Cisco Nexus 3000 Series NX-OS Command References.

The following example shows ports 1 through 30 being enabled on Nexus 3232C cluster switches C1 and C2:

```
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
```

Step 16. Enable the first cluster port, e3a, on each node: `network port modify -node` *node-name* `-port` *port-name* `-up-admin true`

```
cluster::*> network port modify -node n1 -port e3a -up-admin true
cluster::*> network port modify -node n2 -port e3a -up-admin true
```

Step 17. Verify that the clusters are up on both nodes: `network port show -role cluster`

```
cluster::*> network port show –role cluster
  (network port show)
Node: n1
                                                                        Ignore
                                                  Speed(Mbps) Health   Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status   Status
--------- ------------ ---------------- ---- ---- ----------- -------- -----
e3a       Cluster      Cluster          up   9000 auto/100000  -
e3b       Cluster      Cluster          up   9000 auto/100000  -        -

Node: n2
                                                                        Ignore
                                                  Speed(Mbps) Health   Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status   Status
--------- ------------ ---------------- ---- ---- ----------- -------- -----
e3a       Cluster      Cluster          up   9000 auto/100000  -
e3b       Cluster      Cluster          up   9000 auto/100000  -

4 entries were displayed.
```

Step 18. For each node, revert all of the migrated cluster interconnect LIFs: `network interface revert -vserver cluster -lif` *lif-name*

You must revert each LIF to its home port individually as shown in the following example:

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

Step 19. Verify that all the LIFs are now reverted to their home ports: `network interface show -role cluster`The **IsHome** column should display a value of true for all of the ports listed in the **CurrentPort** column. If the displayed value is false, the port has not been reverted.

```
cluster::*> network interface show -role cluster
 (network interface show)
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
Cluster
            n1_clus1   up/up      10.10.0.1/24       n1            e3a     true
            n1_clus2   up/up      10.10.0.2/24       n1            e3b     true
            n2_clus1   up/up      10.10.0.3/24       n2            e3a     true
            n2_clus2   up/up      10.10.0.4/24       n2            e3b     true
4 entries were displayed.
```

Step 20. 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
            e3a    C1                   Ethernet1/7      N3K-C3232C
            e3b    n2                   e3b              DM7100F
```

```
n2          /cdp
             e3a    C1                  Ethernet1/8      N3K-C3232C
             e3b    n1                  e3b              DM7100F
```

Step 21. Migrate clus2 to port e3a on the console of each node: `network interface migrate cluster -lif` *lif-name* `-source-node` *source-node-name* `-destination-node` *destination-node-name* `-destination-port` *destination-port-name*

You must migrate each LIF to its home port individually as shown in the following example:

```
cluster::*> network interface migrate -vserver cluster -lif n1_clus2 -source-node n1
–destination-node n1 -destination-port e3a
cluster::*> network interface migrate -vserver cluster -lif n2_clus2 -source-node n2 –destination-node n2
-destination-port e3a
```

Step 22. Shut down cluster ports clus2 LIF on both nodes: `network port modify`

The following example shows the specified ports being set to false, shutting the ports down on both nodes:

```
cluster::*> network port modify -node n1 -port e3b -up-admin false
cluster::*> network port modify -node n2 -port e3b -up-admin false
```

Step 23. Verify the cluster LIF status: `network interface show`

```
cluster::*> network interface show -role cluster
 (network interface show)
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
Cluster
            n1_clus1   up/up      10.10.0.1/24       n1            e3a     true
            n1_clus2   up/up      10.10.0.2/24       n1            e3a     false
            n2_clus1   up/up      10.10.0.3/24       n2            e3a     true
            n2_clus2   up/up      10.10.0.4/24       n2            e3a     false
4 entries were displayed.
```

Step 24. Disconnect the cable from e3b on node n1.You can refer to the running configuration and connect the first 100 GbE port on switch C2 (port 1/7 in this example) to e3b on node n1, using the appropriate cabling for the Nexus 3232C switch model.

Step 25. Disconnect the cable from e3b on node n2.You can refer to the running configuration and connect e3b to the next available 100 GbE port on C2, port 1/8, using the appropriate cabling for the Nexus 3232C switch model.

Step 26. 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:

