



Switch documentation for ONTAP hardware systems

Cluster and storage switches

NetApp
April 25, 2025

Table of Contents

- Switch documentation for ONTAP hardware systems 1
- Get started 2
 - What's new for switches 2
 - New switch support 2
 - Learn about Cluster, Storage, and Shared switches 2
 - Cluster switches 2
 - Storage switches 2
 - Shared switches 2
 - End-of-availability 2
 - Get up and running with Cluster, Storage, and Shared switches 3
- Cluster switches 6
 - Broadcom-supported BES-53248 6
 - Overview 6
 - Install hardware 10
 - Configure software 15
 - Upgrade switch 105
 - Migrate switches 149
 - Replace switches 184
 - Cisco Nexus 9336C-FX2 or 9336C-FX2-T 211
 - Overview 211
 - Install hardware 216
 - Configure software 229
 - Migrate switches 296
 - Replace switches 348
 - NVIDIA SN2100 380
 - Overview 380
 - Install hardware 383
 - Configure software 392
 - Migrate switches 455
 - Replace switches 494
- Storage switches 522
 - Cisco Nexus 9336C-FX2 522
 - Overview 522
 - Install hardware 525
 - Configure software 536
 - Replace a Cisco Nexus 9336C-FX2 storage switch 574
 - NVIDIA SN2100 579
 - Overview 579
 - Install hardware 582
 - Configure software 591
 - Migrate switches 610
 - Replace a NVIDIA SN2100 storage switch 620
- Shared switches 628

Cisco Nexus 9336C-FX2	628
Overview	628
Install hardware	630
Configure software	640
Migrate switches	703
Replace a Cisco Nexus 9336C-FX2 shared switch	742
Monitor switch health	761
Overview of switch health monitor	761
Configure switch health monitoring	761
Configuration overview	761
Configure log collection	761
Optional: Configure SNMPv3 for your switch	768
Check switch health	786
Health check overview	786
Manage the monitoring of Ethernet Switches	786
Verify the monitoring of Ethernet switches	787
Troubleshoot alerts	788
Log collection	789
Log collection overview	789
Troubleshoot log collection	789
End-of-availability switches	790
End-of-availability	790
End-of-sale and End-of-life announcements	790
Cisco Nexus 3232C	790
Overview	790
Install hardware	793
Configure software	804
Migrate switches	855
Replace switches	951
Cisco Nexus 3132Q-V	992
Overview	992
Install hardware	995
Configure software	1002
Migrate switches	1050
Replace switches	1152
Cisco Nexus 92300YC	1196
Overview	1196
Install hardware	1200
Configure software	1210
Migrate switches	1221
Replace switches	1260
NetApp CN1610	1291
Overview of installation and configuration for NetApp CN1610 switches	1291
Install and configure workflow for NetApp CN1610 switches	1292
Documentation requirements for NetApp CN1610 switches	1292

Install and configure	1293
Migrate switches	1329
Replace switches	1356
Legal notices	1382
Copyright	1382
Trademarks	1382
Patents	1382
Privacy policy	1382

Switch documentation for ONTAP hardware systems

Get started

What's new for switches

Learn about the new switches for FAS and AFF systems. For additional switch information, see the [NetApp Hardware Universe](#).

New switch support

Unresolved directive in whats-new.adoc - include::.../_include/new-switch-support.adoc[]

Learn about Cluster, Storage, and Shared switches

NetApp offers cluster, storage, and shared switches that deliver internal communications with the ability to non-disruptively move data and network interfaces across the cluster.

The “front-end” switches provide connectivity to host storage, while the “back-end” cluster switches provide connections between two or more NetApp controllers.



Only NetApp-validated back-end switches (ordered from NetApp) are supported.

Cluster switches

Cluster switches allow you to build ONTAP clusters with more than two nodes. NetApp-supported cluster switches include:

- Broadcom BES-53248
- Cisco Nexus 9336C-FX2
- NVIDIA SN2100

Storage switches

Storage switches allow you to route data between servers and storage arrays in a Storage Area Network (SAN). NetApp-supported storage switches include:

- Cisco Nexus 9336C-FX2
- NVIDIA SN2100

Shared switches

Shared switches allow you to combine cluster and storage functionality into a shared switch configuration, by supporting the use of shared cluster and storage RCFs. The NetApp-supported shared switch is:

- Cisco Nexus 9336C-FX2

End-of-availability

The following storage switches are no longer available for purchase, but are still supported:

- Cisco Nexus 3232C
- Cisco Nexus 3132Q-V
- Cisco Nexus 92300YC
- NetApp CN1610

Get up and running with Cluster, Storage, and Shared switches

To get up and running with cluster, storage, and shared switches, you install hardware components and configure your switch.

Deploying the switch involves the following workflow.

1

Install the switch hardware

Install your switches in the rack or cabinet. Access the following instructions for your switch model.

Cluster switches

- [Install BES-53248 switch](#)
- [Install Cisco Nexus 9336C-FX2 switch](#)
- [Install NVIDIA SN2100 switch](#)

Storage switches

- [Install Cisco Nexus 9336C-FX2 switch](#)
- [Install NVIDIA SN2100 switch](#)

Shared switches

- [Install Cisco Nexus 9336C-FX2 switch](#)

2

Install the storage system and storage shelves

Install your storage system and storage shelves in the cabinet or rack. Access the install and setup instructions for your platform model.

- [AFF systems](#)
- [ASA systems](#)
- [ASA r2 systems](#)
- [FAS systems](#)

3

Cable the switches to the controllers

The ONTAP systems install and setup instructions include instructions for cabling the controller ports to the switch. However, if you need lists of supported cables and transceivers, and detailed information about the host ports for your switch, access the following instructions for your switch model.

Cluster switches

- [Cable BES-53248 switch](#)
- [Cable Cisco Nexus 9336C-FX2 switch](#)
- [Cable NVIDIA SN2100 switch](#)

Storage switches

- [Cable Cisco Nexus 9336C-FX2 switch](#)
- [Cable NVIDIA SN2100 switch](#)

Shared switches

- [Cable Cisco Nexus 9336C-FX2 switch](#)

4

Configure switch

Perform an initial setup of your switches. Access the following instructions for your switch model.

Cluster switches

- [Configure BES-53248 switch](#)
- [Configure Cisco Nexus 9336C-FX2 switch](#)
- [Configure NVIDIA SN2100 switch](#)

Storage switches

- [Configure Cisco Nexus 9336C-FX2 switch](#)
- [Configure NVIDIA SN2100 switch](#)

Shared switches

- [Configure Cisco Nexus 9336C-FX2 switch](#)

5

Install switch software

To install and configure the software on your switch, follow the software install workflow for your switch model.

Cluster switches

- [Install software for BES-53248 switches](#)
- [Install software for Cisco Nexus 9336C-FX2 switch](#)
- [Install software for NVIDIA SN2100 switch](#)

Storage switches

- [Install software for Cisco Nexus 9336C-FX2 switch](#)
- [Install software for NVIDIA SN2100 switch](#)

Shared switches

- [Install software for Cisco Nexus 9336C-FX2 switch](#)

6

Complete system setup

After you have configured your switches and installed the required software, access the install and setup instructions for your ONTAP systems platform model to complete your system setup.

- [AFF systems](#)
- [ASA systems](#)
- [ASA r2 systems](#)
- [FAS systems](#)

7

Complete ONTAP configuration

After you have installed and set up your ONTAP systems controllers and switches, you must complete configuring your storage in ONTAP. Access the following instructions according to your deployment configuration.

- For ONTAP deployments, see [Configure ONTAP](#).
- For ONTAP with MetroCluster deployments, see [Configure Metrocluster with ONTAP](#).

Cluster switches

Broadcom-supported BES-53248

Overview

Overview of installation and configuration for BES-53248 switches

The BES-53248 is a bare metal switch designed to work in ONTAP clusters ranging from two to 24 nodes.

Initial configuration overview

To initially configure a BES-53248 cluster switch on systems running ONTAP, follow these steps:

1. [Install the hardware for the BES-53248 cluster switch.](#)

Instructions are available in the *Broadcom-supported BES-53248 Cluster Switch Installation Guide*.

2. [Configure the BES-53248 cluster switch.](#)

Perform an initial setup of the BES-53248 cluster switch.

3. [Install the EFOS software.](#)

Download and install the Ethernet Fabric OS (EFOS) software on the BES-53248 cluster switch.

4. [Install licenses for BES-53248 cluster switches.](#)

Optionally, add new ports by purchasing and installing more licenses. The switch base model is licensed for 16 10GbE or 25GbE ports and two 100GbE ports.

5. [Install the Reference Configuration File \(RCF\).](#)

Install or upgrade the RCF on the BES-53248 cluster switch, and then verify the ports for an additional license after the RCF is applied.

6. [Enable SSH on BES-53248 cluster switches.](#)

If you use the Ethernet Switch Health Monitor (CSHM) and log collection features, enable SSH on the switches.

Additional information

Before you begin installation or maintenance, be sure to review the following:

- [Configuration requirements](#)
- [Components and part numbers](#)
- [Required documentation](#)

Configuration requirements for BES-53248 cluster switches

For BES-53248 switch installation and maintenance, be sure to review EFOS and ONTAP support and configuration requirements.

EFOS and ONTAP support

See the [NetApp Hardware Universe](#) and [Broadcom switches compatibility matrix](#) for EFOS and ONTAP compatibility information with BES-53248 switches. EFOS and ONTAP support can vary by the specific machine type of the BES-53248 switch. For details of all BES-53248 switch machine types, see [Components and part numbers for BES-53248 cluster switches](#).

Configuration requirements

To configure a cluster, you need the appropriate number and type of cables and cable connectors for the cluster switches. Depending on the type of cluster switch you are initially configuring, you need to connect to the switch console port with the included console cable.

Cluster switch port assignments

You can use the Broadcom-supported BES-53248 cluster switch port assignments table as a guide to configuring your cluster.

Switch ports	Ports usage
01-16	10/25GbE cluster port nodes, base configuration
17-48	10/25GbE cluster port nodes, with licenses
49-54	40/100GbE cluster port nodes, with licenses, added right to left
55-56	100GbE cluster Inter-Switch Link (ISL) ports, base configuration

See the [Hardware Universe](#) for more information on switch ports.

Port group speed constraint

- On BES-53248 cluster switches, the 48 10/25GbE (SFP28/SFP+) ports are combined into 12 x 4-port groups as follows: Ports 1-4, 5-8, 9-12, 13-16, 17-20, 21-24, 25-28, 29-32, 33-36, 37-40, 41-44, and 45-48.
- The SFP28/SFP+ port speed must be the same (10GbE or 25GbE) across all ports in the 4-port group.

Additional requirements

- If you purchase additional licenses, see [Activate newly licenses ports](#) for details on how to activate them.
- If SSH is active, you must re-enable it manually after running the command `erase startup-config` and rebooting the switch.

Components and part numbers for BES-53248 cluster switches

For BES-53248 switch installation and maintenance, be sure to review the list of components and part numbers.

The following table lists the part number, description, and minimum EFOS and ONTAP versions for the BES-53248 cluster switch components, including rack-mount rail kit details.



A minimum EFOS version of **3.10.0.3** is required for part numbers **X190005-B** and **X190005R-B**.

Part number	Description	Minimum EFOS version	Minimum ONTAP version
X190005-B	BES-53248-B/IX8, CLSW, 16PT10/25GB, PTSX (PTSX = Port Side Exhaust)	3.10.0.3	9.8
X190005R-B	BES-53248-B/IX8, CLSW, 16PT10/25GB, PSIN (PSIN = Port Side Intake)	3.10.0.3	9.8
X190005	BES-53248, CLSW, 16Pt10/25GB, PTSX, BRDCM SUPP	3.4.4.6	9.5P8
X190005R	BES-53248, CLSW, 16Pt10/25GB, PSIN, BRDCM SUPP	3.4.4.6	9.5P8
X-RAIL-4POST-190005	Rack mount rail kit Ozeki 4 post 19"	N/A	N/A



Note the following information with regards to machine types:

Machine type	Minimum EFOS version
BES-53248A1	3.4.4.6
BES-53248A2	3.10.0.3
BES-53248A3	3.10.0.3

You can determine your specific machine type by using the command: `show version`

Show example

```
(cs1)# show version

Switch: cs1

System Description..... EFOS, 3.10.0.3, Linux
5.4.2-b4581018, 2016.05.00.07
Machine Type..... BES-53248A3
Machine Model..... BES-53248
Serial Number..... QTWCU225xxxxx
Part Number..... 1IX8BZxxxxx
Maintenance Level..... a3a
Manufacturer..... QTMC
Burned In MAC Address..... C0:18:50:F4:3x:xx
Software Version..... 3.10.0.3
Operating System..... Linux 5.4.2-b4581018
Network Processing Device..... BCM56873_A0
.
.
.
```

Documentation requirements for BES-53248 cluster switches

For BES-53248 switch installation and maintenance, be sure to review the specific switch and controller documentation.

Broadcom documentation

To set up the BES-53248 cluster switch, you need the following documents available from the Broadcom Support Site: [Broadcom Ethernet Switch Product Line](#)

Document title	Description
<i>EFOS Administrator's Guide v3.4.3</i>	Provides examples of how to use the BES-53248 switch in a typical network.
<i>EFOS CLI Command Reference v3.4.3</i>	Describes the command-line interface (CLI) commands you use to view and configure the BES-53248 software.
<i>EFOS Getting Started Guide v3.4.3</i>	Provides detailed information about for the BES-53248 switch.
<i>EFOS SNMP Reference Guide v3.4.3</i>	Provides examples of how to use the BES-53248 switch in a typical network.

Document title	Description
<i>EFOS Scaling Parameters and Values v3.4.3</i>	Describes the default scaling parameters with which EFOS software is delivered and validated on the supported platforms.
<i>EFOS Functional Specifications v3.4.3</i>	Describes the specifications for the EFOS software on the supported platforms.
<i>EFOS Release Notes v3.4.3</i>	Provides release-specific information about BES-53248 software.
<i>Cluster Network and Management Network Compatibility Matrix</i>	Provides information on network compatibility. The matrix is available from the BES-53248 switch download site at Broadcom cluster switches .

ONTAP systems documentation and KB articles

To set up an ONTAP system, you need the following documents from the NetApp Support Site at mysupport.netapp.com or the Knowledgebase (KB) site at kb.netapp.com.

Name	Description
NetApp Hardware Universe	Describes the power and site requirements for all NetApp hardware, including system cabinets, and provides information on the relevant connectors and cable options to use along with their part numbers.
<i>Controller-specific Installation and Setup Instructions</i>	Describes how to install NetApp hardware.
ONTAP 9	Provides detailed information about all aspects of the ONTAP 9 release.
<i>How to add additional port licensing for the Broadcom-supported BES-53248 switch</i>	Provides detailed information on adding port licenses. Go to the KB article .

Install hardware

Install the hardware for the BES-53248 cluster switch

To install the BES-53248 hardware, refer to Broadcom's documentation.

Steps

1. Review the [configuration requirements](#).
2. Follow the instructions in the [Broadcom-supported BES-53248 Cluster Switch Installation Guide](#).

What's next?

[Configure the switch](#).

Configure the BES-53248 cluster switch

Follow these steps to perform an initial setup of the BES-53248 cluster switch.

Before you begin

- Hardware is installed, as described in [Install the hardware](#).
- You have reviewed the following:
 - [Configuration requirements](#)
 - [Components and part numbers](#)
 - [Documentation requirements](#)

About the examples

The examples in the configuration procedures use the following switch and node nomenclature:

- The NetApp switch names are `cs1` and `cs2`. The upgrade starts on the second switch, `cs2`.
- The cluster LIF names are `node1_clus1` and `node1_clus2` for node1, and `node2_clus1` and `node2_clus2` for node2.
- The IPspace name is `Cluster`.
- The `cluster1: :>` prompt indicates the name of the cluster.
- The cluster ports on each node are named `e0a` and `e0b`. See the [NetApp Hardware Universe](#) for the actual cluster ports supported on your platform.
- The Inter-Switch Links (ISLs) supported for the NetApp switches are ports `0/55` and `0/56`.
- The node connections supported for the NetApp switches are ports `0/1` through `0/16` with default licensing.
- The examples use two nodes, but you can have up to 24 nodes in a cluster.

Steps

1. Connect the serial port to a host or serial port.
2. Connect the management port (the RJ-45 wrench port on the left side of the switch) to the same network where your TFTP server is located.
3. At the console, set the host-side serial settings:
 - 115200 baud
 - 8 data bits
 - 1 stop bit
 - parity: none
 - flow control: none
4. Log in to the switch as `admin` and press **Enter** when prompted for a password. The default switch name is **routing**. At the prompt, enter `enable`. This gives you access to Privileged EXEC mode for switch configuration.

```
User: admin
Password:
(Routing) > enable
Password:
(Routing) #
```

5. Change the switch name to **cs2**.

```
(Routing) # hostname cs2
(cs2) #
```

6. To set a static IP address, use the `serviceport protocol`, `network protocol`, and `serviceport ip` commands as shown in the example.

The `serviceport` is set to use DHCP by default. The IP address, subnet mask, and default gateway address are assigned automatically.

```
(cs2) # serviceport protocol none
(cs2) # network protocol none
(cs2) # serviceport ip ipaddr netmask gateway
```

7. Verify the results using the command:

```
show serviceport
```

Show example

```
(cs2) # show serviceport
Interface Status..... Up
IP Address..... 172.19.2.2
Subnet Mask..... 255.255.255.0
Default Gateway..... 172.19.2.254
IPv6 Administrative Mode..... Enabled
IPv6 Prefix is .....
fe80::dac4:97ff:fe71:123c/64
IPv6 Default Router.....
fe80::20b:45ff:fea9:5dc0
Configured IPv4 Protocol..... DHCP
Configured IPv6 Protocol..... None
IPv6 AutoConfig Mode..... Disabled
Burned In MAC Address..... D8:C4:97:71:12:3C
```


8. Configure the domain and name server:

```
ip domain name <domain_name>
ip name server <server_name>
```

```
(cs2)# configure
(cs2) (Config)# ip domain name company.com
(cs2) (Config)# ip name server 10.10.99.1 10.10.99.2
(cs2) (Config)# exit
(cs2)#
```

9. Configure the NTP server.

EFOS 3.10.0.3 and later

Configure the time zone and time synchronization (NTP):

```
sntp server <server_name>
clock
```

```
(cs2)# configure
(cs2) (Config)# ntp server 10.99.99.5
(cs2) (Config)# clock timezone -7
(cs2) (Config)# exit
(cs2)#
```

EFOS 3.9.0.2 and earlier

Configure the time zone and time synchronization (SNTP):

```
sntp client mode <client_mode>
sntp server <server_name>
clock
```

```
(cs2)# configure
(cs2) (Config)# sntp client mode unicast
(cs2) (Config)# sntp server 10.99.99.5
(cs2) (Config)# clock timezone -7
(cs2) (Config)# exit
(cs2)#
```

10. Configure the time manually if you did not configure an NTP server in the previous step.

EFOS 3.10.0.3 and later

Configure the time manually.

```
clock
```

```
(cs2)# configure
(cs2)(Config)# clock summer-time recurring 1 sun mar 02:00 1 sun nov
02:00 offset 60 zone EST
(cs2)(Config)# clock timezone -5 zone EST
(cs2)(Config)# clock set 07:00:00
(cs2)(Config)# clock set 10/20/2023
(cs2)(Config)# show clock

07:00:11 EST(UTC-5:00) Oct 20 2023
No time source

(cs2)(Config)# exit
(cs2)#
```

EFOS 3.9.0.2 and earlier

Configure the time manually.

```
clock
```

```
(cs2)# configure
(cs2)(Config)# no sntp client mode
(cs2)(Config)# clock summer-time recurring 1 sun mar 02:00 1 sun nov
02:00 offset 60 zone EST
(cs2)(Config)# clock timezone -5 zone EST
(cs2)(Config)# clock set 07:00:00
(cs2)(Config)# clock set 10/20/2023
(cs2)(Config)# show clock

07:00:11 EST(UTC-5:00) Oct 20 2023
No time source

(cs2)(Config)# exit
(cs2)#
```

11. Save the running configuration to the startup configuration:

```
write memory
```

```
(cs2)# write memory
```

```
This operation may take a few minutes.  
Management interfaces will not be available during this time.
```

```
Are you sure you want to save? (y/n) y
```

```
Config file 'startup-config' created successfully.
```

```
Configuration Saved!
```

What's next?

[Install the EFOS software](#)

Configure software

Software install workflow for BES-53248 switches

To initially install and configure the software for a BES-53248 cluster switch, follow these steps:

1. [Install the EFOS software.](#)

Download and install the Ethernet Fabric OS (EFOS) software on the BES-53248 cluster switch.

2. [Install licenses for BES-53248 cluster switches.](#)

Optionally, add new ports by purchasing and installing more licenses. The switch base model is licensed for 16 10GbE or 25GbE ports and two 100GbE ports.

3. [Install the Reference Configuration File \(RCF\).](#)

Install or upgrade the RCF on the BES-53248 cluster switch, and then verify the ports for an additional license after the RCF is applied.

4. [Enable SSH on BES-53248 cluster switches.](#)

If you use the Ethernet Switch Health Monitor (CSHM) and log collection features, enable SSH on the switches.

Install the EFOS software

Follow these steps to install the Ethernet Fabric OS (EFOS) software on the BES-53248 cluster switch.

EFOS software includes a set of advanced networking features and protocols for developing Ethernet and IP infrastructure systems. This software architecture is suitable for any network organizational device using applications that require thorough packet inspection or separation.

Prepare for installation

Before you begin

- This procedure is only suitable for new installations.
- Download the applicable Broadcom EFOS software for your cluster switches from the [Broadcom Ethernet Switch Support](#) site.
- Ensure the [BES-53248 cluster switch is configured](#).

Install the software

Use one of the following methods to install the EFOS software:

- [Method 1: Install EFOS](#). Use for most cases.
- [Method 2: Install EFOS in ONIE mode](#). Use if one EFOS version is FIPS compliant and the other EFOS version is non-FIPS compliant.