```
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
```

Step 27. Enable the second cluster port, e3b, on each node: `network port modify`

The following example shows the second cluster port e3b being brought up on each node:

```
cluster::*> network port modify -node n1 -port e3b -up-admin true
cluster::*> network port modify -node n2 -port e3b -up-admin true
```

Step 28. 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
```

Step 29. Verify that all of the cluster interconnect ports are now reverted to their home ports: `network interface show -role cluster`

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)
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
Cluster
            n1_clus1   up/up      10.10.0.1/24       n1            e3a     true
            n1_clus2   up/up      10.10.0.2/24       n1            e3b     true
            n2_clus1   up/up      10.10.0.3/24       n2            e3a     true
            n2_clus2   up/up      10.10.0.4/24       n2            e3b     true
4 entries were displayed.
```

Step 30. Verify that all of the cluster interconnect ports are in the up state: `network port show -role cluster`

Step 31. Display the cluster switch port numbers through which each cluster port is connected to each node: `network device-discovery show`

```
cluster::*> network device-discovery show
            Local  Discovered
Node        Port   Device              Interface        Platform
----------- ------ ------------------ ---------------- ----------------
n1          /cdp
            e3a    C1                  Ethernet1/7      N3K-C3232C
            e3b    C2                  Ethernet1/7      N3K-C3232C
n2          /cdp
            e3a    C1                  Ethernet1/8      N3K-C3232C
            e3b    C2                  Ethernet1/8      N3K-C3232C
```

Step 32. 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      NX3232CV
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     10.10.1.102      NX3232CV
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.
```

Step 33. Verify that switchless cluster detection changed the switchless cluster option to disabled: `network options switchless-cluster show`

Step 34. Ping the remote cluster interfaces and perform an RPC server check: `cluster ping-cluster -node` *node-name*

```
cluster::*> cluster ping-cluster -node n1
Host is n1 Getting addresses from network interface table...
Cluster n1_clus1 n1        e3a    10.10.0.1
```

```
Cluster n1_clus2 n1        e3b    10.10.0.2
Cluster n2_clus1 n2        e3a    10.10.0.3
Cluster n2_clus2 n2        e3b    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 9000 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)
```

Step 35. 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 Lenovo support.

Step 36. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

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

Chapter 4. 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.

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.

Procedure summary

- **Display and migrate the cluster ports to switch C2 (Step 2 Steps 1-7 on page 30)**
- **Reconnect ISL cables from switch CL2 to switch C2, then migrate ISLs to switch CL1 and C2 (Step 9 Steps 8-14 on page 34)**
- **Revert all LIFs to originally assigned ports (Step 16 Steps 15-18 on page 35)**
- **Verify all ports and LIF are correctly migrated (Step 20 Steps 19-21 on page 37)**

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 CL2 or 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 Lenovo Data Center Support 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 Lenovo Press to verify the correct cluster ports for your platform.

Lenovo Press

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.

Step 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`

In `MAINT=xh`, *x* is the duration of the maintenance window in units of hours.

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

Step 2. Display information about the devices in your configuration: `network device-discovery show`

```
cluster::> network device-discovery show
            Local  Discovered
Node        Port   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
            e0d    CL1                  Ethernet1/1/2    N3K-C3232C

n2         /cdp
            e0a    CL1                  Ethernet1/1/3    N3K-C3232C
            e0b    CL2                  Ethernet1/1/3    N3K-C3232C
            e0c    CL2                  Ethernet1/1/4    N3K-C3232C
            e0d    CL1                  Ethernet1/1/4    N3K-C3232C

n3         /cdp
            e4a    CL1                  Ethernet1/7      N3K-C3232C
            e4e    CL2                  Ethernet1/7      N3K-C3232C

n4         /cdp
            e4a    CL1                  Ethernet1/8      N3K-C3232C
            e4e    CL2                  Ethernet1/8      N3K-C3232C
```

```
12 entries were displayed
```

Step 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
                                                                        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: 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.
```

b. Display information about the logical interfaces (LIFs): `network interface show -role cluster`

```
cluster::*> network interface show -role cluster
             Logical    Status     Network           Current       Current Is
Vserver     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            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.
```

Step 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 Lenovo Data Center Support.

b. Download the appropriate version of the RCF.

c. To continue with download, provide the system model and SN for the DM Series Storage that the switch is attached to.

d. Download the correct NX-OS version for your switch from the CISCO support site.

Step 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
```