Method 1: Install EFOS

Perform the following steps to install the EFOS software.

Steps

1. Log in to the switch serial console port or connect with SSH.
2. Use the `ping` command to verify connectivity to the server hosting EFOS, licenses, and the RCF file.

Show example

This example verifies that the switch is connected to 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.
```

3. Download the image file to the switch.

Check the following table for information on supported copy protocols:

Protocol	Prerequisite
Trivial File Transfer Protocol (TFTP)	None
SSH File Transfer Protocol (SFTP)	Your software package must support secure management
FTP	Password required
XMODEM	None

YMODEM	None
ZMODEM	None
Secure Copy Protocol (SCP)	Your software package must support secure management
HTTP	CLI-based file transfers supported on selected platforms when a native WGET utility is available
HTTPS	CLI-based file transfers supported on selected platforms when a native WGET utility is available

Copying the image file to the active image means that when you reboot, that image establishes the running EFOS version. The previous image remains available as a backup.

Show example

```
(cs2)# copy sftp://root@172.19.2.1//tmp/EFOS-3.10.0.3.stk active
Remote Password:**

Mode..... SFTP
Set Server IP..... 172.19.2.1
Path..... //tmp/
Filename..... EFOS-3.10.0.3.stk
Data Type..... Code
Destination Filename..... active

Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y
SFTP Code transfer starting...

File transfer operation completed successfully.
```

4. Display the boot images for the active and backup configuration:

```
show bootvar
```

Show example

```
(cs2)# show bootvar
```

```
Image Descriptions
```

```
active :
```

```
backup :
```

```
Images currently available on Flash
```

```
-----  
unit      active      backup      current-active      next-active  
-----  
1         3.7.0.4      3.7.0.4      3.7.0.4              3.10.0.3
```

5. Reboot the switch:

```
reload
```

Show example

```
(cs2)# reload
```

```
The system has unsaved changes.
```

```
Would you like to save them now? (y/n) y
```

```
Config file 'startup-config' created successfully .
```

```
Configuration Saved!
```

```
System will now restart!
```

6. Log in again and verify the new version of the EFOS software:

```
show version
```

Show example

```
(cs2)# show version

Switch: 1

System Description..... BES-53248A1,
3.10.0.3, Linux 4.4.211-28a6fe76, 2016.05.00.04
Machine Type..... BES-53248A1,
Machine Model..... BES-53248
Serial Number..... QTFCU38260023
Maintenance Level..... A
Manufacturer..... 0xbc00
Burned In MAC Address..... D8:C4:97:71:0F:40
Software Version..... 3.10.0.3
Operating System..... Linux 4.4.211-
28a6fe76
Network Processing Device..... BCM56873_A0
CPLD Version..... 0xff040c03

Additional Packages..... BGP-4
..... QOS
..... Multicast
..... IPv6
..... Routing
..... Data Center
..... OpEN API
..... Prototype Open API
```

7. Complete the installation.
Follow these five steps to reconfigure the switch:
 - a. [Install licenses](#)
 - b. [Install the RCF file](#)
 - c. [Enable SSH](#)
 - d. [Enable log collection](#)
 - e. [Configure SNMPv3 for monitoring](#)
8. Repeat steps 1 to 7 on the partner switch.

Method 2: Install EFOS in ONIE mode

You can perform the following steps if one EFOS version is FIPS compliant and the other EFOS version is non-FIPS compliant. These steps can be used to install the non-FIPS or FIPS compliant EFOS 3.7.x.x image from ONIE if the switch fails to boot.

Show example

```
+-----+
-+
|*ONIE: Install OS
|
|  ONIE: Rescue
|
|  ONIE: Uninstall OS
|
|  ONIE: Update ONIE
|
|  ONIE: Embed ONIE
|
|  DIAG: Diagnostic Mode
|
|  DIAG: Burn-In Mode
|
|
|
|
|
|
|
|
|
|
+-----+
-+
```

The switch boots into ONIE installation mode.

3. Stop the ONIE discovery and configure the Ethernet interface.

When the following message appears, press **Enter** to invoke the ONIE console:

```
Please press Enter to activate this console. Info: eth0: Checking
link... up.
ONIE:/ #
```



The ONIE discovery continues and messages are printed to the console.

```
Stop the ONIE discovery
ONIE:/ # onie-discovery-stop
discover: installer mode detected.
Stopping: discover... done.
ONIE:/ #
```

4. Configure the Ethernet interface of the switch management port and add the route using `ifconfig eth0 <ipAddress> netmask <netmask> up` and `route add default gw <gatewayAddress>`

```
ONIE:/ # ifconfig eth0 10.10.10.10 netmask 255.255.255.0 up
ONIE:/ # route add default gw 10.10.10.1
```

5. Verify that the server hosting the ONIE installation file is reachable:

```
ping
```

Show example

```
ONIE:/ # ping 50.50.50.50
PING 50.50.50.50 (50.50.50.50): 56 data bytes
64 bytes from 50.50.50.50: seq=0 ttl=255 time=0.429 ms
64 bytes from 50.50.50.50: seq=1 ttl=255 time=0.595 ms
64 bytes from 50.50.50.50: seq=2 ttl=255 time=0.369 ms
^C
--- 50.50.50.50 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.369/0.464/0.595 ms
ONIE:/ #
```

6. Install the new switch software:

```
ONIE:/ # onie-nos-install http://50.50.50.50/Software/onie-installer-x86\_64
```

Show example

```
ONIE:/ # onie-nos-install http://50.50.50.50/Software/onie-
installer-x86_64
discover: installer mode detected.
Stopping: discover... done.
Info: Fetching http://50.50.50.50/Software/onie-installer-3.7.0.4
...
Connecting to 50.50.50.50 (50.50.50.50:80)
installer          100% |*****| 48841k
0:00:00 ETA
ONIE: Executing installer: http://50.50.50.50/Software/onie-
installer-3.7.0.4
Verifying image checksum ... OK.
Preparing image archive ... OK.
```

The software installs and then reboots the switch. Let the switch reboot normally into the new EFOS version.

7. Log in and verify that the new switch software is installed:

```
show bootvar
```

Show example

```
(cs2)# show bootvar
Image Descriptions
active :
backup :
Images currently available on Flash
----  -----  -----  -----  -----
unit   active      backup    current-active  next-active
----  -----  -----  -----  -----
  1    3.7.0.4     3.7.0.4   3.7.0.4         3.10.0.3
(cs2) #
```

8. Complete the installation.
The switch reboots with no configuration applied and resets to factory defaults. Follow these six steps to reconfigure the switch:
 - a. [Configure switch](#)
 - b. [Install licenses](#)
 - c. [Install the RCF file](#)
 - d. [Enable SSH](#)

- e. [Enable log collection](#)
 - f. [Configure SNMPv3 for monitoring](#)
9. Repeat steps 1 to 8 on the partner switch.

Install the Reference Configuration File (RCF) and license file

Beginning with EFOS 3.12.0.1, you can install the Reference Configuration File (RCF) and license file after configuring the BES-53248 cluster switch.



All ports are configured when you install the RCF, but you need to install your license to activate the configured ports.

Review requirements

Before you begin

Verify that the following are in place:

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).
- The current RCF, available from the [Broadcom Cluster Switches](#) page.
- A boot configuration in the RCF that reflects the desired boot images, required if you are installing only EFOS 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.
- A console connection to the switch, required when installing the RCF from a factory-default state. This requirement is optional if you have used the Knowledge Base article [How to clear configuration on a Broadcom interconnect switch while retaining remote connectivity](#) to clear the configuration, beforehand.

Suggested documentation

Consult the switch compatibility table for the supported ONTAP and RCF versions. See the [EFOS Software download](#) page. Note that there can be command dependencies between the command syntax in the RCF and that found in versions of EFOS.

Install the configuration file

About the examples

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

- The names of the two BES-53248 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.
- The examples in this procedure use four nodes. These nodes use two 10GbE cluster interconnect ports e0a and e0b. See the [Hardware Universe](#) to verify the correct cluster ports on your platforms.



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

About this task

The procedure requires the use of both ONTAP commands and Broadcom switch commands; ONTAP commands are used unless otherwise indicated.

No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure migrates all the cluster LIFs to the operational partner switch while performing the steps on the target switch.



Before installing a new switch software version and RCFs, use the Knowledge Base article [How to clear configuration on a Broadcom interconnect switch while retaining remote connectivity](#). If you must erase the switch settings completely, then you need to perform the basic configuration again. You must be connected to the switch using the serial console because a complete configuration erasure resets the configuration of the management network.

Step 1: Prepare for the installation

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
```

where *x* is the duration of the maintenance window in hours.



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

The following command suppresses automatic case creation for two hours:

```
cluster1::*> system node autosupport invoke -node \* -type all -message MAINT=2h
```

2. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (***>**) appears.

3. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
              e0a    cs1                      0/2          BES-
53248
              e0b    cs2                      0/2          BES-
53248
cluster1-02/cdp
              e0a    cs1                      0/1          BES-
53248
              e0b    cs2                      0/1          BES-
53248
cluster1-03/cdp
              e0a    cs1                      0/4          BES-
53248
              e0b    cs2                      0/4          BES-
53248
cluster1-04/cdp
              e0a    cs1                      0/3          BES-
53248
              e0b    cs2                      0/3          BES-
53248
cluster1::*>
```

4. 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
```

Show example

```
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					
e0b	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					
e0b	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
```

```
Health Health Speed (Mbps)
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
```


Show example

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

Current Vserver Port	Logical Interface Home	Current Is	Status Admin/Oper	Network Address/Mask	Node
Cluster					
cluster1-01	cluster1-01_clus1	e0a true	up/up	169.254.3.4/23	
cluster1-01	cluster1-01_clus2	e0b true	up/up	169.254.3.5/23	
cluster1-02	cluster1-02_clus1	e0a true	up/up	169.254.3.8/23	
cluster1-02	cluster1-02_clus2	e0b true	up/up	169.254.3.9/23	
cluster1-03	cluster1-03_clus1	e0a true	up/up	169.254.1.3/23	
cluster1-03	cluster1-03_clus2	e0b true	up/up	169.254.1.1/23	
cluster1-04	cluster1-04_clus1	e0a true	up/up	169.254.1.6/23	
cluster1-04	cluster1-04_clus2	e0b true	up/up	169.254.1.7/23	

5. Verify that the cluster displays information for both cluster switches.

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command:

```
system switch ethernet show -is-monitoring-enabled-operational true
```

```
cluster1::*> system switch ethernet show -is-monitoring-enabled  
-operational true
```

Switch	Type	Address	Model
cs1	cluster-network	10.228.143.200	BES-
53248			
	Serial Number:	QTWCU22510008	
	Is Monitored:	true	
	Reason:	None	
	Software Version:	3.10.0.3	
	Version Source:	CDP/ISDP	
cs2	cluster-network	10.228.143.202	BES-
53248			
	Serial Number:	QTWCU22510009	
	Is Monitored:	true	
	Reason:	None	
	Software Version:	3.10.0.3	
	Version Source:	CDP/ISDP	

```
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command:

```
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                    10.228.143.200                     BES-
53248
    Serial Number: QTWCU22510008
    Is Monitored: true
    Reason: None
    Software Version: 3.10.0.3
    Version Source: CDP/ISDP

cs2                                   cluster-network                    10.228.143.202                     BES-
53248
    Serial Number: QTWCU22510009
    Is Monitored: true
    Reason: None
    Software Version: 3.10.0.3
    Version Source: CDP/ISDP
cluster1::*>

```

6. Disable auto-revert on the cluster LIFs.

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

Step 2: Configure ports

1. On switch cs2, confirm the list of ports that are connected to the nodes in the cluster.

```
show isdp neighbor
```

2. On cluster switch cs2, shut down the ports connected to the cluster ports of the nodes. For example, if ports 0/1 to 0/16 are connected to ONTAP nodes:

```

(cs2)> enable
(cs2)# configure
(cs2) (Config)# interface 0/1-0/16
(cs2) (Interface 0/1-0/16)# shutdown
(cs2) (Interface 0/1-0/16)# exit
(cs2) (Config)#

```

- Verify that the cluster LIFs have migrated to the ports hosted on cluster switch cs1. This might take a few seconds.

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

Show example

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

Current Is	Logical	Status	Network	Current
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		

```
cluster1::*>
```

- Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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
```

5. If you have not already done so, save the current switch configuration by copying the output of the following command to a log file:

```
show running-config
```

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



When updating or applying a new RCF, you must erase the switch settings and perform basic configuration. You must be connected to the switch using the serial console to erase switch settings. This requirement is optional if you have used the Knowledge Base article [How to clear the configuration on a Broadcom interconnect switch while retaining remote connectivity](#) to clear the configuration, beforehand.



Clearing the configuration does not delete licenses.

- a. SSH into the switch.

Only proceed when all the cluster LIFs have been removed from the ports on the switch and the switch is prepared to have the configuration cleared.

- b. Enter privilege mode:

```
(cs2)> enable
(cs2)#
```

- c. Copy and paste the following commands to remove the previous RCF configuration (depending on the previous RCF version used, some commands might generate an error if a particular setting is not present):

```
clear config interface 0/1-0/56
y
clear config interface lag 1
y
configure
```

```
deleteport 1/1 all
no policy-map CLUSTER
no policy-map WRED_25G
no policy-map WRED_100G
no policy-map InShared
no policy-map InMetroCluster
no policy-map InCluster
no policy-map InClusterRdma
no class-map CLUSTER
no class-map HA
no class-map RDMA
no class-map c5
no class-map c4
no class-map CLUSTER
no class-map CLUSTER_RDMA
no class-map StorageSrc
no class-map StorageDst
no class-map RdmaSrc
no class-map RdmaDst
no classofservice dot1p-mapping
no random-detect queue-parms 0
no random-detect queue-parms 1
no random-detect queue-parms 2
no random-detect queue-parms 3
no random-detect queue-parms 4
no random-detect queue-parms 5
no random-detect queue-parms 6
no random-detect queue-parms 7
no cos-queue min-bandwidth
no cos-queue random-detect 0
no cos-queue random-detect 1
no cos-queue random-detect 2
no cos-queue random-detect 3
no cos-queue random-detect 4
no cos-queue random-detect 5
no cos-queue random-detect 6
no cos-queue random-detect 7
exit
vlan database
no vlan 17
no vlan 18
exit
```

d. Save the running configuration to the startup configuration:

```
(cs2)# write memory
```

```
This operation may take a few minutes.  
Management interfaces will not be available during this time.
```

```
Are you sure you want to save? (y/n) y
```

```
Config file 'startup-config' created successfully.
```

```
Configuration Saved!
```

e. Perform a reboot of the switch:

```
(cs2)# reload
```

```
Are you sure you would like to reset the system? (y/n) y
```

f. Log in to the switch again using SSH to complete the RCF installation.

7. Record any customizations that were made in the previous RCF and apply these to the new RCF. For example, setting port speeds or hard-coding FEC mode.
8. Copy the RCF to the bootflash of switch cs2 using one of the following transfer protocols: FTP, HTTP, TFTP, SFTP, or SCP.

This example shows HTTP being used to copy an RCF to the bootflash on switch cs2:

Show example

```
(cs2)# copy http://<ip-to-webserver>/path/to/BES-53248-RCF-v1.12-Cluster-HA.txt nvram:reference-config

Mode..... HTTP
Set Server IP..... 172.19.2.1
Path..... <ip-to-
webserver>/path/to/
Filename..... BES-53248-RCF-v1.12-
Cluster-HA.txt
Data Type..... Unknown

Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y
File transfer in progress.
Management access will be blocked for the duration of the transfer.
Please wait...
HTTP Unknown file type transfer starting...
Validating configuration script
.....
....
Configuration script validated.
File transfer operation completed successfully.
```

9. Verify that the script was downloaded and saved under the file name you gave it:

```
script list
```

```
(cs2)# script list

Configuration Script Name          Size(Bytes)  Date of
Modification
-----
Reference-config.scr              2680        2024 05 31
21:54:22
1 configuration script(s) found.
2045 Kbytes free.
```

10. Apply the script to the switch:

```
script apply
```


Show example

```
(cs2)# script apply reference-config.scr

Are you sure you want to apply the configuration script? (y/n) y

The system has unsaved changes.
Would you like to save them now? (y/n) y
Config file 'startup-config' created successfully.
Configuration Saved!
...
...
Configuration script 'reference-config.scr' applied.
```

11. Install the license file.

Show example

```
(cs2)# copy http://<ip-to-webserver>/path/to/BES-53248-LIC.dat
nvram:license-key 1
Mode..... HTTP
Set Server IP..... 172.19.2.1
Path..... <ip-to-
webserver>/path/to/
Filename..... BES-53248-LIC.dat
Data Type..... license

Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y

File transfer in progress. Management access will be blocked for the
duration of the transfer.

Please wait...

License Key transfer operation completed successfully.

System reboot is required.
(cs2)# write memory

This operation may take a few minutes.

Management interfaces will not be available during this time.
Are you sure you want to save? (y/n) y

Config file 'startup-config' created successfully.

Configuration Saved!

(cs2)# reload
Are you sure you would like to reset the system? (y/n) y
...
...
```

12. Examine the banner output from the `show clibanner` command. You must read and follow these instructions to verify the proper configuration and operation of the switch.

Show example

```
(cs2)# show clibanner
```

```
Banner Message configured :
```

```
=====
```

```
BES-53248 Reference Configuration File v1.12 for Cluster/HA/RDMA
```

```
Switch    : BES-53248
```

```
Filename  : BES-53248-RCF-v1.12-Cluster.txt
```

```
Date      : 11-04-2024
```

```
Version   : v1.12
```

```
Port Usage:
```

```
Ports 01 - 16: 10/25GbE Cluster Node Ports, base config
```

```
Ports 17 - 48: 10/25GbE Cluster Node Ports, with licenses
```

```
Ports 49 - 54: 40/100GbE Cluster Node Ports, with licenses, added  
right to left
```

```
Ports 55 - 56: 100GbE Cluster ISL Ports, base config
```

```
NOTE:
```

```
- The 48 SFP28/SFP+ ports are organized into 4-port groups in terms  
of port speed:
```

```
Ports 1-4, 5-8, 9-12, 13-16, 17-20, 21-24, 25-28, 29-32, 33-36,  
37-40, 41-44, 45-48
```

```
The port speed should be the same (10GbE or 25GbE) across all  
ports in a 4-port group
```

```
- If additional licenses are purchased, follow the 'Additional Node  
Ports
```

```
activated with Licenses' section for instructions
```

```
- If SSH is active, it will have to be re-enabled manually after  
'erase startup-config'
```

```
command has been executed and the switch rebooted"
```

13. On the switch, verify that the additional licensed ports appear after the RCF is applied:

```
show port all | exclude Detach
```

Show example

```
(cs2)# show port all | exclude Detach
```

LACP	Actor	Admin	Physical	Physical	Link	Link
Intf	Type	Mode	Mode	Status	Status	Trap
Mode	Timeout					
0/1	Enable long	Enable	Auto		Down	Enable
0/2	Enable long	Enable	Auto		Down	Enable
0/3	Enable long	Enable	Auto		Down	Enable
0/4	Enable long	Enable	Auto		Down	Enable
0/5	Enable long	Enable	Auto		Down	Enable
0/6	Enable long	Enable	Auto		Down	Enable
0/7	Enable long	Enable	Auto		Down	Enable
0/8	Enable long	Enable	Auto		Down	Enable
0/9	Enable long	Enable	Auto		Down	Enable
0/10	Enable long	Enable	Auto		Down	Enable
0/11	Enable long	Enable	Auto		Down	Enable
0/12	Enable long	Enable	Auto		Down	Enable
0/13	Enable long	Enable	Auto		Down	Enable
0/14	Enable long	Enable	Auto		Down	Enable
0/15	Enable long	Enable	Auto		Down	Enable
0/16	Enable long	Enable	Auto		Down	Enable
0/49	Enable long	Enable	40G Full		Down	Enable
0/50	Enable long	Enable	40G Full		Down	Enable

```
0/51          Enable    100G Full          Down    Enable
Enable long
0/52          Enable    100G Full          Down    Enable
Enable long
0/53          Enable    100G Full          Down    Enable
Enable long
0/54          Enable    100G Full          Down    Enable
Enable long
0/55          Enable    100G Full          Down    Enable
Enable long
0/56          Enable    100G Full          Down    Enable
Enable long
```

14. On the switch, verify that your changes have been made:

```
show running-config
```

```
(cs2)# show running-config
```

15. Save the running configuration so that it becomes the startup configuration when you reboot the switch:

```
write memory
```

```
(cs2)# write memory
This operation may take a few minutes.
Management interfaces will not be available during this time.

Are you sure you want to save? (y/n) y

Config file 'startup-config' created successfully.

Configuration Saved!
```

16. Reboot the switch and verify that the running configuration is correct:

```
reload
```

```
(cs2)# reload
```

```
Are you sure you would like to reset the system? (y/n) y
```

```
System will now restart!
```

17. On cluster switch cs2, bring up the ports connected to the cluster ports of the nodes. For example, if ports 0/1 to 0/16 are connected to ONTAP nodes:

```
(cs2)> enable
```

```
(cs2)# configure
```

```
(cs2) (Config)# interface 0/1-0/16
```

```
(cs2) (Interface 0/1-0/16)# no shutdown
```

```
(cs2) (Interface 0/1-0/16)# exit
```

```
(cs2) (Config)#
```

18. Verify the ports on switch cs2:

```
show interfaces status all | exclude Detach
```

Show example

```
(cs1)# show interfaces status all | exclude Detach
```

Media	Flow	Link	Physical	Physical	
Port	Name	State	Mode	Status	Type
Control	VLAN				
-----	-----	-----	-----	-----	
-----	-----	-----			
.					
.					
.					
0/16	10/25GbE Node Port	Down	Auto		
Inactive	Trunk				
0/17	10/25GbE Node Port	Down	Auto		
Inactive	Trunk				
0/18	10/25GbE Node Port	Up	25G Full	25G Full	
25GBase-SR	Inactive Trunk				
0/19	10/25GbE Node Port	Up	25G Full	25G Full	
25GBase-SR	Inactive Trunk				
.					
.					
.					
0/50	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/51	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/52	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/53	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/54	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/55	Cluster ISL Port	Up	Auto	100G Full	
Copper	Inactive Trunk				
0/56	Cluster ISL Port	Up	Auto	100G Full	
Copper	Inactive Trunk				

19. Verify the health of cluster ports on the cluster.

a. Verify that e0b ports are up and healthy across all nodes in the cluster:

```
network port show -role cluster
```

Show example

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

Node: cluster1-01

Ignore

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

Node: cluster1-02

Ignore

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

Node: cluster1-03

Ignore

Health Health
Port IPspace Broadcast Domain Link MTU Admin/Oper
Status Status
Speed (Mbps)
-----
e0a Cluster Cluster up 9000 auto/100000
healthy false
e0b Cluster Cluster up 9000 auto/100000
healthy false
```



```
Node: cluster1-04
```

```
Ignore
```

```
Health Health                               Speed (Mbps)
Port     IPspace      Broadcast Domain Link MTU  Admin/Oper
Status   Status
-----
e0a      Cluster     Cluster           up   9000  auto/100000
healthy  false
e0b      Cluster     Cluster           up   9000  auto/100000
healthy  false
```

b. Verify the switch health from the cluster:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
          e0a   cs1                        0/2
BES-53248
          e0b   cs2                        0/2
BES-53248
cluster01-2/cdp
          e0a   cs1                        0/1
BES-53248
          e0b   cs2                        0/1
BES-53248
cluster01-3/cdp
          e0a   cs1                        0/4
BES-53248
          e0b   cs2                        0/4
BES-53248
cluster1-04/cdp
          e0a   cs1                        0/3
BES-53248
          e0b   cs2                        0/2
BES-53248
```

20. Verify that the cluster displays information for both cluster switches.

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command:

```
system switch ethernet show -is-monitoring-enabled-operational true
```

```
cluster1::*> system switch ethernet show -is-monitoring-enabled
-operational true
```

Switch	Type	Address	Model
cs1 53248	cluster-network	10.228.143.200	BES-
Serial Number: QTWCU22510008			
Is Monitored: true			
Reason: None			
Software Version: 3.10.0.3			
Version Source: CDP/ISDP			
cs2 53248	cluster-network	10.228.143.202	BES-
Serial Number: QTWCU22510009			
Is Monitored: true			
Reason: None			
Software Version: 3.10.0.3			
Version Source: CDP/ISDP			

```
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command:

```
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                     10.228.143.200                       BES-
53248
      Serial Number: QTWCU22510008
      Is Monitored: true
      Reason: None
      Software Version: 3.10.0.3
      Version Source: CDP/ISDP

cs2                                    cluster-network                     10.228.143.202                       BES-
53248
      Serial Number: QTWCU22510009
      Is Monitored: true
      Reason: None
      Software Version: 3.10.0.3
      Version Source: CDP/ISDP
cluster1::*>

```

21. On cluster switch cs1, shut down the ports connected to the cluster ports of the nodes.

The following example uses the interface example output:

```

(cs1)> enable
(cs1)# configure
(cs1) (Config)# interface 0/1-0/16
(cs1) (Interface 0/1-0/16)# shutdown

```

22. 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

```

Show example

```
cluster1::*> network interface show -role cluster
          Logical          Status      Network          Current
Current  Is
Vserver  Interface          Admin/Oper  Address/Mask     Node
Port     Home
-----
Cluster
cluster1-01  cluster1-01_clus1  up/up      169.254.3.4/23
           e0a         false
cluster1-01  cluster1-01_clus2  up/up      169.254.3.5/23
           e0b         true
cluster1-02  cluster1-02_clus1  up/up      169.254.3.8/23
           e0a         false
cluster1-02  cluster1-02_clus2  up/up      169.254.3.9/23
           e0b         true
cluster1-03  cluster1-03_clus1  up/up      169.254.1.3/23
           e0a         false
cluster1-03  cluster1-03_clus2  up/up      169.254.1.1/23
           e0b         true
cluster1-04  cluster1-04_clus1  up/up      169.254.1.6/23
           e0a         false
cluster1-04  cluster1-04_clus2  up/up      169.254.1.7/23
           e0b         true
cluster1::*>
```

23. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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
```

24. Repeat steps 4 to 19 on switch cs1.

25. Enable auto-revert on the cluster LIFs:

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

26. Reboot switch cs1. This triggers 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  
The system has unsaved changes.  
Would you like to save them now? (y/n) y  
Config file 'startup-config' created successfully.  
Configuration Saved! System will now restart!
```

Step 3: Verify the configuration

1. On switch cs1, verify that the switch ports connected to the cluster ports are **up**:

```
show interfaces status all | exclude Detach
```

Show example

```
(cs1)# show interfaces status all | exclude Detach
```

Media	Flow	Link	Physical	Physical	
Port	Name	State	Mode	Status	Type
Control	VLAN				
-----	-----	-----	-----	-----	
-----	-----	-----			
.					
.					
.					
0/16	10/25GbE Node Port	Down	Auto		
Inactive	Trunk				
0/17	10/25GbE Node Port	Down	Auto		
Inactive	Trunk				
0/18	10/25GbE Node Port	Up	25G Full	25G Full	
25GBase-SR	Inactive Trunk				
0/19	10/25GbE Node Port	Up	25G Full	25G Full	
25GBase-SR	Inactive Trunk				
.					
.					
.					
0/50	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/51	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/52	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/53	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/54	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/55	Cluster ISL Port	Up	Auto	100G Full	
Copper	Inactive Trunk				
0/56	Cluster ISL Port	Up	Auto	100G Full	
Copper	Inactive Trunk				

2. Verify that the ISL between switches cs1 and cs2 is functional:

```
show port-channel 1/1
```

Show example

```
(cs1)# show port-channel 1/1
Local Interface..... 1/1
Channel Name..... Cluster-ISL
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Port-channel Min-links..... 1
Load Balance Option..... 7
(Enhanced hashing mode)
Mbr      Device/      Port      Port
Ports   Timeout      Speed     Active
-----  -
0/55    actor/long    Auto     True
        partner/long
0/56    actor/long    Auto     True
        partner/long
```

3. Verify that the cluster LIFs have reverted to their home port:

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


Show example

```
cluster1::*> network interface show -role cluster
          Logical          Status      Network          Current
Current Is
Vserver   Interface              Admin/Oper  Address/Mask     Node
Port      Home
-----
Cluster
cluster1-01 cluster1-01_clus1 up/up      169.254.3.4/23
          e0a          true
cluster1-01 cluster1-01_clus2 up/up      169.254.3.5/23
          e0b          true
cluster1-02 cluster1-02_clus1 up/up      169.254.3.8/23
          e0a          true
cluster1-02 cluster1-02_clus2 up/up      169.254.3.9/23
          e0b          true
cluster1-03 cluster1-03_clus1 up/up      169.254.1.3/23
          e0a          true
cluster1-03 cluster1-03_clus2 up/up      169.254.1.1/23
          e0b          true
cluster1-04 cluster1-04_clus1 up/up      169.254.1.6/23
          e0a          true
cluster1-04 cluster1-04_clus2 up/up      169.254.1.7/23
          e0b          true
```

4. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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
```

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
-----	-----	-----
-----	-----	-----
cluster1-01		
3/5/2022 19:21:18 -06:00	cluster1-01_clus2	cluster01-
02_clus1 none		
3/5/2022 19:21:20 -06:00	cluster1-01_clus2	cluster01-
02_clus2 none		
cluster1-02		
3/5/2022 19:21:18 -06:00	cluster1-02_clus2	cluster1-02_clus1
none		
3/5/2022 19:21:20 -06:00	cluster1-02_clus2	cluster1-02_clus2
none		

All ONTAP releases

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

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

```

cluster1::~*> cluster ping-cluster -node local
Host is 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 e0b
Cluster cluster1-02_clus1 169.254.3.8 cluster1-02 e0a
Cluster cluster1-02_clus2 169.254.3.9 cluster1-02 e0b
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)

```

6. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

What's next?

[Enable SSH.](#)

Install licenses for BES-53248 cluster switches

The BES-53248 cluster switch base model is licensed for 16 10GbE or 25GbE ports and two 100GbE ports. You can add new ports by purchasing more licenses.



For EFOS 3.12 and later, follow the installation steps in [Install the Reference Configuration File \(RCF\) and license file.](#)

Review available licenses

The following licenses are available for use on the BES-53248 cluster switch:

License type	License details	Supported firmware version
SW-BES-53248A2-8P-2P	Broadcom 8PT-10G25G + 2PT-40G100G License Key, X190005/R	EFOS 3.4.4.6 and later
SW-BES-53248A2-8P-1025G	Broadcom 8 Port 10G25G License Key, X190005/R	EFOS 3.4.4.6 and later
SW-BES53248A2-6P-40-100G	Broadcom 6 Port 40G100G License Key, X190005/R	EFOS 3.4.4.6 and later

Legacy licenses

The following table lists the legacy licenses that were available for use on the BES-53248 cluster switch:

License type	License details	Supported firmware version
SW-BES-53248A1-G1-8P-LIC	Broadcom 8P 10-25,2P40-100 License Key, X190005/R	EFOS 3.4.3.3 and later
SW-BES-53248A1-G1-16P-LIC	Broadcom 16P 10-25,4P40-100 License Key, X190005/R	EFOS 3.4.3.3 and later
SW-BES-53248A1-G1-24P-LIC	Broadcom 24P 10-25,6P40-100 License Key, X190005/R	EFOS 3.4.3.3 and later

License type	License details	Supported firmware version
SW-BES54248-40-100G-LIC	Broadcom 6Port 40G100G License Key, X190005/R	EFOS 3.4.4.6 and later
SW-BES53248-8P-10G25G-LIC	Broadcom 8Port 10G25G License Key, X190005/R	EFOS 3.4.4.6 and later
SW-BES53248-16P-1025G-LIC	Broadcom 16Port 10G25G License Key, X190005/R	EFOS 3.4.4.6 and later
SW-BES53248-24P-1025G-LIC	Broadcom 24Port 10G25G License Key, X190005/R	EFOS 3.4.4.6 and later



A license is not required for the base configuration.

Install license files

Follow these steps to install licenses for BES-53248 cluster switches.

Steps

1. Connect the cluster switch to the management network.
2. Use the `ping` command to verify connectivity to the server hosting EFOS, licenses, and the RCF file.

Show example

This example verifies that the switch is connected to 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.
```

3. Check the current license usage on switch cs2:

```
show license
```

Show example

```
(cs2)# show license
Reboot needed..... No
Number of active licenses..... 0

License Index  License Type      Status
-----
-----

No license file found.
```

4. Install the license file.

Repeat this step to load more licenses and to use different key index numbers.

Show example

The following example uses SFTP to copy a license file to a key index 1.

```
(cs2)# copy sftp://root@172.19.2.1/var/lib/tftpboot/license.dat
nvram:license-key 1
Remote Password:**

Mode..... SFTP
Set Server IP..... 172.19.2.1
Path..... /var/lib/tftpboot/
Filename..... license.dat
Data Type..... license

Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y

File transfer in progress. Management access will be blocked for the
duration of the transfer. Please wait...

License Key transfer operation completed successfully. System reboot
is required.
```

5. Display all current license information and note the license status before switch cs2 is rebooted:

```
show license
```

Show example

```
(cs2)# show license
```

```
Reboot needed..... Yes  
Number of active licenses..... 0
```

License Index	License Type	Status
1	Port	License valid but not applied

6. Display all licensed ports:

```
show port all | exclude Detach
```

The ports from the additional license files are not displayed until after the switch is rebooted.

Show example




```
(cs2)# show port all | exclude Detach
```

Actor	Admin	Physical	Physical	Link	Link	LACP	
Intf	Type	Mode	Mode	Status	Status	Trap	Mode
Timeout							
0/1		Disable	Auto		Down	Enable	
Enable long							
0/2		Disable	Auto		Down	Enable	
Enable long							
0/3		Disable	Auto		Down	Enable	
Enable long							
0/4		Disable	Auto		Down	Enable	
Enable long							
0/5		Disable	Auto		Down	Enable	
Enable long							
0/6		Disable	Auto		Down	Enable	
Enable long							
0/7		Disable	Auto		Down	Enable	
Enable long							
0/8		Disable	Auto		Down	Enable	
Enable long							
0/9		Disable	Auto		Down	Enable	
Enable long							
0/10		Disable	Auto		Down	Enable	
Enable long							
0/11		Disable	Auto		Down	Enable	
Enable long							
0/12		Disable	Auto		Down	Enable	
Enable long							
0/13		Disable	Auto		Down	Enable	
Enable long							
0/14		Disable	Auto		Down	Enable	
Enable long							
0/15		Disable	Auto		Down	Enable	
Enable long							
0/16		Disable	Auto		Down	Enable	
Enable long							
0/55		Disable	Auto		Down	Enable	
Enable long							
0/56		Disable	Auto		Down	Enable	
Enable long							

7. Reboot the switch:

```
reload
```

Show example

```
(cs2)# reload

The system has unsaved changes.
Would you like to save them now? (y/n) y

Config file 'startup-config' created successfully .

Configuration Saved!
Are you sure you would like to reset the system? (y/n) y
```

8. Check that the new license is active and note that the license has been applied:

```
show license
```

Show example

```
(cs2)# show license

Reboot needed..... No
Number of installed licenses..... 1
Total Downlink Ports enabled..... 16
Total Uplink Ports enabled..... 8

License Index  License Type          Status
-----
-----
1              Port                License applied
```

9. Check that all new ports are available:

```
show port all | exclude Detach
```

Show example

```
(cs2)# show port all | exclude Detach
```

Actor	Admin	Physical	Physical	Link	Link	LACP
Intf	Type	Mode	Mode	Status	Status	Trap
Timeout						Mode
0/1		Disable	Auto		Down	Enable
Enable long						
0/2		Disable	Auto		Down	Enable
Enable long						
0/3		Disable	Auto		Down	Enable
Enable long						
0/4		Disable	Auto		Down	Enable
Enable long						
0/5		Disable	Auto		Down	Enable
Enable long						
0/6		Disable	Auto		Down	Enable
Enable long						
0/7		Disable	Auto		Down	Enable
Enable long						
0/8		Disable	Auto		Down	Enable
Enable long						
0/9		Disable	Auto		Down	Enable
Enable long						
0/10		Disable	Auto		Down	Enable
Enable long						
0/11		Disable	Auto		Down	Enable
Enable long						
0/12		Disable	Auto		Down	Enable
Enable long						
0/13		Disable	Auto		Down	Enable
Enable long						
0/14		Disable	Auto		Down	Enable
Enable long						
0/15		Disable	Auto		Down	Enable
Enable long						
0/16		Disable	Auto		Down	Enable
Enable long						
0/49		Disable	100G Full		Down	Enable
Enable long						
0/50		Disable	100G Full		Down	Enable
Enable long						

0/51	Disable	100G	Full	Down	Enable
Enable long					
0/52	Disable	100G	Full	Down	Enable
Enable long					
0/53	Disable	100G	Full	Down	Enable
Enable long					
0/54	Disable	100G	Full	Down	Enable
Enable long					
0/55	Disable	100G	Full	Down	Enable
Enable long					
0/56	Disable	100G	Full	Down	Enable
Enable long					



When installing additional licenses, you must configure the new interfaces manually. Do not re-apply an RCF to an existing working production switch.

Troubleshoot install issues

Where problems arise when installing a license, run the following debug commands before running the `copy` command again.

Debug commands to use: `debug transfer` and `debug license`

Show example

```
(cs2)# debug transfer
Debug transfer output is enabled.
(cs2)# debug license
Enabled capability licensing debugging.
```

When you run the `copy` command with the `debug transfer` and `debug license` options enabled, the log output is returned.

Show example

```
transfer.c(3083):Transfer process key or certificate file type = 43
transfer.c(3229):Transfer process key/certificate cmd = cp
/mnt/download//license.dat.1 /mnt/fastpath/ >/dev/null 2>&1CAPABILITY
LICENSING :
Fri Sep 11 13:41:32 2020: License file with index 1 added.
CAPABILITY LICENSING : Fri Sep 11 13:41:32 2020: Validating hash value
29de5e9a8af3e510f1f16764a13e8273922d3537d3f13c9c3d445c72a180a2e6.
CAPABILITY LICENSING : Fri Sep 11 13:41:32 2020: Parsing JSON buffer {
  "license": {
    "header": {
      "version": "1.0",
      "license-key": "964B-2D37-4E52-BA14",
      "serial-number": "QTFCU38290012",
      "model": "BES-53248"
    },
    "description": "",
    "ports": "0+6"
  }
}.
CAPABILITY LICENSING : Fri Sep 11 13:41:32 2020: License data does not
contain 'features' field.
CAPABILITY LICENSING : Fri Sep 11 13:41:32 2020: Serial number
QTFCU38290012 matched.
CAPABILITY LICENSING : Fri Sep 11 13:41:32 2020: Model BES-53248
matched.
CAPABILITY LICENSING : Fri Sep 11 13:41:32 2020: Feature not found in
license file with index = 1.
CAPABILITY LICENSING : Fri Sep 11 13:41:32 2020: Applying license file
1.
```

Check for the following in the debug output:

- Check that the Serial number matches: Serial number QTFCU38290012 matched.
- Check that the switch Model matches: Model BES-53248 matched.
- Check that the specified license index was not used previously. Where a license index is already used, the following error is returned: License file /mnt/download//license.dat.1 already exists.
- A port license is not a feature license. Therefore, the following statement is expected: Feature not found in license file with index = 1.

Use the `copy` command to back up port licenses to the server:

```
(cs2) # copy nvram:license-key 1
scp://<UserName>@<IP_address>/saved_license_1.dat
```



If you need to downgrade the switch software from version 3.4.4.6, the licenses are removed. This is expected behavior.

You must install an appropriate older license before reverting to an older version of the software.

Activate newly licensed ports

To activate newly licensed ports, you need to edit the latest version of the RCF and uncomment the applicable port details.

The default license activates ports 0/1 to 0/16 and 0/55 to 0/56 while the newly licensed ports will be between ports 0/17 to 0/54 depending on the type and number of licenses available. For example, to activate the SW-BES54248-40-100G-LIC license, you must uncomment the following section in the RCF:

Show example

```
.
.
!
! 2-port or 6-port 40/100GbE node port license block
!
interface 0/49
no shutdown
description "40/100GbE Node Port"
!speed 100G full-duplex
speed 40G full-duplex
service-policy in WRED_100G
spanning-tree edgeport
mtu 9216
switchport mode trunk
datacenter-bridging
priority-flow-control mode on
priority-flow-control priority 5 no-drop
exit
exit
!
interface 0/50
no shutdown
description "40/100GbE Node Port"
!speed 100G full-duplex
speed 40G full-duplex
service-policy in WRED_100G
spanning-tree edgeport
mtu 9216
switchport mode trunk
datacenter-bridging
priority-flow-control mode on
priority-flow-control priority 5 no-drop
exit
exit
!
interface 0/51
no shutdown
description "40/100GbE Node Port"
speed 100G full-duplex
!speed 40G full-duplex
service-policy in WRED_100G
spanning-tree edgeport
mtu 9216
switchport mode trunk
```

```
datacenter-bridging
priority-flow-control mode on
priority-flow-control priority 5 no-drop
exit
exit
!
interface 0/52
no shutdown
description "40/100GbE Node Port"
speed 100G full-duplex
!speed 40G full-duplex
service-policy in WRED_100G
spanning-tree edgeport
mtu 9216
switchport mode trunk
datacenter-bridging
priority-flow-control mode on
priority-flow-control priority 5 no-drop
exit
exit
!
interface 0/53
no shutdown
description "40/100GbE Node Port"
speed 100G full-duplex
!speed 40G full-duplex
service-policy in WRED_100G
spanning-tree edgeport
mtu 9216
switchport mode trunk
datacenter-bridging
priority-flow-control mode on
priority-flow-control priority 5 no-drop
exit
exit
!
interface 0/54
no shutdown
description "40/100GbE Node Port"
speed 100G full-duplex
!speed 40G full-duplex
service-policy in WRED_100G
spanning-tree edgeport
mtu 9216
switchport mode trunk
datacenter-bridging
```



```
priority-flow-control mode on
priority-flow-control priority 5 no-drop
exit
exit
!
.
.
```



For high-speed ports between 0/49 to 0/54 inclusive, uncomment each port but only uncomment one **speed** line in the RCF for each of these ports, either: **speed 100G full-duplex** or **speed 40G full-duplex** as shown in the example.

For low-speed ports between 0/17 to 0/48 inclusive, uncomment the entire 8-port section when an appropriate license has been activated.

What's next?

[Install the Reference Configuration File \(RCF\)](#) or [Upgrade the RCF](#).

Install the Reference Configuration File (RCF)

You can install the Reference Configuration File (RCF) after configuring the BES-53248 cluster switch and after applying the new licenses.



For EFOS 3.12 and later, follow the installation steps in [Install the Reference Configuration File \(RCF\) and license file](#).

Review requirements

Before you begin

Verify that the following are in place:

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).
- The current RCF file, available from the [Broadcom Cluster Switches](#) page.
- A boot configuration in the RCF that reflects the desired boot images, required if you are installing only EFOS 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.
- A console connection to the switch, required when installing the RCF from a factory-default state. This requirement is optional if you have used the Knowledge Base article [How to clear configuration on a Broadcom interconnect switch while retaining remote connectivity](#) to clear the configuration, beforehand.

Suggested documentation

Consult the switch compatibility table for the supported ONTAP and RCF versions. See the [EFOS Software download](#) page. Note that there can be command dependencies between the command syntax in the RCF and that found in versions of EFOS.

Install the configuration file

About the examples

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

- The names of the two BES-53248 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.
- The examples in this procedure use four nodes. These nodes use two 10GbE cluster interconnect ports e0a and e0b. See the [Hardware Universe](#) to verify the correct cluster ports on your platforms.