Step 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)
```

```
            Logical    Status     Network            Current       Current Is
Vserver     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            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.
```

Step 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
```

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

Step 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)#
```

Step 10. Remove all the cables attached to the cluster switch CL2 and reconnect them to the replacement switch C2 for all the nodes.

Step 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.

Step 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)#
```

Step 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)
```

Step 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)
```

Step 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
```

Step 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
```

Step 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)
            Logical    Status     Network            Current       Current Is
Vserver     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.
```

Step 18. Verify that the cluster ports are connected: `network port show -role cluster`

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

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

Step 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
            Local  Discovered
Node        Port   Device              Interface        Platform
```

```
----------- ------ ------------------ --------------- ----------------
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
                                                                        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: 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

            Logical    Status     Network            Current       Current Is
```

```
Vserver     Interface  Admin/Oper Address/Mask       Node           Port    Home
----------- ---------- ---------- ------------------ -------------- ------- ----
Cluster
            nm1_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 3 entries were displayed.
```

Step 21. Delete the replaced cluster switch CL2 if it has not been removed automatically: `system cluster-switch delete -device` *cluster-switch-name*

Step 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.

```
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

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.
```

Step 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 Lenovo technical support.

Step 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
```

Chapter 5. Replacing a Cisco Nexus 3232C storage switch

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

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

- Your system can support Cisco Nexus 3232C storage switches.
- 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. Cisco Nexus 3000 Series Switches
- You must have downloaded the applicable RCFs.

Steps to replace a Cisco Nexus 3232C storage switch

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

The existing network configuration must have the following characteristics:

- Management connectivity must exist on both switches.
- The lastest RCF file for Cisco Ether Switch can be downloaded on Lenovo Data Center Support.

Note: Make sure that all troubleshooting steps have been completed to confirm that your switch needs replacing.

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.

Procedure summary:

- **Confirm the switch to be replaced is S2**
- **Verify all device configuration on switch S1 (Steps 1-5)**
- **Disconnect the cables from switch S2 (Step 6)**
- **Reconnect the cables to switch NS2 (Step 7)**
- **Verify all device configurations on switch NS2 (Steps 8-10)**

- This procedure replaces the second Nexus 3232C storage switch S2 with the new 3232C switch NS2

  Note: You may not see S2 relative information if that switch was totally damaged.
- The two nodes are node1 and node2.

Step 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`

In `MAINT=xh`, *x* is the duration of the maintenance window in units of hours.

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

Step 2. Check on the health status of the storage node ports to make sure that there is connection to storage switch S1: `storage port show -port-type ENET`

```
storage::*> storage port show -port-type ENET
                                    Speed                      VLAN
Node                Port Type  Mode    (Gb/s) State    Status      ID
------------------- ---- ----- ------- ------ -------- --------- ----
node1
                    e4a  ENET  storage    100 enabled  online      30
                    e4b  ENET  storage    100 enabled  online     30
node2
                    e4a  ENET  storage    100 enabled  online      30
                    e4b  ENET  storage    100 enabled  online     30
```

Step 3. Verify that storage switch S1 is available: `network device-discovery show`

```
storage::*> network device-discovery show
Node/       Local  Discovered
Protocol    Port   Device (LLDP: ChassisID)  Interface          Platform
----------- ------ ------------------------- ----------------- ----------------
node1/cdp
            e3a    node2                     e3a                DM7100F
            e3b    node2                     e3b                DM7100F
            e4a    S1                        Ethernet1/1       NX3232C
            e4b    S2                        Ethernet1/1       NX3232C
node1/lldp
            e3a    node2                     e3a                -
            e3b    node2                     e3b                -
            e4a    S1                        Ethernet1/1       -
            e4b    S2                        Ethernet1/1       -
node2/cdp
            e3a    node1                     e3a                DM7100F
            e3b    node1                     e3b                DM7100F
            e4a    S1                        Ethernet1/2       NX3232C
            e4b    S2                        Ethernet1/2       NX3232C
node2/lldp
            e3a    node1                     e3a                -
            e3b    node1                     e3b                -
            e4a    S1                        Ethernet1/2       -
            e4b    S2                        Ethernet1/2       -
```

Step 4. Run the `show lldp neighbors` command on the working switch to confirm that you can see both nodes and all shelves: `show lldp neighbors`