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

About this task

The procedure requires the use of both ONTAP commands and Broadcom switch commands; ONTAP commands are used unless otherwise indicated.

No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure migrates all the cluster LIFs to the operational partner switch while performing the steps on the target switch.



Before installing a new switch software version and RCFs, use the Knowledge Base article [How to clear configuration on a Broadcom interconnect switch while retaining remote connectivity](#). If you must erase the switch settings completely, then you need to perform the basic configuration again. You must be connected to the switch using the serial console because a complete configuration erasure resets the configuration of the management network.

Step 1: Prepare for the installation

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
```

where *x* is the duration of the maintenance window in hours.



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

The following command suppresses automatic case creation for two hours:

```
cluster1::*> system node autosupport invoke -node \* -type all -message MAINT=2h
```

2. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (*>) appears.

3. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
              e0a    cs1                       0/2          BES-
53248
              e0b    cs2                       0/2          BES-
53248
cluster1-02/cdp
              e0a    cs1                       0/1          BES-
53248
              e0b    cs2                       0/1          BES-
53248
cluster1-03/cdp
              e0a    cs1                       0/4          BES-
53248
              e0b    cs2                       0/4          BES-
53248
cluster1-04/cdp
              e0a    cs1                       0/3          BES-
53248
              e0b    cs2                       0/3          BES-
53248
cluster1::*>
```

4. 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
```

Show example

```
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					
e0b	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					
e0b	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
```

```
Health Health Speed (Mbps)
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
```

Show example

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

Current	Logical	Status	Network	
Vserver	Current Is			
Port	Interface	Admin/Oper	Address/Mask	Node
Home				
Cluster				
cluster1-01	cluster1-01_clus1	up/up	169.254.3.4/23	
	e0a true			
cluster1-01	cluster1-01_clus2	up/up	169.254.3.5/23	
	e0b true			
cluster1-02	cluster1-02_clus1	up/up	169.254.3.8/23	
	e0a true			
cluster1-02	cluster1-02_clus2	up/up	169.254.3.9/23	
	e0b true			
cluster1-03	cluster1-03_clus1	up/up	169.254.1.3/23	
	e0a true			
cluster1-03	cluster1-03_clus2	up/up	169.254.1.1/23	
	e0b true			
cluster1-04	cluster1-04_clus1	up/up	169.254.1.6/23	
	e0a true			
cluster1-04	cluster1-04_clus2	up/up	169.254.1.7/23	
	e0b true			

5. Verify that the cluster displays information for both cluster switches.

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command:

```
system switch ethernet show -is-monitoring-enabled-operational true
```

```
cluster1::*> system switch ethernet show -is-monitoring-enabled  
-operational true
```

Switch	Type	Address	Model
cs1 53248	cluster-network	10.228.143.200	BES-
Serial Number: QTWCU22510008			
Is Monitored: true			
Reason: None			
Software Version: 3.10.0.3			
Version Source: CDP/ISDP			
cs2 53248	cluster-network	10.228.143.202	BES-
Serial Number: QTWCU22510009			
Is Monitored: true			
Reason: None			
Software Version: 3.10.0.3			
Version Source: CDP/ISDP			

```
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command:

```
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                    10.228.143.200                    BES-
53248
    Serial Number: QTWCU22510008
    Is Monitored: true
    Reason: None
    Software Version: 3.10.0.3
    Version Source: CDP/ISDP

cs2                                    cluster-network                    10.228.143.202                    BES-
53248
    Serial Number: QTWCU22510009
    Is Monitored: true
    Reason: None
    Software Version: 3.10.0.3
    Version Source: CDP/ISDP
cluster1::*>

```

6. Disable auto-revert on the cluster LIFs.

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

Step 2: Configure ports

1. On switch cs2, confirm the list of ports that are connected to the nodes in the cluster.

```
show isdp neighbor
```

2. On cluster switch cs2, shut down the ports connected to the cluster ports of the nodes. For example, if ports 0/1 to 0/16 are connected to ONTAP nodes:

```

(cs2)> enable
(cs2)# configure
(cs2) (Config)# interface 0/1-0/16
(cs2) (Interface 0/1-0/16)# shutdown
(cs2) (Interface 0/1-0/16)# exit
(cs2) (Config)#

```

3. Verify that the cluster LIFs have migrated to the ports hosted on cluster switch cs1. This might take a few seconds.

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

Show example

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

Current Is	Logical	Status	Network	Current
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			

```
cluster1::*>
```

4. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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
```

5. If you have not already done so, save the current switch configuration by copying the output of the following command to a log file:

```
show running-config
```

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



When updating or applying a new RCF, you must erase the switch settings and perform basic configuration. You must be connected to the switch using the serial console to erase switch settings. This requirement is optional if you have used the Knowledge Base article [How to clear the configuration on a Broadcom interconnect switch while retaining remote connectivity](#) to clear the configuration, beforehand.



Clearing the configuration does not delete licenses.

- a. SSH into the switch.

Only proceed when all the cluster LIFs have been removed from the ports on the switch and the switch is prepared to have the configuration cleared.

- b. Enter privilege mode:

```
(cs2)> enable
(cs2)#
```

- c. Copy and paste the following commands to remove the previous RCF configuration (depending on the previous RCF version used, some commands might generate an error if a particular setting is not present):

```
clear config interface 0/1-0/56
y
clear config interface lag 1
y
configure
deleport 1/1 all
no policy-map CLUSTER
no policy-map WRED_25G
no policy-map WRED_100G
no class-map CLUSTER
no class-map HA
no class-map RDMA
no classofservice dot1p-mapping
no random-detect queue-parms 0
no random-detect queue-parms 1
no random-detect queue-parms 2
no random-detect queue-parms 3
no random-detect queue-parms 4
no random-detect queue-parms 5
no random-detect queue-parms 6
no random-detect queue-parms 7
no cos-queue min-bandwidth
no cos-queue random-detect 0
no cos-queue random-detect 1
no cos-queue random-detect 2
no cos-queue random-detect 3
no cos-queue random-detect 4
no cos-queue random-detect 5
no cos-queue random-detect 6
no cos-queue random-detect 7
exit
vlan database
no vlan 17
no vlan 18
exit
```

d. Save the running configuration to the startup configuration:

```
(cs2)# write memory
```

```
This operation may take a few minutes.  
Management interfaces will not be available during this time.
```

```
Are you sure you want to save? (y/n) y
```

```
Config file 'startup-config' created successfully.
```

```
Configuration Saved!
```

e. Perform a reboot of the switch:

```
(cs2)# reload
```

```
Are you sure you would like to reset the system? (y/n) y
```

f. Log in to the switch again using SSH to complete the RCF installation.

7. Note the following:

- a. If additional port licenses have been installed on the switch, you must modify the RCF to configure the additional licensed ports. See [Activate newly licensed ports](#) for details.
- b. Record any customizations that were made in the previous RCF and apply these to the new RCF. For example, setting port speeds or hard-coding FEC mode.

EFOS version 3.12.x and later

- Copy the RCF to the bootflash of switch cs2 using one of the following transfer protocols: FTP, TFTP, SFTP, or SCP.

This example shows SFTP being used to copy an RCF to the bootflash on switch cs2:

```
(cs2)# copy sftp://172.19.2.1/BES-53248-RCF-v1.9-Cluster-HA.txt
nvram:reference-config.scr
Remote Password:**
Mode..... TFTP
Set Server IP..... 172.19.2.1
Path..... /
Filename..... BES-53248_RCF_v1.9-
Cluster-HA.txt
Data Type..... Config Script
Destination Filename..... reference-config.scr
Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y
TFTP Code transfer starting...
File transfer operation completed successfully.
```

- Verify that the script was downloaded and saved under the file name you gave it:

```
script list
```

```
(cs2)# script list

Configuration Script Name                Size(Bytes)  Date of
Modification
-----
reference-config.scr                    2680         2024 05 31
21:54:22
2 configuration script(s) found.
2042 Kbytes free.
```

- Apply the script to the switch:

```
script apply
```

```
(cs2)# script apply reference-config.scr
```

```
Are you sure you want to apply the configuration script? (y/n) y
```

```
The system has unsaved changes.
```

```
Would you like to save them now? (y/n) y
```

```
Config file 'startup-config' created successfully.
```

```
Configuration Saved!
```

```
Configuration script 'reference-config.scr' applied.
```

All other EFOS versions

8. Copy the RCF to the bootflash of switch cs2 using one of the following transfer protocols: FTP, TFTP, SFTP, or SCP.

This example shows SFTP being used to copy an RCF to the bootflash on switch cs2:

```
(cs2)# copy sftp://172.19.2.1/tmp/BES-53248_RCF_v1.9-Cluster-HA.txt  
nvram:script BES-53248_RCF_v1.9-Cluster-HA.scr
```

```
Remote Password:**
```

```
Mode..... SFTP
```

```
Set Server IP..... 172.19.2.1
```

```
Path..... //tmp/
```

```
Filename..... BES-53248_RCF_v1.9-  
Cluster-HA.txt
```

```
Data Type..... Config Script
```

```
Destination Filename..... BES-53248_RCF_v1.9-  
Cluster-HA.scr
```

```
Management access will be blocked for the duration of the transfer
```

```
Are you sure you want to start? (y/n) y
```

```
SFTP Code transfer starting...
```

```
File transfer operation completed successfully.
```

9. Verify that the script was downloaded and saved to the file name you gave it:

```
script list
```

```
(cs2)# script list
```

```
Configuration Script Name          Size(Bytes)  Date of
Modification
-----
-----
BES-53248_RCF_v1.9-Cluster-HA.scr  2241        2020 09 30
05:41:00

1 configuration script(s) found.
```

10. Apply the script to the switch:

```
script apply
```

```
(cs2)# script apply BES-53248_RCF_v1.9-Cluster-HA.scr
```

```
Are you sure you want to apply the configuration script? (y/n) y
```

```
The system has unsaved changes.
```

```
Would you like to save them now? (y/n) y
```

```
Config file 'startup-config' created successfully.
```

```
Configuration Saved!
```

```
Configuration script 'BES-53248_RCF_v1.9-Cluster-HA.scr' applied.
```

11. Examine the banner output from the `show clibanner` command. You must read and follow these instructions to verify the proper configuration and operation of the switch.

Show example

```
(cs2)# show clibanner
```

```
Banner Message configured :
```

```
=====
```

```
BES-53248 Reference Configuration File v1.9 for Cluster/HA/RDMA
```

```
Switch    : BES-53248
```

```
Filename  : BES-53248-RCF-v1.9-Cluster.txt
```

```
Date      : 10-26-2022
```

```
Version   : v1.9
```

```
Port Usage:
```

```
Ports 01 - 16: 10/25GbE Cluster Node Ports, base config
```

```
Ports 17 - 48: 10/25GbE Cluster Node Ports, with licenses
```

```
Ports 49 - 54: 40/100GbE Cluster Node Ports, with licenses, added  
right to left
```

```
Ports 55 - 56: 100GbE Cluster ISL Ports, base config
```

```
NOTE:
```

```
- The 48 SFP28/SFP+ ports are organized into 4-port groups in terms  
of port
```

```
speed:
```

```
Ports 1-4, 5-8, 9-12, 13-16, 17-20, 21-24, 25-28, 29-32, 33-36, 37-  
40, 41-44,  
45-48
```

```
The port speed should be the same (10GbE or 25GbE) across all ports  
in a 4-port
```

```
group
```

```
- If additional licenses are purchased, follow the 'Additional Node  
Ports
```

```
activated with Licenses' section for instructions
```

```
- If SSH is active, it will have to be re-enabled manually after  
'erase
```

```
startup-config'
```

```
command has been executed and the switch rebooted
```

12. On the switch, verify that the additional licensed ports appear after the RCF is applied:

```
show port all | exclude Detach
```

Show example

```
(cs2)# show port all | exclude Detach
```

Intf	Admin	Physical	Physical	Link	Link
Mode	Mode	Mode	Status	Status	Trap
Actor	Mode	Mode	Status	Status	Trap
Type	Mode	Mode	Status	Status	Trap
Timeout	Mode	Mode	Status	Status	Trap
0/1	Enable	Auto		Down	Enable
Enable long					
0/2	Enable	Auto		Down	Enable
Enable long					
0/3	Enable	Auto		Down	Enable
Enable long					
0/4	Enable	Auto		Down	Enable
Enable long					
0/5	Enable	Auto		Down	Enable
Enable long					
0/6	Enable	Auto		Down	Enable
Enable long					
0/7	Enable	Auto		Down	Enable
Enable long					
0/8	Enable	Auto		Down	Enable
Enable long					
0/9	Enable	Auto		Down	Enable
Enable long					
0/10	Enable	Auto		Down	Enable
Enable long					
0/11	Enable	Auto		Down	Enable
Enable long					
0/12	Enable	Auto		Down	Enable
Enable long					
0/13	Enable	Auto		Down	Enable
Enable long					
0/14	Enable	Auto		Down	Enable
Enable long					
0/15	Enable	Auto		Down	Enable
Enable long					
0/16	Enable	Auto		Down	Enable
Enable long					
0/49	Enable	40G Full		Down	Enable
Enable long					
0/50	Enable	40G Full		Down	Enable
Enable long					

```
0/51          Enable    100G Full          Down    Enable
Enable long
0/52          Enable    100G Full          Down    Enable
Enable long
0/53          Enable    100G Full          Down    Enable
Enable long
0/54          Enable    100G Full          Down    Enable
Enable long
0/55          Enable    100G Full          Down    Enable
Enable long
0/56          Enable    100G Full          Down    Enable
Enable long
```

13. Verify on the switch that your changes have been made:

```
show running-config
```

```
(cs2)# show running-config
```

14. Save the running configuration so that it becomes the startup configuration when you reboot the switch:

```
write memory
```

```
(cs2)# write memory
This operation may take a few minutes.
Management interfaces will not be available during this time.

Are you sure you want to save? (y/n) y

Config file 'startup-config' created successfully.

Configuration Saved!
```

15. Reboot the switch and verify that the running configuration is correct:

```
reload
```

```
(cs2)# reload
```

```
Are you sure you would like to reset the system? (y/n) y
```

```
System will now restart!
```

16. On cluster switch cs2, bring up the ports connected to the cluster ports of the nodes. For example, if ports 0/1 to 0/16 are connected to ONTAP nodes:

```
(cs2)> enable
```

```
(cs2)# configure
```

```
(cs2) (Config)# interface 0/1-0/16
```

```
(cs2) (Interface 0/1-0/16)# no shutdown
```

```
(cs2) (Interface 0/1-0/16)# exit
```

```
(cs2) (Config)#
```

17. Verify the ports on switch cs2:

```
show interfaces status all | exclude Detach
```

Show example

```
(cs1)# show interfaces status all | exclude Detach
```

Media	Flow	Link	Physical	Physical	
Port	Name	State	Mode	Status	Type
Control	VLAN				
-----	-----	-----	-----	-----	
-----	-----	-----			
.					
.					
.					
0/16	10/25GbE Node Port	Down	Auto		
Inactive	Trunk				
0/17	10/25GbE Node Port	Down	Auto		
Inactive	Trunk				
0/18	10/25GbE Node Port	Up	25G Full	25G Full	
25GBase-SR	Inactive Trunk				
0/19	10/25GbE Node Port	Up	25G Full	25G Full	
25GBase-SR	Inactive Trunk				
.					
.					
.					
0/50	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/51	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/52	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/53	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/54	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/55	Cluster ISL Port	Up	Auto	100G Full	
Copper	Inactive Trunk				
0/56	Cluster ISL Port	Up	Auto	100G Full	
Copper	Inactive Trunk				

18. Verify the health of cluster ports on the cluster.

a. Verify that e0b ports are up and healthy across all nodes in the cluster:

```
network port show -role cluster
```

Show example

```
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					
e0b	Cluster	Cluster		up	9000	auto/100000
healthy	false					

Node: cluster1-04

Ignore

Health	Health					Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
e0a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e0b	Cluster	Cluster		up	9000	auto/100000
healthy	false					

b. Verify the switch health from the cluster:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
           e0a   cs1                       0/2
BES-53248
           e0b   cs2                       0/2
BES-53248
cluster01-2/cdp
           e0a   cs1                       0/1
BES-53248
           e0b   cs2                       0/1
BES-53248
cluster01-3/cdp
           e0a   cs1                       0/4
BES-53248
           e0b   cs2                       0/4
BES-53248
cluster1-04/cdp
           e0a   cs1                       0/3
BES-53248
           e0b   cs2                       0/2
BES-53248
```

19. Verify that the cluster displays information for both cluster switches.

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command:

```
system switch ethernet show -is-monitoring-enabled-operational true
```

```
cluster1::*> system switch ethernet show -is-monitoring-enabled  
-operational true
```

Switch	Type	Address	Model
cs1 53248	cluster-network	10.228.143.200	BES-
Serial Number: QTWCU22510008			
Is Monitored: true			
Reason: None			
Software Version: 3.10.0.3			
Version Source: CDP/ISDP			
cs2 53248	cluster-network	10.228.143.202	BES-
Serial Number: QTWCU22510009			
Is Monitored: true			
Reason: None			
Software Version: 3.10.0.3			
Version Source: CDP/ISDP			

```
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command:

```
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                    10.228.143.200                       BES-
53248
    Serial Number: QTWCU22510008
    Is Monitored: true
    Reason: None
    Software Version: 3.10.0.3
    Version Source: CDP/ISDP

cs2                                   cluster-network                    10.228.143.202                       BES-
53248
    Serial Number: QTWCU22510009
    Is Monitored: true
    Reason: None
    Software Version: 3.10.0.3
    Version Source: CDP/ISDP
cluster1::*>

```

20. On cluster switch cs1, shut down the ports connected to the cluster ports of the nodes.

The following example uses the interface example output:

```

(cs1)> enable
(cs1)# configure
(cs1) (Config)# interface 0/1-0/16
(cs1) (Interface 0/1-0/16)# shutdown

```

21. 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

```

Show example

```
cluster1::*> network interface show -role cluster
          Logical          Status      Network          Current
Current  Is
Vserver  Interface            Admin/Oper  Address/Mask     Node
Port     Home
-----
Cluster
cluster1-01  cluster1-01_clus1  up/up      169.254.3.4/23
            e0a          false
cluster1-01  cluster1-01_clus2  up/up      169.254.3.5/23
            e0b          true
cluster1-02  cluster1-02_clus1  up/up      169.254.3.8/23
            e0a          false
cluster1-02  cluster1-02_clus2  up/up      169.254.3.9/23
            e0b          true
cluster1-03  cluster1-03_clus1  up/up      169.254.1.3/23
            e0a          false
cluster1-03  cluster1-03_clus2  up/up      169.254.1.1/23
            e0b          true
cluster1-04  cluster1-04_clus1  up/up      169.254.1.6/23
            e0a          false
cluster1-04  cluster1-04_clus2  up/up      169.254.1.7/23
            e0b          true
cluster1::*>
```

22. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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
```

23. Repeat steps 4 to 19 on switch cs1.

24. Enable auto-revert on the cluster LIFs:

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

25. Reboot switch cs1. This triggers 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  
The system has unsaved changes.  
Would you like to save them now? (y/n) y  
Config file 'startup-config' created successfully.  
Configuration Saved! System will now restart!
```

Step 3: Verify the configuration

1. On switch cs1, verify that the switch ports connected to the cluster ports are **up**:

```
show interfaces status all | exclude Detach
```

Show example

```
(cs1)# show interfaces status all | exclude Detach
```

Media	Flow	Link	Physical	Physical	
Port	Name	State	Mode	Status	Type
Control	VLAN				
-----	-----	-----	-----	-----	
-----	-----	-----			
.					
.					
.					
0/16	10/25GbE Node Port	Down	Auto		
Inactive	Trunk				
0/17	10/25GbE Node Port	Down	Auto		
Inactive	Trunk				
0/18	10/25GbE Node Port	Up	25G Full	25G Full	
25GBase-SR	Inactive Trunk				
0/19	10/25GbE Node Port	Up	25G Full	25G Full	
25GBase-SR	Inactive Trunk				
.					
.					
.					
0/50	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/51	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/52	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/53	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/54	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/55	Cluster ISL Port	Up	Auto	100G Full	
Copper	Inactive Trunk				
0/56	Cluster ISL Port	Up	Auto	100G Full	
Copper	Inactive Trunk				

2. Verify that the ISL between switches cs1 and cs2 is functional:

```
show port-channel 1/1
```

Show example

```
(cs1)# show port-channel 1/1
Local Interface..... 1/1
Channel Name..... Cluster-ISL
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Port-channel Min-links..... 1
Load Balance Option..... 7
(Enhanced hashing mode)
Mbr      Device/      Port      Port
Ports   Timeout      Speed     Active
-----  -
0/55    actor/long    Auto     True
        partner/long
0/56    actor/long    Auto     True
        partner/long
```

3. Verify that the cluster LIFs have reverted to their home port:

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

Show example

```
cluster1::*> network interface show -role cluster
          Logical          Status      Network          Current
Current Is
Vserver   Interface          Admin/Oper Address/Mask      Node
Port      Home
-----
Cluster
cluster1-01 cluster1-01_clus1 up/up      169.254.3.4/23
           e0a      true
cluster1-01 cluster1-01_clus2 up/up      169.254.3.5/23
           e0b      true
cluster1-02 cluster1-02_clus1 up/up      169.254.3.8/23
           e0a      true
cluster1-02 cluster1-02_clus2 up/up      169.254.3.9/23
           e0b      true
cluster1-03 cluster1-03_clus1 up/up      169.254.1.3/23
           e0a      true
cluster1-03 cluster1-03_clus2 up/up      169.254.1.1/23
           e0b      true
cluster1-04 cluster1-04_clus1 up/up      169.254.1.6/23
           e0a      true
cluster1-04 cluster1-04_clus2 up/up      169.254.1.7/23
           e0b      true
```

4. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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
```

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
-----	-----	-----
-----	-----	-----
cluster1-01		
3/5/2022 19:21:18 -06:00	cluster1-01_clus2	cluster01-
02_clus1 none		
3/5/2022 19:21:20 -06:00	cluster1-01_clus2	cluster01-
02_clus2 none		
cluster1-02		
3/5/2022 19:21:18 -06:00	cluster1-02_clus2	cluster1-02_clus1
none		
3/5/2022 19:21:20 -06:00	cluster1-02_clus2	cluster1-02_clus2
none		

All ONTAP releases

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

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



```

cluster1::~*> cluster ping-cluster -node local
Host is 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 e0b
Cluster cluster1-02_clus1 169.254.3.8 cluster1-02 e0a
Cluster cluster1-02_clus2 169.254.3.9 cluster1-02 e0b
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)

```

6. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

What's next?

[Enable SSH.](#)

Enable SSH on BES-53248 cluster switches

If you are using the Ethernet Switch Health Monitor (CSHM) and log collection features, you must generate the SSH keys and then enable SSH on the cluster switches.

Steps

1. Verify that SSH is disabled:

```
show ip ssh
```

Show example

```
(switch)# show ip ssh

SSH Configuration

Administrative Mode: ..... Disabled
SSH Port: ..... 22
Protocol Level: ..... Version 2
SSH Sessions Currently Active: ..... 0
Max SSH Sessions Allowed: ..... 5
SSH Timeout (mins): ..... 5
Keys Present: ..... DSA(1024) RSA(1024)
ECDSA(521)
Key Generation In Progress: ..... None
SSH Public Key Authentication Mode: ..... Disabled
SCP server Administrative Mode: ..... Disabled
```

2. Generate the SSH keys:

```
crypto key generate
```

Show example

```
(switch)# config

(switch) (Config)# crypto key generate rsa

Do you want to overwrite the existing RSA keys? (y/n): y

(switch) (Config)# crypto key generate dsa

Do you want to overwrite the existing DSA keys? (y/n): y

(switch) (Config)# crypto key generate ecdsa 521

Do you want to overwrite the existing ECDSA keys? (y/n): y

(switch) (Config)# aaa authorization commands "noCmdAuthList" none
(switch) (Config)# exit
(switch)# ip ssh server enable
(switch)# ip scp server enable
(switch)# ip ssh pubkey-auth
(switch)# write mem

This operation may take a few minutes.
Management interfaces will not be available during this time.
Are you sure you want to save? (y/n) y

Config file 'startup-config' created successfully.

Configuration Saved!
```



Make sure that SSH is disabled before modifying the keys otherwise a warning is reported on the switch.

3. Encrypt the SSH keys (for **FIPS-mode only**):



In FIPS mode, the keys are required to be encrypted with a passphrase for security. In the absence of an encrypted key, the application fails to start. The keys are created and encrypted using the following commands:

Show example

```
(switch) configure  
(switch) (Config)# crypto key encrypt write rsa passphrase  
<passphrase>
```

The key will be encrypted and saved on NVRAM.
This will result in saving all existing configuration also.
Do you want to continue? (y/n): **y**

Config file 'startup-config' created successfully.

```
(switch) (Config)# crypto key encrypt write dsa passphrase  
<passphrase>
```

The key will be encrypted and saved on NVRAM.
This will result in saving all existing configuration also.
Do you want to continue? (y/n): **y**

Config file 'startup-config' created successfully.

```
(switch) (Config)# crypto key encrypt write ecdsa passphrase  
<passphrase>
```

The key will be encrypted and saved on NVRAM.
This will result in saving all existing configuration also.
Do you want to continue? (y/n): **y**

Config file 'startup-config' created successfully.

```
(switch) (Config)# end  
(switch)# write memory
```

This operation may take a few minutes.
Management interfaces will not be available during this time.
Are you sure you want to save? (y/n) **y**

Config file 'startup-config' created successfully.

Configuration Saved!

4. Reboot the switch:

```
reload
```

5. Verify that SSH is enabled:

```
show ip ssh
```

Show example

```
(switch)# show ip ssh

SSH Configuration

Administrative Mode: ..... Enabled
SSH Port: ..... 22
Protocol Level: ..... Version 2
SSH Sessions Currently Active: ..... 0
Max SSH Sessions Allowed: ..... 5
SSH Timeout (mins): ..... 5
Keys Present: ..... DSA(1024) RSA(1024)
ECDSA(521)
Key Generation In Progress: ..... None
SSH Public Key Authentication Mode: ..... Enabled
SCP server Administrative Mode: ..... Enabled
```

What's next?

[Configure switch health monitoring.](#)

Upgrade switch

Overview of the upgrade process for BES-53248 cluster switches

Follow these steps to upgrade the EFOS software and reference configuration files (RCFs) on Broadcom BES-54328 cluster switches, as applicable.

1. [Upgrade your EFOS version](#)

Download and install the Ethernet Fabric OS (EFOS) software on the BES-53248 cluster switch.

2. [Upgrade your RCF version](#)

Upgrade the RCF on the BES-53248 cluster switch, and then verify the ports for an additional license after the RCF is applied.

3. [Verify the ONTAP cluster network after upgrade](#)

Verify the health of the ONTAP cluster network after an upgrade of the EFOS software or RCF for BES-53248 cluster switches.

Upgrade the EFOS software

Follow these steps to upgrade the EFOS software on the BES-53248 cluster switch.

EFOS software includes a set of advanced networking features and protocols for developing Ethernet and IP infrastructure systems. This software architecture is suitable for any network organizational device using applications that require thorough packet inspection or separation.

Prepare for upgrade

Before you begin

- Download the applicable Broadcom EFOS software for your cluster switches from the [Broadcom Ethernet Switch Support](#) site.
- Review the following notes regarding EFOS versions.

Note the following:

- When upgrading from EFOS 3.4.x.x to EFOS 3.7.x.x or later, the switch must be running EFOS 3.4.4.6 (or later 3.4.x.x release). If you are running a release prior to that, then upgrade the switch to EFOS 3.4.4.6 (or later 3.4.x.x release) first, then upgrade the switch to EFOS 3.7.x.x or later.
- The configuration for EFOS 3.4.x.x and 3.7.x.x or later are different. Changing the EFOS version from 3.4.x.x to 3.7.x.x or later, or vice versa, requires the switch to be reset to factory defaults and the RCF files for the corresponding EFOS version to be (re)applied. This procedure requires access through the serial console port.
- Beginning with EFOS version 3.7.x.x or later, a non-FIPS compliant and a FIPS compliant version is available. Different steps apply when moving from a non-FIPS compliant to a FIPS compliant version or vice versa. Changing EFOS from a non-FIPS compliant to a FIPS compliant version or vice versa will reset the switch to factory defaults. This procedure requires access through the serial console port.

Procedure	Current EFOS version	New EFOS version	High level steps
-----------	----------------------	------------------	------------------

Steps to upgrade EFOS between two (non) FIPS compliant versions	3.4.x.x	3.4.x.x	Upgrade the new EFOS image using Method 1: Upgrade EFOS . The configuration and license information is retained.
	3.4.4.6 (or later 3.4.x.x)	3.7.x.x or later non-FIPS compliant	Upgrade EFOS using Method 1: Upgrade EFOS . Reset the switch to factory defaults and apply the RCF file for EFOS 3.7.x.x or later.
	3.7.x.x or later non-FIPS compliant	3.4.4.6 (or later 3.4.x.x)	Downgrade EFOS using Method 1: Upgrade EFOS . Reset the switch to factory defaults and apply the RCF file for EFOS 3.4.x.x
		3.7.x.x or later non-FIPS compliant	Upgrade the new EFOS image using Method 1: Upgrade EFOS . The configuration and license information is retained.
	3.7.x.x or later FIPS compliant	3.7.x.x or later FIPS compliant	Upgrade the new EFOS image using Method 1: Upgrade EFOS . The configuration and license information is retained.
Steps to upgrade to/from a FIPS compliant EFOS version	Non-FIPS compliant	FIPS compliant	Upgrade of the EFOS image using Method 2: Upgrade EFOS using the ONIE OS installation . The switch configuration and license information will be lost.
	FIPS compliant	Non-FIPS compliant	

To check if your version of EFOS is FIPS compliant or non-FIPS compliant, use the `show fips status` command. In the following examples, **IP_switch_a1** is using FIPS compliant EFOS and **IP_switch_a2** is using non-FIPS compliant EFOS.

- On switch **IP_switch_a1** (FIPS compliant EFOS):

```
IP_switch_a1 # show fips status

System running in FIPS mode
```

- On switch **IP_switch_a2** (non-FIPS compliant EFOS):

```
IP_switch_a2 # show fips status
                ^
% Invalid input detected at ^ marker.
```

Upgrade the software

Use one of the following methods:

- [Method 1: Upgrade EFOS](#). Use for most cases (see the table above).
- [Method 2: Upgrade EFOS using the ONIE OS installation](#). Use if one EFOS version is FIPS compliant and the other EFOS version is non-FIPS compliant.



Upgrade EFOS on one switch at a time to ensure continued cluster network operation.

Method 1: Upgrade EFOS

Perform the following steps to upgrade the EFOS software.



Note that after upgrading BES-53248 cluster switches from EFOS 3.3.x.x or 3.4.x.x to EFOS 3.7.0.4 or 3.8.0.2, Inter-Switch Links (ISLs) and port channels are marked in the **Down** state. This is expected behavior and it's safe to continue with the upgrade unless you are having issues with auto-reverting LIFs. See the Knowledge Base article: [BES-53248 Cluster Switch NDU failed upgrade to EFOS 3.7.0.4 and later](#) for further details.

Steps

1. Connect the BES-53248 cluster switch to the management network.
2. Use the `ping` command to verify connectivity to the server hosting EFOS, licenses, and the RCF file.

This example verifies that the switch is connected to 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.
```

3. Disable auto-revert on the cluster LIFs.

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

4. Display the boot images for the active and backup configuration:

```
show bootvar
```

Show example

```
(cs2)# show bootvar

Image Descriptions

active :
backup :

Images currently available on Flash
-----
unit      active      backup      current-active      next-active
-----
1         3.7.0.4     3.4.4.6     3.7.0.4             3.7.0.4
```

5. Download the image file to the switch.

Copying the image file to the active image means that when you reboot, that image establishes the running EFOS version. The previous image remains available as a backup.

```
(cs2)# copy sftp://root@172.19.2.1//tmp/EFOS-3.10.0.3.stk backup
Remote Password:**

Mode..... SFTP
Set Server IP..... 172.19.2.1
Path..... //tmp/
Filename..... EFOS-3.10.0.3.stk
Data Type..... Code
Destination Filename..... backup

Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y
SFTP Code transfer starting...

File transfer operation completed successfully.
```

6. Display the boot images for the active and backup configuration:

```
show bootvar
```

Show example

```
(cs2)# show bootvar

Image Descriptions

active :
backup :

Images currently available on Flash
-----
unit      active      backup      current-active      next-active
-----
1         3.7.0.4     3.7.0.4     3.7.0.4             3.10.0.3
```

7. Boot the system from the backup configuration:

```
boot system backup
```

```
(cs2)# boot system backup
Activating image backup ..
```

8. Display the boot images for the active and backup configuration:

```
show bootvar
```

Show example

```
(cs2)# show bootvar
```

```
Image Descriptions
```

```
active :
```

```
backup :
```

```
Images currently available on Flash
```

```
-----  
unit      active      backup      current-active      next-active  
-----  
1         3.10.0.3      3.10.0.3      3.10.0.3            3.11.0.2
```

9. Save the running configuration to the startup configuration:

```
write memory
```

Show example

```
(cs2)# write memory
```

```
This operation may take a few minutes.
```

```
Management interfaces will not be available during this time.
```

```
Are you sure you want to save? (y/n) y
```

```
Config file 'startup-config' created successfully.
```

```
Configuration Saved!
```

10. Reboot the switch:

```
reload
```

Show example

```
(cs2)# reload
```

```
The system has unsaved changes.
```

```
Would you like to save them now? (y/n) y
```

```
Config file 'startup-config' created successfully.
```

```
Configuration Saved!
```

```
System will now restart!
```

11. Log in again and verify the new version of the EFOS software:

```
show version
```

Show example

```
(cs2)# show version
```

```
Switch: 1
```

```
System Description..... BES-53248A1,  
3.10.0.3, Linux 4.4.211-28a6fe76, 2016.05.00.04
```

```
Machine Type..... BES-53248A1,
```

```
Machine Model..... BES-53248
```

```
Serial Number..... QTFCU38260023
```

```
Maintenance Level..... A
```

```
Manufacturer..... 0xbc00
```

```
Burned In MAC Address..... D8:C4:97:71:0F:40
```

```
Software Version..... 3.10.0.3
```

```
Operating System..... Linux 4.4.211-  
28a6fe76
```

```
Network Processing Device..... BCM56873_A0
```

```
CPLD Version..... 0xff040c03
```

```
Additional Packages..... BGP-4
```

```
..... QOS
```

```
..... Multicast
```

```
..... IPv6
```

```
..... Routing
```

```
..... Data Center
```

```
..... OpEN API
```

```
..... Prototype Open API
```

12. Repeat steps 5 through to 11 on the switch cs1.
13. Enable auto-revert on the cluster LIFs.

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

14. Verify that the cluster LIFs have reverted to their home port:

```
network interface show -role Cluster
```

For further details, see [Revert a LIF to its home port](#).

Method 2: Upgrade EFOS using the ONIE OS installation

You can perform the following steps if one EFOS version is FIPS compliant and the other EFOS version is non-FIPS compliant. These steps can be used to upgrade the non-FIPS or FIPS compliant EFOS 3.7.x.x image from ONIE if the switch fails to boot.



This functionality is only available for EFOS 3.7.x.x or later non-FIPS compliant.



If you upgrade EFOS using the ONIE OS installation, the configuration is reset to factory defaults and licenses are deleted. You must set up the switch and install licenses and a supported RCF to return the switch to normal operation.

Steps

1. Disable auto-revert on the cluster LIFs.

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

2. Boot the switch into ONIE installation mode.

During boot, select ONIE when you see the prompt:

```
+-----+
| EFOS                                     |
| *ONIE                                   |
|                                         |
|                                         |
|                                         |
|                                         |
|                                         |
|                                         |
|                                         |
|                                         |
|                                         |
+-----+
```

After you select **ONIE**, the switch loads and presents you with several choices. Select **Install OS**.

```
+-----+
| *ONIE: Install OS                       |
| ONIE: Rescue                             |
| ONIE: Uninstall OS                       |
| ONIE: Update ONIE                       |
| ONIE: Embed ONIE                         |
| DIAG: Diagnostic Mode                    |
| DIAG: Burn-In Mode                       |
|                                         |
|                                         |
|                                         |
|                                         |
+-----+
```

The switch boots into ONIE installation mode.

3. Stop the ONIE discovery and configure the Ethernet interface.

When the following message appears, press **Enter** to invoke the ONIE console:

```
Please press Enter to activate this console. Info: eth0: Checking
link... up.
ONIE:/ #
```



The ONIE discovery continues and messages are printed to the console.

```
Stop the ONIE discovery
ONIE:/ # onie-discovery-stop
discover: installer mode detected.
Stopping: discover... done.
ONIE:/ #
```

4. Configure the Ethernet interface and add the route using `ifconfig eth0 <ipAddress> netmask <netmask> up` and `route add default gw <gatewayAddress>`

```
ONIE:/ # ifconfig eth0 10.10.10.10 netmask 255.255.255.0 up
ONIE:/ # route add default gw 10.10.10.1
```

5. Verify that the server hosting the ONIE installation file is reachable:

```
ping
```

Show example

```
ONIE:/ # ping 50.50.50.50
PING 50.50.50.50 (50.50.50.50): 56 data bytes
64 bytes from 50.50.50.50: seq=0 ttl=255 time=0.429 ms
64 bytes from 50.50.50.50: seq=1 ttl=255 time=0.595 ms
64 bytes from 50.50.50.50: seq=2 ttl=255 time=0.369 ms
^C
--- 50.50.50.50 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.369/0.464/0.595 ms
ONIE:/ #
```

6. Install the new switch software:

```
ONIE:/ # onie-nos-install http://50.50.50.50/Software/onie-installer-x86\_64
```

Show example

```
ONIE:/ # onie-nos-install http://50.50.50.50/Software/onie-
installer-x86_64
discover: installer mode detected.
Stopping: discover... done.
Info: Fetching http://50.50.50.50/Software/onie-installer-3.7.0.4
...
Connecting to 50.50.50.50 (50.50.50.50:80)
installer          100% |*****| 48841k
0:00:00 ETA
ONIE: Executing installer: http://50.50.50.50/Software/onie-
installer-3.7.0.4
Verifying image checksum ... OK.
Preparing image archive ... OK.
```

The software installs and then reboots the switch. Let the switch reboot normally into the new EFOS version.

7. Verify that the new switch software is installed:

```
show bootvar
```

Show example

```
(cs2)# show bootvar
Image Descriptions
active :
backup :
Images currently available on Flash
-----
unit   active      backup      current-active  next-active
-----
  1    3.7.0.4      3.7.0.4     3.7.0.4         3.10.0.3
(cs2) #
```

8. Complete the installation. The switch reboots with no configuration applied and resets to factory defaults. Complete the following steps to reconfigure the switch:

- a. [Install licenses](#)
- b. [Install the RCF](#)
- c. [Enable SSH](#)
- d. [Enable log collection](#)

e. [Configure SNMPv3 for monitoring](#)

9. Repeat steps 2 through to 8 on the switch cs1.
10. Enable auto-revert on the cluster LIFs.

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

11. Verify that the cluster LIFs have reverted to their home port:

```
network interface show -role Cluster
```

For further details, see [Revert a LIF to its home port](#).

Upgrade the Reference Configuration File (RCF)

You can upgrade the Reference Configuration File (RCF) after upgrading the BES-53248 cluster switch EFOS and after applying any new licenses.

What you'll need

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).
- The current RCF file, available from the [Broadcom Cluster Switches](#) page.
- A boot configuration in the RCF that reflects the desired boot images, required if you are installing only EFOS 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.
- A console connection to the switch, required when installing the RCF from a factory-default state. This requirement is optional if you have used the Knowledge Base article [How to clear configuration on a Broadcom interconnect switch while retaining remote connectivity](#) to clear the configuration, beforehand.

Suggested documentation

- Consult the switch compatibility table for the supported ONTAP and RCF versions. See the [EFOS Software download](#) page. Note that there can be command dependencies between the command syntax in the RCF and that found in versions of EFOS.
- Refer to the appropriate software and upgrade guides available on the [Broadcom](#) site for complete documentation on the BES-53248 switch upgrade and downgrade procedures.

About the examples

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

- The names of the two BES-53248 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.

- The examples in this procedure use four nodes. These nodes use two 10GbE cluster interconnect ports e0a and e0b. See the [Hardware Universe](#) to verify the correct cluster ports on your platforms.



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

About this task

The procedure requires the use of both ONTAP commands and Broadcom switch commands; ONTAP commands are used unless otherwise indicated.

No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure migrates all the cluster LIFs to the operational partner switch while performing the steps on the target switch.



Before installing a new switch software version and RCFs, use the Knowledge Base article [How to clear configuration on a Broadcom interconnect switch while retaining remote connectivity](#). If you must erase the switch settings completely, then you will need to perform the basic configuration again. You must be connected to the switch using the serial console, since a complete configuration erasure resets the configuration of the management network.

Step 1: Prepare for upgrade

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
```

where *x* is the duration of the maintenance window in hours.



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

The following command suppresses automatic case creation for two hours:

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

2. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (*>) appears.

3. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
              e0a    cs1                        0/2        BES-
53248
              e0b    cs2                        0/2        BES-
53248
cluster1-02/cdp
              e0a    cs1                        0/1        BES-
53248
              e0b    cs2                        0/1        BES-
53248
cluster1-03/cdp
              e0a    cs1                        0/4        BES-
53248
              e0b    cs2                        0/4        BES-
53248
cluster1-04/cdp
              e0a    cs1                        0/3        BES-
53248
              e0b    cs2                        0/3        BES-
53248
cluster1::*>
```

4. 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
```

Show example

```
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					
e0b	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					
e0b	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
```

```
Health Health Speed (Mbps)
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
```

Show example

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

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

Cluster				
cluster1-01	cluster1-01_clus1	up/up	169.254.3.4/23	
	e0a true			
cluster1-01	cluster1-01_clus2	up/up	169.254.3.5/23	
	e0b true			
cluster1-02	cluster1-02_clus1	up/up	169.254.3.8/23	
	e0a true			
cluster1-02	cluster1-02_clus2	up/up	169.254.3.9/23	
	e0b true			
cluster1-03	cluster1-03_clus1	up/up	169.254.1.3/23	
	e0a true			
cluster1-03	cluster1-03_clus2	up/up	169.254.1.1/23	
	e0b true			
cluster1-04	cluster1-04_clus1	up/up	169.254.1.6/23	
	e0a true			
cluster1-04	cluster1-04_clus2	up/up	169.254.1.7/23	
	e0b true			

5. Verify that the cluster displays information for both cluster switches.

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command:

```
system switch ethernet show -is-monitoring-enabled-operational true
```

```
cluster1::*> system switch ethernet show -is-monitoring-enabled  
-operational true
```

Switch	Type	Address	Model
cs1 53248	cluster-network	10.228.143.200	BES-
Serial Number: QTWCU22510008			
Is Monitored: true			
Reason: None			
Software Version: 3.10.0.3			
Version Source: CDP/ISDP			
cs2 53248	cluster-network	10.228.143.202	BES-
Serial Number: QTWCU22510009			
Is Monitored: true			
Reason: None			
Software Version: 3.10.0.3			
Version Source: CDP/ISDP			

```
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command:

```
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                     10.228.143.200                       BES-
53248
      Serial Number: QTWCU22510008
      Is Monitored: true
      Reason: None
      Software Version: 3.10.0.3
      Version Source: CDP/ISDP

cs2                                    cluster-network                     10.228.143.202                       BES-
53248
      Serial Number: QTWCU22510009
      Is Monitored: true
      Reason: None
      Software Version: 3.10.0.3
      Version Source: CDP/ISDP
cluster1::*>

```

6. Disable auto-revert on the cluster LIFs.

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

Step 2: Configure ports

1. On switch cs2, confirm the list of ports that are connected to the nodes in the cluster.