```
S1# show lldp neighbors
Capability codes:
  (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
  (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other
Device ID               Local Intf      Hold-time  Capability  Port ID
node1                    Eth1/1          121        S           e4a
node2                    Eth1/2          121        S           e4a
SHFGD2008000011          Eth1/5          121        S           e0a
SHFGD2008000011          Eth1/6          120        S           e0a
SHFGD2008000022          Eth1/7          120        S           e0a
SHFGD2008000022          Eth1/8          120        S           e0a
```

Step 5. Verify the shelf ports in the storage system: `storage shelf port show -fields remote-device,remote-port`

```
storage::*> storage shelf port show -fields remote-device,remote-port

shelf  id  remote-port  remote-device
-----  --  -----------  -------------
3.20   0   Ethernet1/5  S1
3.20   1   Ethernet1/5  S2
3.20   2   Ethernet1/6  S1
3.20   3   Ethernet1/6  S2
3.30   0   Ethernet1/7  S1
3.20   1   Ethernet1/7  S2
3.30   2   Ethernet1/8  S1
3.20   3   Ethernet1/8  S2
```

Step 6. Remove all cables attached to storage switch S2.

Step 7. Reconnect all cables to the replacement switch NS2.

Step 8. Recheck the health status of the storage node ports: `storage port show -port-type ENET`

```
storage::*> storage port show -port-type ENET
                                     Speed                          VLAN
Node                Port Type  Mode    (Gb/s) State    Status        ID
------------------- ---- ----- ------- ------ -------- ------------ ---
node1
                    e4a  ENET  storage    100 enabled  online         30
                    e4b  ENET  storage    100 enabled  online         30
node2
                    e4a  ENET  storage    100 enabled  online         30
                    e4b  ENET  storage    100 enabled  online         30
```

Step 9. Verify that both switches are available: `network device-discovery show`

```
storage::*> network device-discovery show
Protocol    Port   Device (LLDP: ChassisID)  Interface          Platform
----------- ------ ------------------------- ----------------  --------
node1/cdp
            e3a    node2                     e3a                DM7100F
            e3b    node2                     e3b                DM7100F
            e4a    S1                        Ethernet1/1       NX3232C
            e4b    NS2                       Ethernet1/1       NX3232C
node1/lldp
            e3a    node2                     e3a                -
            e3b    node2                     e3b                -
            e4a    S1                        Ethernet1/1       -
            e4b    NS2                       Ethernet1/1       -
node2/cdp
            e3a    node1                     e3a                DM7100F
            e3b    node1                     e3b                DM7100F
            e4a    S1                        Ethernet1/2       NX3232C
            e4b    NS2                       Ethernet1/2       NX3232C
node2/lldp
            e3a    node1                     e3a                -
            e3b    node1                     e3b                -
            e4a    S1                        Ethernet1/2       -
            e4b    NS2                       Ethernet1/2       -
```

Step 10. Verify the shelf ports in the storage system: `storage shelf port show -fields remote-device,remote-port`

```
storage::*> storage shelf port show -fields remote-device,remote-port
shelf id remote-port remote-device
----- -- ----------- -------------
3.20  0  Ethernet1/5 S1
3.20  1  Ethernet1/5 NS2
3.20  2  Ethernet1/6 S1
3.20  3  Ethernet1/6 NS2
3.30  0  Ethernet1/7 S1
3.20  1  Ethernet1/7 NS2
3.30  2  Ethernet1/8 S1
3.20  3  Ethernet1/8 NS2
```

Step 11. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=END`

Appendix A. Contacting Support

You can contact Support to obtain help for your issue.

You can receive hardware service through a Lenovo Authorized Service Provider. To locate a service provider authorized by Lenovo to provide warranty service, go to https://datacentersupport.lenovo.com/serviceprovider and use filter searching for different countries. For Lenovo support telephone numbers, see https://datacentersupport.lenovo.com/supportphonelist for your region support details.

Appendix B. Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area.

Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document is not an offer and does not provide a license under any patents or patent applications. You can send inquiries in writing to the following:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Trademarks