```
show isdp neighbor
```

2. On switch cs2, shut down the ports connected to the cluster ports of the nodes. For example, if ports 0/1 to 0/16 are connected to ONTAP nodes:

```

(cs2)> enable
(cs2)# configure
(cs2) (Config)# interface 0/1-0/16
(cs2) (Interface 0/1-0/16)# shutdown
(cs2) (Interface 0/1-0/16)# exit
(cs2) (Config)#

```


3. Verify that the cluster LIFs have migrated to the ports hosted on cluster switch cs1. This might take a few seconds.

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

Show example

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

Current Is	Logical	Status	Network	Current
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		

```
cluster1::*>
```

4. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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
```

5. If you have not already done so, save the current switch configuration by copying the output of the following command to a log file:

```
show running-config
```

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



When updating or applying a new RCF, you must erase the switch settings and perform basic configuration. You must be connected to the switch using the serial console to erase switch settings. This requirement is optional if you have used the Knowledge Base article [How to clear the configuration on a Broadcom interconnect switch while retaining remote connectivity](#) to clear the configuration, beforehand.



Clearing the configuration does not delete licenses.

- a. SSH into the switch.

Only proceed when all the cluster LIFs have been removed from the ports on the switch and the switch is prepared to have the configuration cleared.

- b. Enter privilege mode:

```
(cs2)> enable
(cs2) #
```

- c. Copy and paste the following commands to remove the previous RCF configuration (depending on the previous RCF version used, some commands might generate an error if a particular setting is not present):

```
clear config interface 0/1-0/56
Y
clear config interface lag 1
Y
configure
```

```
deleteport 1/1 all
no policy-map CLUSTER
no policy-map WRED_25G
no policy-map WRED_100G
no policy-map InShared
no policy-map InMetroCluster
no policy-map InCluster
no policy-map InClusterRdma
no class-map CLUSTER
no class-map HA
no class-map RDMA
no class-map c5
no class-map c4
no class-map CLUSTER
no class-map CLUSTER_RDMA
no class-map StorageSrc
no class-map StorageDst
no class-map RdmaSrc
no class-map RdmaDstA
no classofservice dot1p-mapping
no random-detect queue-parms 0
no random-detect queue-parms 1
no random-detect queue-parms 2
no random-detect queue-parms 3
no random-detect queue-parms 4
no random-detect queue-parms 5
no random-detect queue-parms 6
no random-detect queue-parms 7
no cos-queue min-bandwidth
no cos-queue random-detect 0
no cos-queue random-detect 1
no cos-queue random-detect 2
no cos-queue random-detect 3
no cos-queue random-detect 4
no cos-queue random-detect 5
no cos-queue random-detect 6
no cos-queue random-detect 7
exit
vlan database
no vlan 17
no vlan 18
exit
show-running-config
```

d. Save the running configuration to the startup configuration:

write memory

```
(cs2)# write memory
```

```
This operation may take a few minutes.  
Management interfaces will not be available during this time.
```

```
Are you sure you want to save? (y/n) y
```

```
Config file 'startup-config' created successfully.  
Configuration Saved!
```

e. Perform a reboot of the switch:

reload

```
(cs2)# reload
```

```
Are you sure you would like to reset the system? (y/n) y
```

f. Log in to the switch again using SSH to complete the RCF installation.

7. Note the following:

- a. If additional port licenses have been installed on the switch, you must modify the RCF to configure the additional licensed ports. See [Activate newly licensed ports](#) for details.
- b. Record any customizations that were made in the previous RCF and apply these to the new RCF. For example, setting port speeds or hard-coding FEC mode.

EFOS version 3.12.x and later

- Copy the RCF to the bootflash of switch cs2 using one of the following transfer protocols: FTP, TFTP, SFTP, or SCP.

This example shows SFTP being used to copy an RCF to the bootflash on switch cs2:

```
(cs2)# copy sftp://172.19.2.1/BES-53248-RCF-v1.9-Cluster-HA.txt
nvram:reference-config.scr
Remote Password:**
Mode..... TFTP
Set Server IP..... 172.19.2.1
Path..... /
Filename..... BES-53248_RCF_v1.9-
Cluster-HA.txt
Data Type..... Config Script
Destination Filename..... reference-config.scr
Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y
TFTP Code transfer starting...
File transfer operation completed successfully.
```

- Verify that the script was downloaded and saved under the file name you gave it:

```
script list
```

```
(cs2)# script list

Configuration Script Name                Size(Bytes)  Date of
Modification
-----
reference-config.scr                    2680        2024 05 31
21:54:22
2 configuration script(s) found.
2042 Kbytes free.
```

- Apply the script to the switch:

```
script apply
```

```
(cs2)# script apply reference-config.scr

Are you sure you want to apply the configuration script? (y/n) y

The system has unsaved changes.
Would you like to save them now? (y/n) y
Config file 'startup-config' created successfully.
Configuration Saved!

Configuration script 'reference-config.scr' applied.
```

All other EFOS versions

8. Copy the RCF to the bootflash of switch cs2 using one of the following transfer protocols: FTP, TFTP, SFTP, or SCP.

This example shows SFTP being used to copy an RCF to the bootflash on switch cs2:

```
(cs2)# copy sftp://172.19.2.1/tmp/BES-53248_RCF_v1.9-Cluster-HA.txt
nvram:script BES-53248_RCF_v1.9-Cluster-HA.scr
Remote Password:**
Mode..... SFTP
Set Server IP..... 172.19.2.1
Path..... //tmp/
Filename..... BES-53248_RCF_v1.9-
Cluster-HA.txt
Data Type..... Config Script
Destination Filename..... BES-53248_RCF_v1.9-
Cluster-HA.scr
Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y
SFTP Code transfer starting...
File transfer operation completed successfully.
```

9. Verify that the script was downloaded and saved to the file name you gave it:

```
script list
```

```
(cs2)# script list
```

Configuration Script Name Modification	Size(Bytes)	Date of Modification
----- -----	-----	-----
BES-53248_RCF_v1.9-Cluster-HA.scr 05:41:00	2241	2020 09 30

```
1 configuration script(s) found.
```

10. Apply the script to the switch:

```
script apply
```

```
(cs2)# script apply BES-53248_RCF_v1.9-Cluster-HA.scr
```

```
Are you sure you want to apply the configuration script? (y/n) y
```

```
The system has unsaved changes.
```

```
Would you like to save them now? (y/n) y
```

```
Config file 'startup-config' created successfully.
```

```
Configuration Saved!
```

```
Configuration script 'BES-53248_RCF_v1.9-Cluster-HA.scr' applied.
```

11. Examine the banner output from the `show clibanner` command. You must read and follow these instructions to ensure the proper configuration and operation of the switch.

```
show clibanner
```

Show example

```
(cs2)# show clibanner
```

```
Banner Message configured :
```

```
=====
```

```
BES-53248 Reference Configuration File v1.9 for Cluster/HA/RDMA
```

```
Switch    : BES-53248
```

```
Filename  : BES-53248-RCF-v1.9-Cluster.txt
```

```
Date      : 10-26-2022
```

```
Version   : v1.9
```

```
Port Usage:
```

```
Ports 01 - 16: 10/25GbE Cluster Node Ports, base config
```

```
Ports 17 - 48: 10/25GbE Cluster Node Ports, with licenses
```

```
Ports 49 - 54: 40/100GbE Cluster Node Ports, with licenses, added  
right to left
```

```
Ports 55 - 56: 100GbE Cluster ISL Ports, base config
```

```
NOTE:
```

```
- The 48 SFP28/SFP+ ports are organized into 4-port groups in terms  
of port
```

```
speed:
```

```
Ports 1-4, 5-8, 9-12, 13-16, 17-20, 21-24, 25-28, 29-32, 33-36, 37-  
40, 41-44,  
45-48
```

```
The port speed should be the same (10GbE or 25GbE) across all ports  
in a 4-port
```

```
group
```

```
- If additional licenses are purchased, follow the 'Additional Node  
Ports
```

```
activated with Licenses' section for instructions
```

```
- If SSH is active, it will have to be re-enabled manually after  
'erase
```

```
startup-config'
```

```
command has been executed and the switch rebooted
```

12. On the switch, verify that the additional licensed ports appear after the RCF is applied:

```
show port all | exclude Detach
```


Show example

```
(cs2)# show port all | exclude Detach
```

LACP	Actor	Admin	Physical	Physical	Link	Link
Intf	Type	Mode	Mode	Status	Status	Trap
Mode	Timeout					
0/1	Enable long	Enable	Auto		Down	Enable
0/2	Enable long	Enable	Auto		Down	Enable
0/3	Enable long	Enable	Auto		Down	Enable
0/4	Enable long	Enable	Auto		Down	Enable
0/5	Enable long	Enable	Auto		Down	Enable
0/6	Enable long	Enable	Auto		Down	Enable
0/7	Enable long	Enable	Auto		Down	Enable
0/8	Enable long	Enable	Auto		Down	Enable
0/9	Enable long	Enable	Auto		Down	Enable
0/10	Enable long	Enable	Auto		Down	Enable
0/11	Enable long	Enable	Auto		Down	Enable
0/12	Enable long	Enable	Auto		Down	Enable
0/13	Enable long	Enable	Auto		Down	Enable
0/14	Enable long	Enable	Auto		Down	Enable
0/15	Enable long	Enable	Auto		Down	Enable
0/16	Enable long	Enable	Auto		Down	Enable
0/49	Enable long	Enable	40G Full		Down	Enable
0/50	Enable long	Enable	40G Full		Down	Enable

```

0/51          Enable    100G Full          Down    Enable
Enable long
0/52          Enable    100G Full          Down    Enable
Enable long
0/53          Enable    100G Full          Down    Enable
Enable long
0/54          Enable    100G Full          Down    Enable
Enable long
0/55          Enable    100G Full          Down    Enable
Enable long
0/56          Enable    100G Full          Down    Enable
Enable long

```

13. Verify on the switch that your changes have been made.

```
show running config
```

14. Save the running configuration so that it becomes the startup configuration when you reboot the switch:

```
write memory
```

Show example

```

(cs2)# write memory
This operation may take a few minutes.
Management interfaces will not be available during this time.

Are you sure you want to save? (y/n) y

Config file 'startup-config' created successfully.
Configuration Saved!

```

15. Reboot the switch and verify that the running configuration is correct.

```
reload
```

```

(cs2)# reload
Are you sure you would like to reset the system? (y/n) y
System will now restart!

```

16. On cluster switch cs2, bring up the ports connected to the cluster ports of the nodes.

```
(cs2)> enable
(cs2)# configure
(cs2) (Config)# interface 0/1-0/16
(cs2) (Interface 0/1-0/16)# no shutdown
(cs2) (Config)# exit
```

17. Save the running configuration to the startup configuration:

```
write memory
```

Show example

```
(cs2)# write memory

This operation may take a few minutes.
Management interfaces will not be available during this time.

Are you sure you want to save? (y/n) y

Config file 'startup-config' created successfully.
Configuration Saved!
```

18. Verify the ports on switch cs2:

```
show interfaces status all | exclude Detach
```

Show example

```
(cs1)# show interfaces status all | exclude Detach
```

Media	Flow	Link	Physical	Physical	
Port	Name	State	Mode	Status	Type
Control	VLAN				
-----	-----	-----	-----	-----	
-----	-----	-----			
.					
.					
.					
0/16	10/25GbE Node Port	Down	Auto		
Inactive	Trunk				
0/17	10/25GbE Node Port	Down	Auto		
Inactive	Trunk				
0/18	10/25GbE Node Port	Up	25G Full	25G Full	
25GBase-SR	Inactive Trunk				
0/19	10/25GbE Node Port	Up	25G Full	25G Full	
25GBase-SR	Inactive Trunk				
.					
.					
.					
0/50	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/51	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/52	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/53	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/54	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/55	Cluster ISL Port	Up	Auto	100G Full	
Copper	Inactive Trunk				
0/56	Cluster ISL Port	Up	Auto	100G Full	
Copper	Inactive Trunk				

19. Verify the health of cluster ports on the cluster.

a. Verify that e0b ports are up and healthy across all nodes in the cluster:

```
network port show -role cluster
```

Show example

```
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						
e0b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

```
Node: cluster1-04
```

```
Ignore
```

```
Health Health Speed (Mbps)
Port IPspace Broadcast Domain Link MTU Admin/Oper
Status Status
-----
e0a Cluster Cluster up 9000 auto/100000
healthy false
e0b Cluster Cluster up 9000 auto/100000
healthy false
```

b. Verify the switch health from the cluster:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
          e0a   cs1                        0/2
BES-53248
          e0b   cs2                        0/2
BES-53248
cluster01-2/cdp
          e0a   cs1                        0/1
BES-53248
          e0b   cs2                        0/1
BES-53248
cluster01-3/cdp
          e0a   cs1                        0/4
BES-53248
          e0b   cs2                        0/4
BES-53248
cluster1-04/cdp
          e0a   cs1                        0/3
BES-53248
          e0b   cs2                        0/2
BES-53248
```

20. Verify that the cluster displays information for both cluster switches.

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command:

```
system switch ethernet show -is-monitoring-enabled-operational true
```

```
cluster1::*> system switch ethernet show -is-monitoring-enabled  
-operational true
```

Switch	Type	Address	Model
cs1 53248	cluster-network	10.228.143.200	BES-
Serial Number: QTWCU22510008			
Is Monitored: true			
Reason: None			
Software Version: 3.10.0.3			
Version Source: CDP/ISDP			
cs2 53248	cluster-network	10.228.143.202	BES-
Serial Number: QTWCU22510009			
Is Monitored: true			
Reason: None			
Software Version: 3.10.0.3			
Version Source: CDP/ISDP			

```
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command:

```
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                    10.228.143.200                       BES-
53248
    Serial Number: QTWCU22510008
    Is Monitored: true
    Reason: None
    Software Version: 3.10.0.3
    Version Source: CDP/ISDP

cs2                                   cluster-network                    10.228.143.202                       BES-
53248
    Serial Number: QTWCU22510009
    Is Monitored: true
    Reason: None
    Software Version: 3.10.0.3
    Version Source: CDP/ISDP
cluster1::*>

```

21. Repeat steps 1 to 20 on switch cs1.
22. Enable auto-revert on the cluster LIFs:

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

23. . Verify that the cluster LIFs have reverted to their home port:

```
network interface show -role Cluster
```

For further details, see [Revert a LIF to its home port](#).

Step 3: Verify the configuration

1. On switch cs1, verify that the switch ports connected to the cluster ports are **up**:

```
show interfaces status all
```

Show example

```
(cs1)# show interfaces status all | exclude Detach
```

Media	Flow	Link	Physical	Physical	
Port	Name	State	Mode	Status	Type
Control	VLAN				
-----	-----	-----	-----	-----	
-----	-----	-----			
.					
.					
.					
0/16	10/25GbE Node Port	Down	Auto		
Inactive	Trunk				
0/17	10/25GbE Node Port	Down	Auto		
Inactive	Trunk				
0/18	10/25GbE Node Port	Up	25G Full	25G Full	
25GBase-SR	Inactive Trunk				
0/19	10/25GbE Node Port	Up	25G Full	25G Full	
25GBase-SR	Inactive Trunk				
.					
.					
.					
0/50	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/51	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/52	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/53	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/54	40/100GbE Node Port	Down	Auto		
Inactive	Trunk				
0/55	Cluster ISL Port	Up	Auto	100G Full	
Copper	Inactive Trunk				
0/56	Cluster ISL Port	Up	Auto	100G Full	
Copper	Inactive Trunk				

2. Verify that the ISL between switches cs1 and cs2 is functional:

```
show port-channel 1/1
```

Show example

```
(cs1)# show port-channel 1/1
Local Interface..... 1/1
Channel Name..... Cluster-ISL
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Port-channel Min-links..... 1
Load Balance Option..... 7
(Enhanced hashing mode)
Mbr      Device/      Port      Port
Ports   Timeout      Speed     Active
-----  -
0/55    actor/long    Auto     True
        partner/long
0/56    actor/long    Auto     True
        partner/long
```

3. Verify that the cluster LIFs have reverted to their home port:

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

Show example

```
cluster1::*> network interface show -role cluster
          Logical          Status      Network          Current
Current Is
Vserver   Interface              Admin/Oper  Address/Mask     Node
Port      Home
-----
Cluster
cluster1-01 cluster1-01_clus1  up/up      169.254.3.4/23
          e0a          true
cluster1-01 cluster1-01_clus2  up/up      169.254.3.5/23
          e0b          true
cluster1-02 cluster1-02_clus1  up/up      169.254.3.8/23
          e0a          true
cluster1-02 cluster1-02_clus2  up/up      169.254.3.9/23
          e0b          true
cluster1-03 cluster1-03_clus1  up/up      169.254.1.3/23
          e0a          true
cluster1-03 cluster1-03_clus2  up/up      169.254.1.1/23
          e0b          true
cluster1-04 cluster1-04_clus1  up/up      169.254.1.6/23
          e0a          true
cluster1-04 cluster1-04_clus2  up/up      169.254.1.7/23
          e0b          true
```

4. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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
```

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
-----	-----	-----
-----	-----	-----
cluster1-01		
3/5/2022 19:21:18 -06:00	cluster1-01_clus2	cluster01-
02_clus1 none		
3/5/2022 19:21:20 -06:00	cluster1-01_clus2	cluster01-
02_clus2 none		
cluster1-02		
3/5/2022 19:21:18 -06:00	cluster1-02_clus2	cluster1-02_clus1
none		
3/5/2022 19:21:20 -06:00	cluster1-02_clus2	cluster1-02_clus2
none		

All ONTAP releases

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

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

```

cluster1::~*> cluster ping-cluster -node local
Host is 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 e0b
Cluster cluster1-02_clus1 169.254.3.8 cluster1-02 e0a
Cluster cluster1-02_clus2 169.254.3.9 cluster1-02 e0b
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)

```

6. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

Verify the ONTAP cluster network after an EFOS software or RCF upgrade of the BES-53248 cluster switches

You can use the following commands to verify the health of the ONTAP cluster network after an upgrade of the EFOS software or RCF for BES-53248 cluster switches.

Steps

1. Display information about the network ports on the cluster using the command:

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

Link must have the value up and Health Status must be healthy.

Show example

The following example shows the output from the command:

```
cluster1::> network port show -ipSpace Cluster

Node: node1

Ignore

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

Node: node2

Ignore

Health
Speed (Mbps) 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
```

2. For each LIF, verify that `Is Home` is true and `Status Admin/Oper` is up on both nodes, using the command:

```
network interface show -vserver Cluster
```


Show example

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

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

Cluster				
e0a	node1_clus1	up/up	169.254.217.125/16	node1
	true			
e0b	node1_clus2	up/up	169.254.205.88/16	node1
	true			
e0a	node2_clus1	up/up	169.254.252.125/16	node2
	true			
e0b	node2_clus2	up/up	169.254.110.131/16	node2
	true			

3. Verify that the Health Status of each node is true using the command:

```
cluster show
```

Show example

```
cluster1::> cluster show
```

Node	Health	Eligibility	Epsilon

node1	true	true	false
node2	true	true	false

What's next?

[Configure switch health monitoring.](#)

Migrate switches

Migrate CN1610 cluster switches to BES-53248 cluster switches

To migrate the CN1610 cluster switches in a cluster to Broadcom-supported BES-53248 cluster switches, review the migration requirements and then follow the migration procedure.

The following cluster switches are supported:

- CN1610
- BES-53248

Review requirements

Verify that your configuration meets the following requirements:

- Some of the ports on BES-53248 switches are configured to run at 10GbE.
- The 10GbE connectivity from nodes to BES-53248 cluster switches have been planned, migrated, and documented.
- The cluster is fully functioning (there should be no errors in the logs or similar issues).
- Initial customization of the BES-53248 switches is complete, so that:
 - BES-53248 switches are running the latest recommended version of EFOS software.
 - Reference Configuration Files (RCFs) have been applied to the switches.
 - Any site customization, such as DNS, NTP, SMTP, SNMP, and SSH, are configured on the new switches.

Node connections

The cluster switches support the following node connections:

- NetApp CN1610: ports 0/1 through 0/12 (10GbE)
- BES-53248: ports 0/1-0/16 (10GbE/25GbE)



Additional ports can be activated by purchasing port licenses.

ISL ports

The cluster switches use the following inter-switch link (ISL) ports:

- NetApp CN1610: ports 0/13 through 0/16 (10GbE)
- BES-53248: ports 0/55-0/56 (100GbE)

The [NetApp Hardware Universe](#) contains information about ONTAP compatibility, supported EFOS firmware, and cabling to BES-53248 cluster switches.

ISL cabling

The appropriate ISL cabling is as follows:

- **Beginning:** For CN1610 to CN1610 (SFP+ to SFP+), four SFP+ optical fiber or copper direct-attach cables.
- **Final:** For BES-53248 to BES-53248 (QSFP28 to QSFP28), two QSFP28 optical transceivers/fiber or copper direct-attach cables.

Migrate the switches

Follow this procedure to migrate CN1610 cluster switches to BES-53248 cluster switches.

About the examples

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

- The examples use two nodes, each deploying two 10 GbE cluster interconnect ports: e0a and e0b.
- The command outputs might vary depending on different releases of ONTAP software.
- The CN1610 switches to be replaced are CL1 and CL2.
- The BES-53248 switches to replace the CN1610 switches are cs1 and cs2.
- The nodes are node1 and node2.
- The switch CL2 is replaced by cs2 first, followed with CL1 by cs1.
- The BES-53248 switches are pre-loaded with the supported versions of Reference Configuration File (RCF) and Ethernet Fabric OS (EFOS) with ISL cables connected on ports 55 and 56.
- The cluster LIF names are node1_clus1 and node1_clus2 for node1, and node2_clus1 and node2_clus2 for node2.

About this task

This procedure covers the following scenario:

- The cluster starts with two nodes connected to two CN1610 cluster switches.
- CN1610 switch CL2 is replaced by BES-53248 switch cs2:
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - Disconnect the cables from all cluster ports on all nodes connected to CL2, and then use supported cables to reconnect the ports to the new cluster switch cs2.
- CN1610 switch CL1 is replaced by BES-53248 switch cs1:
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - Disconnect the cables from all cluster ports on all nodes connected to CL1, and then use supported cables to reconnect the ports to the new cluster switch cs1.



No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure migrates all of the cluster LIFs to the operational partner switch while performing the steps on the target switch.

Step 1: Prepare for migration

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
```

where x is the duration of the maintenance window in hours.



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

The following command suppresses automatic case creation for two hours:

```
cluster1::*> system node autosupport invoke -node * -type all -message  
MAINT=2h
```

2. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (*>) appears.

Step 2: Configure ports and cabling

1. On the new switches, confirm that the ISL is cabled and healthy between switches cs1 and cs2:

```
show port-channel
```

Show example

The following example shows that the ISL ports are **up** on switch cs1:

```
(cs1)# show port-channel 1/1
Local Interface..... 1/1
Channel Name..... Cluster-ISL
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Port channel Min-links..... 1
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr      Device/      Port      Port
Ports   Timeout     Speed     Active
-----
0/55    actor/long   100G Full  True
        partner/long
0/56    actor/long   100G Full  True
        partner/long
(cs1) #
```

The following example shows that the ISL ports are **up** on switch cs2:

```
(cs2)# show port-channel 1/1
Local Interface..... 1/1
Channel Name..... Cluster-ISL
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Port channel Min-links..... 1
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr      Device/      Port      Port
Ports   Timeout     Speed     Active
-----
0/55    actor/long   100G Full  True
        partner/long
0/56    actor/long   100G Full  True
        partner/long
```

2. Display the cluster ports on each node that is connected to the existing cluster switches:

```
network device-discovery show -protocol cdp
```

Show example

The following example displays how many cluster interconnect interfaces have been configured in each node for each cluster interconnect switch:

```
cluster1::*> network device-discovery show -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface
Platform
-----
node2      /cdp
           e0a    CL1                        0/2
CN1610
           e0b    CL2                        0/2
CN1610
node1      /cdp
           e0a    CL1                        0/1
CN1610
           e0b    CL2                        0/1
CN1610
```

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

a. Verify that all the cluster ports are up with a healthy status:

```
network port show -ip-space Cluster
```

Show example

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

Node: node1

Ignore

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

Node: node2

Ignore

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
```

- b. Verify that all the cluster interfaces (LIFs) are on their home ports:

```
network interface show -vserver Cluster
```

Show example

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

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

Cluster				
e0a	node1_clus1	up/up	169.254.209.69/16	node1
e0b	true			
e0a	node1_clus2	up/up	169.254.49.125/16	node1
e0b	true			
e0a	node2_clus1	up/up	169.254.47.194/16	node2
e0b	true			
e0a	node2_clus2	up/up	169.254.19.183/16	node2
e0b	true			

4. Verify that the cluster displays information for both cluster switches:

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command: `system switch ethernet show -is-monitoring-enabled-operational true`

```
cluster1::*> system switch ethernet show -is-monitoring-enabled
-operational true
```

Switch	Type	Address	Model
CL1	cluster-network	10.10.1.101	CN1610
Serial Number: 01234567			
Is Monitored: true			
Reason:			
Software Version: 1.3.0.3			
Version Source: ISDP			
CL2	cluster-network	10.10.1.102	CN1610
Serial Number: 01234568			
Is Monitored: true			
Reason:			
Software Version: 1.3.0.3			
Version Source: ISDP			

```
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command: `system cluster-switch show -is-monitoring-enabled-operational true`

```

cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                                     Type                Address             Model
-----
CL1                                         cluster-network    10.10.1.101       CN1610
    Serial Number: 01234567
    Is Monitored: true
    Reason:
    Software Version: 1.3.0.3
    Version Source: ISDP

CL2                                         cluster-network    10.10.1.102       CN1610
    Serial Number: 01234568
    Is Monitored: true
    Reason:
    Software Version: 1.3.0.3
    Version Source: ISDP
cluster1::*>

```

5. Disable auto-revert on the cluster LIFs.

```

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

```

6. On cluster switch CL2, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs:

```

(CL2)# configure
(CL2) (Config)# interface 0/1-0/16
(CL2) (Interface 0/1-0/16)# shutdown
(CL2) (Interface 0/1-0/16)# exit
(CL2) (Config)# exit
(CL2)#

```

7. Verify that the cluster LIFs have failed over to the ports hosted on cluster switch CL1. This might take a few seconds.

```

network interface show -vserver Cluster

```

Show example

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

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

Cluster				
e0a	node1_clus1	up/up	169.254.209.69/16	node1
	true			
e0a	node1_clus2	up/up	169.254.49.125/16	node1
	false			
e0a	node2_clus1	up/up	169.254.47.194/16	node2
	true			
e0a	node2_clus2	up/up	169.254.19.183/16	node2
	false			

8. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon

node1	true	true	false
node2	true	true	false

9. Move all cluster node connection cables from the old CL2 switch to the new cs2 switch.

10. Confirm the health of the network connections moved to cs2:

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

Show example

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

```
Node: node1
```

```
Ignore
```

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

-----	-----	-----	----	----	-----		
e0a	Cluster	Cluster		up	9000	auto/10000	
healthy	false						
e0b	Cluster	Cluster		up	9000	auto/10000	
healthy	false						

```
Node: node2
```

```
Ignore
```

						Speed(Mbps)	Health
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						

All cluster ports that were moved should be up.

11. Check neighbor information on the cluster ports:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
node2         /cdp
              e0a    CL1                       0/2
CN1610
              e0b    cs2                       0/2          BES-
53248
node1         /cdp
              e0a    CL1                       0/1
CN1610
              e0b    cs2                       0/1          BES-
53248
```

12. Confirm the switch port connections are healthy from switch cs2's perspective:

```
cs2# show port all
cs2# show isdp neighbors
```

13. On cluster switch CL1, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs:

```
(CL1)# configure
(CL1) (Config)# interface 0/1-0/16
(CL1) (Interface 0/1-0/16)# shutdown
(CL1) (Interface 0/13-0/16)# exit
(CL1) (Config)# exit
(CL1)#
```

All cluster LIFs failover to the cs2 switch.

14. Verify that the cluster LIFs have failed over to the ports hosted on switch cs2. This might take a few seconds:

```
network interface show -vserver Cluster
```

Show example

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

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

Cluster				
e0b	node1_clus1	up/up	169.254.209.69/16	node1
	false			
e0b	node1_clus2	up/up	169.254.49.125/16	node1
	true			
e0b	node2_clus1	up/up	169.254.47.194/16	node2
	false			
e0b	node2_clus2	up/up	169.254.19.183/16	node2
	true			

15. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon

node1	true	true	false
node2	true	true	false

16. Move the cluster node connection cables from CL1 to the new cs1 switch.

17. Confirm the health of the network connections moved to cs1:

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

Show example

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

```
Node: node1
```

```
Ignore
```

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

```
Node: node2
```

```
Ignore
```

						Speed(Mbps)	Health
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						

All cluster ports that were moved should be up.

18. Check neighbor information on the cluster ports:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port  Device (LLDP: ChassisID)  Interface
Platform
-----
-----
node1         /cdp
              e0a   cs1                       0/1          BES-
53248
              e0b   cs2                       0/1          BES-
53248
node2         /cdp
              e0a   cs1                       0/2          BES-
53248
              e0b   cs2                       0/2          BES-
53248
```

19. Confirm the switch port connections are healthy from switch cs1's perspective:

```
cs1# show port all
cs1# show isdp neighbors
```

20. Verify that the ISL between cs1 and cs2 is still operational:

```
show port-channel
```


Show example

The following example shows that the ISL ports are **up** on switch cs1:

```
(cs1)# show port-channel 1/1
Local Interface..... 1/1
Channel Name..... Cluster-ISL
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Port channel Min-links..... 1
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr      Device/      Port      Port
Ports   Timeout     Speed     Active
-----
0/55    actor/long   100G Full  True
        partner/long
0/56    actor/long   100G Full  True
        partner/long
(cs1) #
```

The following example shows that the ISL ports are **up** on switch cs2:

```
(cs2)# show port-channel 1/1
Local Interface..... 1/1
Channel Name..... Cluster-ISL
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Port channel Min-links..... 1
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr      Device/      Port      Port
Ports   Timeout     Speed     Active
-----
0/55    actor/long   100G Full  True
        partner/long
0/56    actor/long   100G Full  True
        partner/long
```

21. Delete the replaced CN1610 switches from the cluster's switch table, if they are not automatically removed:

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command: `system switch ethernet delete -device device-name`

```
cluster::*> system switch ethernet delete -device CL1
cluster::*> system switch ethernet delete -device CL2
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command: `system cluster-switch delete -device device-name`

```
cluster::*> system cluster-switch delete -device CL1
cluster::*> system cluster-switch delete -device CL2
```

Step 3: Verify the configuration

1. Enable auto-revert on the cluster LIFs.

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

2. Verify that the cluster LIFs have reverted to their home ports (this might take a minute):

```
network interface show -vserver Cluster
```

If the cluster LIFs have not reverted to their home port, manually revert them:

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

3. Verify that the cluster is healthy:

```
cluster show
```

4. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

Packet	Source	Destination	
Node	Date	LIF	LIF
Loss			
node1	3/5/2022 19:21:18 -06:00	node1_clus2	node2_clus1
node1	3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2	3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node2	3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node node2
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
    Local 169.254.19.183 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)

```

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

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

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

What's next?

[Configure switch health monitoring.](#)

Migrate to a switched NetApp cluster environment

If you have an existing two-node *switchless* cluster environment, you can migrate to a two-node *switched* cluster environment using Broadcom-supported BES-53248 cluster switches, which enables you to scale beyond two nodes in the cluster.

The migration process works for all cluster node ports using optical or Twinax ports, but it is not supported on this switch if nodes are using onboard 10GBASE-T RJ45 ports for the cluster network ports.

Review requirements

Review the following requirements for the cluster environment.

- Be aware that most systems require two dedicated cluster-network ports on each controller.
- Make sure that the BES-53248 cluster switch is set up as described in [Replace requirements](#) before starting this migration process.
- For the two-node switchless configuration, ensure that:
 - The two-node switchless configuration is properly set up and functioning.
 - The nodes are running ONTAP 9.5P8 and later. Support for 40/100 GbE cluster ports starts with EFOS firmware version 3.4.4.6 and later.
 - All cluster ports are in the **up** state.
 - All cluster logical interfaces (LIFs) are in the **up** state and on their home ports.
- For the Broadcom-supported BES-53248 cluster switch configuration, ensure that:
 - The BES-53248 cluster switch is fully functional on both switches.
 - Both switches have management network connectivity.
 - There is console access to the cluster switches.
 - BES-53248 node-to-node switch and switch-to-switch connections are using Twinax or fiber cables.

The [NetApp Hardware Universe](#) contains information about ONTAP compatibility, supported EFOS firmware, and cabling to BES-53248 switches.

- Inter-Switch Link (ISL) cables are connected to ports 0/55 and 0/56 on both BES-53248 switches.
- Initial customization of both the BES-53248 switches is complete, so that:
 - BES-53248 switches are running the latest version of software.
 - BES-53248 switches have optional port licenses installed, if purchased.
 - Reference Configuration Files (RCFs) are applied to the switches.
- Any site customization (SMTP, SNMP, and SSH) are configured on the new switches.

Port group speed constraints

- The 48 10/25GbE (SFP28/SFP+) ports are combined into 12 x 4-port groups as follows: Ports 1-4, 5-8, 9-12, 13-16, 17-20, 21-24, 25-28, 29-32, 33-36, 37-40, 41-44, and 45-48.
- The SFP28/SFP+ port speed must be the same (10GbE or 25GbE) across all ports in the 4-port group.
- If speeds in a 4-port group are different, the switch ports will not operate correctly.

Migrate to the cluster environment

About the examples

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

- The names of the BES-53248 switches are `cs1` and `cs2`.
- The names of the cluster SVMs are `node1` and `node2`.
- The names of the LIFs are `node1_clus1` and `node1_clus2` on node 1, and `node2_clus1` and `node2_clus2` on node 2 respectively.

- The `cluster1::*>` prompt indicates the name of the cluster.
- The cluster ports used in this procedure are `e0a` and `e0b`.

The [NetApp Hardware Universe](#) contains the latest information about the actual cluster ports for your platforms.

Step 1: Prepare for migration

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
```

where `x` is the duration of the maintenance window in hours.



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

The following command suppresses automatic case creation for two hours:

```
cluster1::*> system node autosupport invoke -node \* -type all -message MAINT=2h
```

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

```
set -privilege advanced
```

The advanced prompt (`*>`) appears.

Step 2: Configure ports and cabling

1. Disable all activated node-facing ports (not ISL ports) on both the new cluster switches `cs1` and `cs2`.



You must not disable the ISL ports.

The following example shows that node-facing ports 1 through 16 are disabled on switch `cs1`:

```
(cs1)# configure
(cs1) (Config)# interface 0/1-0/16
(cs1) (Interface 0/1-0/16)# shutdown
(cs1) (Interface 0/1-0/16)# exit
(cs1) (Config)# exit
```

2. Verify that the ISL and the physical ports on the ISL between the two BES-53248 switches `cs1` and `cs2` are up:

```
show port-channel
```

Show example

The following example shows that the ISL ports are up on switch cs1:

```
(cs1)# show port-channel 1/1
Local Interface..... 1/1
Channel Name..... Cluster-ISL
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Port channel Min-links..... 1
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr      Device/      Port      Port
Ports   Timeout     Speed     Active
-----  -
0/55    actor/long   100G Full  True
        partner/long
0/56    actor/long   100G Full  True
        partner/long
(cs1) #
```

The following example shows that the ISL ports are up on switch cs2:

```
(cs2)# show port-channel 1/1
Local Interface..... 1/1
Channel Name..... Cluster-ISL
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Port channel Min-links..... 1
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr      Device/      Port      Port
Ports   Timeout     Speed     Active
-----  -
0/55    actor/long   100G Full  True
        partner/long
0/56    actor/long   100G Full  True
        partner/long
```

3. Display the list of neighboring devices:

```
show isdp neighbors
```

This command provides information about the devices that are connected to the system.

Show example

The following example lists the neighboring devices on switch cs1:

```
(cs1)# show isdp neighbors
```

```
Capability Codes: R - Router, T - Trans Bridge, B - Source Route  
Bridge,
```

```
                S - Switch, H - Host, I - IGMP, r - Repeater
```

Device ID	Intf	Holdtime	Capability	Platform	Port ID
cs2	0/55	176	R	BES-53248	0/55
cs2	0/56	176	R	BES-53248	0/56

The following example lists the neighboring devices on switch cs2:

```
(cs2)# show isdp neighbors
```

```
Capability Codes: R - Router, T - Trans Bridge, B - Source Route  
Bridge,
```

```
                S - Switch, H - Host, I - IGMP, r - Repeater
```

Device ID	Intf	Holdtime	Capability	Platform	Port ID
cs2	0/55	176	R	BES-53248	0/55
cs2	0/56	176	R	BES-53248	0/56

4. Verify that all cluster ports are up:

```
network port show -ipspace Cluster
```


Show example

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

```
Node: node1
```

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

```
Node: node2
```

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

5. Verify that all cluster LIFs are up and operational:

```
network interface show -vserver Cluster
```

Show example

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

Current Is Vserver Port Home	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node
Cluster	node1_clus1	up/up	169.254.209.69/16	node1
e0a	true			
	node1_clus2	up/up	169.254.49.125/16	node1
e0b	true			
	node2_clus1	up/up	169.254.47.194/16	node2
e0a	true			
	node2_clus2	up/up	169.254.19.183/16	node2
e0b	true			

6. Disable auto-revert on the cluster LIFs.

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

7. Disconnect the cable from cluster port e0a on node1, and then connect e0a to port 1 on cluster switch cs1, using the appropriate cabling supported by the BES-53248 switches.

The [NetApp Hardware Universe](#) contains more information about cabling.

8. Disconnect the cable from cluster port e0a on node2, and then connect e0a to port 2 on cluster switch cs1, using the appropriate cabling supported by the BES-53248 switches.
9. Enable all node-facing ports on cluster switch cs1.

The following example shows that ports 1 through 16 are enabled on switch cs1:

```
(cs1)# configure
(cs1) (Config)# interface 0/1-0/16
(cs1) (Interface 0/1-0/16)# no shutdown
(cs1) (Interface 0/1-0/16)# exit
(cs1) (Config)# exit
```

10. Verify that all cluster ports are up:

```
network port show -ipspace Cluster
```

Show example

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

```
Node: node1
```

```
Ignore
```

```
Speed(Mbps) Health
```

```
Health
```

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

```
-----
```

```
e0a      Cluster      Cluster      up  9000  auto/10000
```

```
healthy  false
```

```
e0b      Cluster      Cluster      up  9000  auto/10000
```

```
healthy  false
```

```
Node: node2
```

```
Ignore
```

```
Speed(Mbps) Health
```

```
Health
```

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

```
-----
```

```
e0a      Cluster      Cluster      up  9000  auto/10000
```

```
healthy  false
```

```
e0b      Cluster      Cluster      up  9000  auto/10000
```

```
healthy  false
```

11. Verify that all cluster LIFs are up and operational:

```
network interface show -vserver Cluster
```

Show example

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

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	-----
Cluster					
false	node1_clus1	up/up	169.254.209.69/16	node1	e0a
true	node1_clus2	up/up	169.254.49.125/16	node1	e0b
false	node2_clus1	up/up	169.254.47.194/16	node2	e0a
true	node2_clus2	up/up	169.254.19.183/16	node2	e0b

12. Display information about the status of the nodes in the cluster:

```
cluster show
```

Show example

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

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
-----	-----	-----	-----
node1	true	true	false
node2	true	true	false

13. Disconnect the cable from cluster port e0b on node1, and then connect e0b to port 1 on cluster switch cs2, using the appropriate cabling supported by the BES-53248 switches.
14. Disconnect the cable from cluster port e0b on node2, and then connect e0b to port 2 on cluster switch cs2, using the appropriate cabling supported by the BES-53248 switches.
15. Enable all node-facing ports on cluster switch cs2.

The following example shows that ports 1 through 16 are enabled on switch cs2:

```
(cs2)# configure
(cs2) (Config)# interface 0/1-0/16
(cs2) (Interface 0/1-0/16)# no shutdown
(cs2) (Interface 0/1-0/16)# exit
(cs2) (Config)# exit
```

16. Verify that all cluster ports are up:

```
network port show -ipspace Cluster
```

Show example

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

```
Node: node1
```

```
Ignore
```

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

```
Node: node2
```

```
Ignore
```

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

Step 3: Verify the configuration

1. Enable auto-revert on the cluster LIFs.

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

2. Verify that the cluster LIFs have reverted to their home ports (this might take a minute):

```
network interface show -vserver Cluster
```

If the cluster LIFs have not reverted to their home port, manually revert them:

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

3. Verify that all interfaces display true for Is Home:

```
network interface show -vserver Cluster
```



This might take several minutes to complete.

Show example

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

Current Is Home	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Port
Cluster	node1_clus1	up/up	169.254.209.69/16	node1	e0a
true	node1_clus2	up/up	169.254.49.125/16	node1	e0b
true	node2_clus1	up/up	169.254.47.194/16	node2	e0a
true	node2_clus2	up/up	169.254.19.183/16	node2	e0b
true					

4. Verify that both nodes each have one connection to each switch:

```
show isdp neighbors
```

Show example

The following example shows the appropriate results for both switches:

```
(cs1)# show isdp neighbors
```

```
Capability Codes: R - Router, T - Trans Bridge, B - Source Route  
Bridge,
```

```
                S - Switch, H - Host, I - IGMP, r - Repeater
```

```
Device ID      Intf          Holdtime  Capability  Platform  -- Port  
ID
```

```
-----  
-----  
node1          0/1           175       H           FAS2750   e0a  
node2          0/2           157       H           FAS2750   e0a  
cs2            0/55          178       R           BES-53248 0/55  
cs2            0/56          178       R           BES-53248 0/56
```

```
(cs2)# show isdp neighbors
```

```
Capability Codes: R - Router, T - Trans Bridge, B - Source Route  
Bridge,
```

```
                S - Switch, H - Host, I - IGMP, r - Repeater
```

```
Device ID      Intf          Holdtime  Capability  Platform  Port  
ID
```

```
-----  
-----  
node1          0/1           137       H           FAS2750   e0b  
node2          0/2           179       H           FAS2750   e0b  
cs1            0/55          175       R           BES-53248 0/55  
cs1            0/56          175       R           BES-53248 0/56
```

5. Display information about the discovered network devices in your cluster:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port  Device (LLDP: ChassisID)  Interface
Platform
-----
node2         /cdp
              e0a   cs1                       0/2          BES-
53248
              e0b   cs2                       0/2          BES-
53248
node1         /cdp
              e0a   cs1                       0/1          BES-
53248
              e0b   cs2                       0/1          BES-
53248
```

6. Verify that the settings are disabled:

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



It might take several minutes for the command to complete. Wait for the '3 minute lifetime to expire' announcement.

The false output in the following example shows that the configuration settings are disabled:

```
cluster1::*> network options switchless-cluster show
Enable Switchless Cluster: false
```

7. Verify the status of the node members in the cluster:

```
cluster show
```


Show example

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

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
node1	true	true	false
node2	true	true	false

8. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2_clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node node2
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
    Local 169.254.19.183 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)

```

9. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

Show example

```

cluster1::*> system node autosupport invoke -node \* -type all
-messsage MAINT=END

```

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

What's next?

[Configure switch health monitoring.](#)

Replace switches

Replacement requirements

Before replacing the switch, make sure the following conditions are met in the current environment and on the replacement switch.

Existing cluster and network infrastructure

Make sure that:

- The existing cluster is verified as completely functional, with at least one fully connected cluster switch.
- All cluster ports are **up**.
- All cluster logical interfaces (LIFs) are administratively and operationally **up** and on their home ports.
- The ONTAP `cluster ping-cluster -node node1` command must indicate that the settings, basic connectivity and larger than PMTU communication, are successful on all paths.

BES-53248 replacement cluster switch

Make sure that:

- Management network connectivity on the replacement switch is functional.
- Console access to the replacement switch is in place.
- The node connections are ports 0/1 through 0/16 with default licensing.
- All Inter-Switch Link (ISL) ports are disabled on ports 0/55 and 0/56.
- The desired reference configuration file (RCF) and EFOS operating system switch image are loaded onto the switch.
- Initial customization of the switch is complete, as detailed in [Configure the BES-53248 cluster switch](#).

Any previous site customizations, such as STP, SNMP, and SSH, are copied to the new switch.

Enable console logging

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

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

For more information

- [NetApp Support Site](#)
- [NetApp Hardware Universe](#)

Replace a Broadcom-supported BES-53248 cluster switch

Follow these steps to replace a defective Broadcom-supported BES-53248 cluster switch in a cluster network. This is a nondisruptive procedure (NDU).

About the examples

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

- The names of the existing BES-53248 switches are `cs1` and `cs2`.
- The name of the new BES-53248 switch is `newcs2`.
- The node names are `node1` and `node2`.
- The cluster ports on each node are named `e0a` and `e0b`.
- The cluster LIF names are `node1_clus1` and `node1_clus2` for `node1`, and `node2_clus1` and `node2_clus2` for `node2`.
- The prompt for changes to all cluster nodes is `cluster1::>`

About the topology

This procedure is based on the following cluster network topology:

Show example topology

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

```
Node: node1
```

```
Ignore
```

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

```
Node: node2
```

```
Ignore
```

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

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

	Logical	Status	Network	Current		
Current Is	Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home						
	Cluster	node1_clus1	up/up	169.254.209.69/16	node1	e0a
true		node1_clus2	up/up	169.254.49.125/16	node1	e0b
true						

```
node2_clus1 up/up 169.254.47.194/16 node2 e0a
true
node2_clus2 up/up 169.254.19.183/16 node2 e0b
true
```

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

Node/	Local	Discovered		
Protocol	Port	Device (LLDP: ChassisID)	Interface	Platform

node2	/cdp			
	e0a	cs1	0/2	BES-
53248				
	e0b	cs2	0/2	BES-
53248				
node1	/cdp			
	e0a	cs1	0/1	BES-
53248				
	e0b	cs2	0/1	BES-
53248				

```
(cs1)# show isdp neighbors
```

```
Capability Codes: R - Router, T - Trans Bridge, B - Source Route  
Bridge,
```

```
S - Switch, H - Host, I - IGMP, r - Repeater
```

Device ID Port ID	Intf	Holdtime	Capability	Platform
node1 e0a	0/1	175	H	FAS2750
node2 e0a	0/2	152	H	FAS2750
cs2 0/55	0/55	179	R	BES-53248
cs2 0/56	0/56	179	R	BES-53248

```
(cs2)# show isdp neighbors
```

```
Capability Codes: R - Router, T - Trans Bridge, B - Source Route  
Bridge,
```

```
S - Switch, H - Host, I - IGMP, r - Repeater
```

Device ID Port ID	Intf	Holdtime	Capability	Platform
node1 e0b	0/1	129	H	FAS2750
node2 e0b	0/2	165	H	FAS2750
cs1 0/55	0/55	179	R	BES-53248
cs1 0/56	0/56	179	R	BES-53248

Steps

1. Review the [Replacement requirements](#).
2. 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
```

where *x* is the duration of the maintenance window in hours.



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

3. Install the appropriate Reference Configuration File (RCF) and image on the switch, newcs2, and make any necessary site preparations.

If necessary, verify, download, and install the appropriate versions of the RCF and EFOS software for the new switch. If you have verified that the new switch is correctly set up and does not need updates to the RCF and EFOS software, continue to step 2.

- a. You can download the applicable Broadcom EFOS software for your cluster switches from the [Broadcom Ethernet Switch Support](#) site. Follow the steps on the Download page to download the EFOS file for the version of ONTAP software you are installing.
 - b. The appropriate RCF is available from the [Broadcom Cluster Switches](#) page. Follow the steps on the Download page to download the correct RCF for the version of ONTAP software you are installing.
4. On the new switch, log in as `admin` and shut down all of the ports that will be connected to the node cluster interfaces (ports 1 to 16).



If you purchased additional licenses for additional ports, shut down these ports too.

If the switch that you are replacing is not functional and is powered down, the LIFs on the cluster nodes should have already failed over to the other cluster port for each node.



No password is required to enter `enable` mode.

Show example

```
User: admin
Password:
(newcs2) > enable
(newcs2) # config
(newcs2) (config) # interface 0/1-0/16
(newcs2) (interface 0/1-0/16) # shutdown
(newcs2) (interface 0/1-0/16) # exit
(newcs2) (config) # exit
(newcs2) #
```

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

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

Show example topology

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

Logical Vserver	Interface	Auto-revert
Cluster	node1_clus1	true
Cluster	node1_clus2	true
Cluster	node2_clus1	true
Cluster	node2_clus2	true

6. Shut down the ISL ports 0/55 and 0/56 on the BES-53248 switch cs1:

Show example topology

```
(cs1)# config  
(cs1)(config)# interface 0/55-0/56  
(cs1)(interface 0/55-0/56)# shutdown
```

7. Remove all cables from the BES-53248 cs2 switch, and then connect them to the same ports on the BES-53248 newcs2 switch.
8. Bring up the ISLs ports 0/55 and 0/56 between the cs1 and newcs2 switches, and then verify the port channel operation status.

The Link State for port-channel 1/1 should be **up** and all member ports should be True under the Port Active heading.

Show example

This example enables ISL ports 0/55 and 0/56 and displays the Link State for port-channel 1/1 on switch cs1:

```
(cs1)# config
(cs1)(config)# interface 0/55-0/56
(cs1)(interface 0/55-0/56)# no shutdown
(cs1)(interface 0/55-0/56)# exit
(cs1)# show port-channel 1/1

Local Interface..... 1/1
Channel Name..... Cluster-ISL
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Port-channel Min-links..... 1
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr      Device/      Port      Port
Ports   Timeout     Speed     Active
-----
0/55    actor/long   100G Full  True
        partner/long
0/56    actor/long   100G Full  True
        partner/long
```

9. On the new switch newcs2, re-enable all of the ports that are connected to the node cluster interfaces (ports 1 to 16).



If you purchased additional licenses for additional ports, shut down these ports too.

Show example

```
User:admin
Password:
(newcs2)> enable
(newcs2)# config
(newcs2)(config)# interface 0/1-0/16
(newcs2)(interface 0/1-0/16)# no shutdown
(newcs2)(interface 0/1-0/16)# exit
(newcs2)(config)# exit
```

10. Verify that port e0b is **up**:

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

Show example

The output should be similar to the following:

```
cluster1::> network port show -ipSpace Cluster

Node: node1

Ignore

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

Node: node2

Ignore

Health      Health      Speed (Mbps)
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/auto  -
false
```

11. On the same node as you used in the previous step, wait for the cluster LIF node1_clus2 on node1 to auto-revert.

Show example

In this example, LIF node1_clus2 on node1 is successfully reverted if `Is Home` is `true` and the port is `e0b`.

The following command displays information about the LIFs on both nodes. Bringing up the first node is successful if `Is Home` is `true` for both cluster interfaces and they show the correct port assignments, in this example `e0a` and `e0b` on node1.

```
cluster::> network interface show -vserver Cluster
```

Current Is Vserver Port	Logical Interface Home	Status Admin/Oper	Network Address/Mask	Current Node
Cluster	node1_clus1	up/up	169.254.209.69/16	node1
e0a	true			
	node1_clus2	up/up	169.254.49.125/16	node1
e0b	true			
	node2_clus1	up/up	169.254.47.194/16	node2
e0a	true			
	node2_clus2	up/up	169.254.19.183/16	node2
e0a	false			

12. Display information about the nodes in a cluster:

```
cluster show
```

Show example

This example shows that the node health for `node1` and `node2` in this cluster is `true`:

```
cluster1::> cluster show
```

Node	Health	Eligibility	Epsilon
node1	true	true	true
node2	true	true	true

13. Confirm the following cluster network configuration:

```
network port show
```

network interface show

Show example

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

```
Node: node1
```

```
Ignore
```

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

```
-----  
-----
```

	e0a	Cluster	Cluster	up	9000	auto/10000	
healthy		false					
	e0b	Cluster	Cluster	up	9000	auto/10000	
healthy		false					

```
Node: node2
```

```
Ignore
```

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

```
-----  
-----
```

	e0a	Cluster	Cluster	up	9000	auto/10000	
healthy		false					
	e0b	Cluster	Cluster	up	9000	auto/10000	
healthy		false					

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

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

```
-----  
-----
```

Cluster					
		node1_clus1	up/up	169.254.209.69/16	node1
e0a	true				
		node1_clus2	up/up	169.254.49.125/16	node1
e0b	true				
		node2_clus1	up/up	169.254.47.194/16	node2

```
e0a      true
          node2_clus2  up/up    169.254.19.183/16  node2
e0b      true
4 entries were displayed.
```

14. Verify that the cluster network is healthy:

```
show isdp neighbors
```

Show example

```
(cs1)# show isdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route
Bridge,
S - Switch, H - Host, I - IGMP, r - Repeater
Device ID      Intf      Holdtime    Capability    Platform      Port ID
-----
node1          0/1       175         H             FAS2750       e0a
node2          0/2       152         H             FAS2750       e0a
newcs2         0/55      179         R             BES-53248     0/55
newcs2         0/56      179         R             BES-53248     0/56

(newcs2)# show isdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route
Bridge,
S - Switch, H - Host, I - IGMP, r - Repeater

Device ID      Intf      Holdtime    Capability    Platform      Port ID
-----
node1          0/1       129         H             FAS2750       e0b
node2          0/2       165         H             FAS2750       e0b
cs1            0/55      179         R             BES-53248     0/55
cs1            0/56      179         R             BES-53248     0/56
```

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

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

What's next?

[Configure switch health monitoring.](#)

Replace Broadcom BES-53248 cluster switches with switchless connections

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

are directly connected for ONTAP 9.3 and later.

Review requirements

Guidelines

Review the following guidelines:

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


What you'll need

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

Migrate the switches

About this task

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



About the examples

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

Step 1: Prepare for migration

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

```
set -privilege advanced
```

The advanced prompt `*>` appears.

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

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

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

Show example

The following example output shows if the option is enabled.

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

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

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

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

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

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

Show example

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


Step 2: Configure ports and cabling

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

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

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

using different cluster ports because they vary by system.



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

Show example

```

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

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

Node: node2

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

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

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

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

Show example

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

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

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

4. Disable auto-revert for the cluster LIFs:

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

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

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

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

Show example

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

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

6. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

7. Verify that the cluster is healthy:

```
cluster ring show
```

All units must be either master or secondary.


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



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


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

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



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

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



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

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

The following example shows that the switchless cluster is enabled:

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

10. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```



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


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



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


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

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



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

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



Step 3: Verify the configuration

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

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

Show example

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

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

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

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

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

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

Show example

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

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

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

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

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

```
cluster show
```

Show example

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

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

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

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

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

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

7. Change the privilege level back to admin:

```
set -privilege admin
```

Cisco Nexus 9336C-FX2 or 9336C-FX2-T

Overview

Overview of installation and configuration for Cisco Nexus 9336C-FX2 cluster switches

The Cisco Nexus 9336C-FX2 cluster switch is part of the Cisco Nexus 9000 platform and can be installed in a NetApp system cabinet. Cluster switches allow you to build ONTAP clusters with more than two nodes.

Initial configuration overview

To initially configure a Cisco Nexus 9336C-FX2 switch on systems running ONTAP, follow these steps:

1. [Complete the Cisco Nexus 9336C-FX2 or 9336C-FX2-T cabling worksheet.](#)

The sample cabling worksheet provides examples of recommended port assignments from the switches to the controllers. The blank worksheet provides a template that you can use in setting up your cluster.

2. [Install the switch.](#)

Set up the switch hardware.

3. [Configure the 9336C-FX2 cluster switch.](#)

Set up the Cisco Nexus 9336C-FX2 switch.

4. [Install a Cisco Nexus 9336C-FX2 switch in a NetApp cabinet.](#)

Depending on your configuration, you can install the Cisco Nexus 9336C-FX2 switch and pass-through panel in a NetApp cabinet with the standard brackets that are included with the switch.

5. [Prepare to install the NX-OS software and RCF.](#)

Follow preliminary procedures in preparation for installing the Cisco NX-OS software and reference configuration files (RCFs).

6. [Install or upgrade the NX-OS software.](#)

Install or upgrade the NX-OS software on the Nexus 9336C-FX2 cluster switch.

7. [Install or upgrade the RCF.](#)

Install or upgrade the RCF after setting up the Nexus 9336C-FX2 switch for the first time. You can also use this procedure to upgrade your RCF version.

8. [Verify SSH configuration on the Nexus 9336C-FX2 cluster switches.](#)

If you use the Ethernet Switch Health Monitor (CSHM) and log collection features, verify SSH is enabled on the switches.

Additional information

Before you begin installation or maintenance, be sure to review the following:

- [Configuration requirements](#)
- [Components and part numbers](#)
- [Required documentation](#)
- [Smart Call Home requirements](#)

Configuration requirements for Cisco Nexus 9336C-FX2 cluster switches

For Cisco Nexus 9336C-FX2 switch installation and maintenance, be sure to review configuration and network requirements.

ONTAP support

From ONTAP 9.9.1, you can use Cisco Nexus 9336C-FX2 switches to combine storage and cluster functionality into a shared switch configuration.

If you want to build ONTAP clusters with more than two nodes, you need two supported network switches.



Ethernet switch health monitor does not support ONTAP 9.13.1P8 and earlier and 9.14.1P3 and earlier or NX-OS version 10.3(4a)(M).

Configuration requirements

Make sure that:

- You have the appropriate number and type of cables and cable connectors for your switches. See the [Hardware Universe](#).
- Depending on the type of switch you are initially configuring, you need to connect to the switch console port with the included console cable.

Network requirements

You need the following network information for all switch configurations.

- IP subnet for management network traffic
- Host names and IP addresses for each of the storage system controllers and all applicable switches
- Most storage system controllers are managed through the e0M interface by connecting to the Ethernet service port (wrench icon). On AFF A800 and AFF A700s systems, the e0M interface uses a dedicated Ethernet port.
- Refer to the [Hardware Universe](#) for the latest information.

For more information about the initial configuration of your switch, see the following guide: [Cisco Nexus 9336C-FX2 Installation and Upgrade Guide](#).

Components and part numbers for Cisco Nexus 9336C-FX2 and 9336C-FX2-T cluster switches

For Cisco Nexus 9336C-FX2 and 9336C-FX2-T switch installation and maintenance, ensure to review the list of components and part numbers.

Part number details

The following table lists the part number and description for the 9336C-FX2 and 9336C-FX2-T switches, fans, and power supplies:

Part number	Description
X190200-CS-PE	Cluster Switch, N9336C 36Pt PTSX 10/25/40/100G
X190200-CS-PI	Cluster Switch, N9336C 36Pt PSIN 10/25/40/100G
X190212-CS-PE	Cluster Switch, N9336C 12Pt (9336C-FX2-T) PTSX 10/25/40/100G

Part number	Description
X190212-CS-PI	Cluster Switch, N9336C 12Pt (9336C-FX2-T) PSIN 10/25/40/100G
SW-N9K-FX2-24P-UPG	SW, Cisco 9336CFX2 24-Port POD License
X190210-FE-PE	N9K-9336C, FTE, PTSX, 36PT 10/25/40/100GQSFP28
X190210-FE-PI	N9K-9336C, FTE, PSIN, 36PT 10/25/40/100GQSFP28
X190002	Accessory Kit X190001/X190003
X-NXA-PAC-1100W-PE2	N9K-9336C AC 1100W PSU - Port side exhaust airflow
X-NXA-PAC-1100W-PI2	N9K-9336C AC 1100W PSU - Port side Intake airflow
X-NXA-FAN-65CFM-PE	N9K-9336C 65CFM, Port side exhaust airflow
X-NXA-FAN-65CFM-PI	N9K-9336C 65CFM, Port side intake airflow

Cisco Smart licenses

In order to activate more than 12 ports on your Cisco Nexus 9336C-FX-T cluster switch, you must purchase a Cisco Smart license. Cisco Smart licenses are managed through Cisco Smart accounts.

1. Create a new Smart account, if required. See [Create a new Smart account](#) for details.
2. Request access to an existing Smart account. See [Request access to an existing Smart account](#) for details.



Once you have purchased your Smart license, install the appropriate RCF to enable and configure all 36 available ports for use.

Documentation requirements for Cisco Nexus 9336C-FX2 switches

For Cisco Nexus 9336C-FX2 switch installation and maintenance, be sure to review specific switch and controller documentation to set up your Cisco 9336-FX2 switches and ONTAP cluster.

Switch documentation

To set up the Cisco Nexus 9336C-FX2 switches, you need the following documentation from the [Cisco Nexus 9000 Series Switches Support](#) page:

Document title	Description
<i>Nexus 9000 Series Hardware Installation Guide</i>	Provides detailed information about site requirements, switch hardware details, and installation options.

Document title	Description
<i>Cisco Nexus 9000 Series Switch Software Configuration Guides</i> (choose the guide for the NX-OS release installed on your switches)	Provides initial switch configuration information that you need before you can configure the switch for ONTAP operation.
<i>Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide</i> (choose the guide for the NX-OS release installed on your switches)	Provides information on how to downgrade the switch to ONTAP supported switch software, if necessary.
<i>Cisco Nexus 9000 Series NX-OS Command Reference Master Index</i>	Provides links to the various command references provided by Cisco.
<i>Cisco Nexus 9000 MIBs Reference</i>	Describes the Management Information Base (MIB) files for the Nexus 9000 switches.
<i>Nexus 9000 Series NX-OS System Message Reference</i>	Describes the system messages for Cisco Nexus 9000 series switches, those that are informational, and others that might help diagnose problems with links, internal hardware, or the system software.
<i>Cisco Nexus 9000 Series NX-OS Release Notes</i> (choose the notes for the NX-OS release installed on your switches)	Describes the features, bugs, and limitations for the Cisco Nexus 9000 Series.
Regulatory Compliance and Safety Information for Cisco Nexus 9000 Series	Provides international agency compliance, safety, and statutory information for the Nexus 9000 series switches.

ONTAP systems documentation

To set up an ONTAP system, you need the following documents for your version of the operating system from the [ONTAP 9 Documentation Center](#).

Name	Description
Controller-specific <i>Installation and Setup Instructions</i>	Describes how to install NetApp hardware.
ONTAP documentation	Provides detailed information about all aspects of the ONTAP releases.
Hardware Universe	Provides NetApp hardware configuration and compatibility information.

Rail kit and cabinet documentation

To install a Cisco 9336-FX2 switch in a NetApp cabinet, see the following hardware documentation.

Name	Description
42U System Cabinet, Deep Guide	Describes the FRUs associated with the 42U system cabinet, and provides maintenance and FRU replacement instructions.
Install a Cisco 9336-FX2 switch in a NetApp Cabinet	Describes how to install a Cisco Nexus 9336C-FX2 switch in a four-post NetApp cabinet.

Smart Call Home requirements

To use the Smart Call Home feature, review the following guidelines.

Smart Call Home monitors the hardware and software components on your network. When a critical system configuration occurs, it generates an email-based notification and raises an alert to all the recipients that are configured in your destination profile. To use Smart Call Home, you must configure a cluster network switch to communicate using email with the Smart Call Home system. In addition, you can optionally set up your cluster network switch to take advantage of Cisco's embedded Smart Call Home support feature.

Before you can use Smart Call Home, be aware of the following considerations:

- An email server must be in place.
- The switch must have IP connectivity to the email server.
- The contact name (SNMP server contact), phone number, and street address information must be configured. This is required to determine the origin of messages received.
- A CCO ID must be associated with an appropriate Cisco SMARTnet Service contract for your company.
- Cisco SMARTnet Service must be in place for the device to be registered.

The [Cisco support site](#) contains information about the commands to configure Smart Call Home.

Install hardware

Complete the Cisco Nexus 9336C-FX2 or 9336C-FX2-T cabling worksheet

If you want to document the supported platforms, download a PDF of this page and complete the cabling worksheet.

The sample cabling worksheet provides examples of recommended port assignments from the switches to the controllers. The blank worksheet provides a template that you can use in setting up your cluster.

- [9336C-FX2 sample cabling worksheet](#)
- [9336C-FX2 blank cabling worksheet](#)
- [9336C-FX2-T sample cabling worksheet \(12-port\)](#)
- [9336C-FX2-T blank cabling worksheet \(12-port\)](#)

9336C-FX2 sample cabling worksheet

The sample port definition on each pair of switches is as follows:

Cluster switch A		Cluster switch B	
Switch port	Node and port usage	Switch port	Node and port usage
1	4x10GbE node 1	1	4x10GbE node 1
2	4x10GbE node 2	2	4x10GbE node 2
3	4x10GbE node 3	3	4x10GbE node 3
4	4x25GbE node 4	4	4x25GbE node 4
5	4x25GbE node 5	5	4x25GbE node 5
6	4x25GbE node 6	6	4x25GbE node 6
7	40/100GbE node 7	7	40/100GbE node 7
8	40/100GbE node 8	8	40/100GbE node 8
9	40/100GbE node 9	9	40/100GbE node 9
10	40/100GbE node 10	10	40/100GbE node 10
11	40/100GbE node 11	11	40/100GbE node 11
12	40/100GbE node 12	12	40/100GbE node 12
13	40/100GbE node 13	13	40/100GbE node 13
14	40/100GbE node 14	14	40/100GbE node 14
15	40/100GbE node 15	15	40/100GbE node 15
16	40/100GbE node 16	16	40/100GbE node 16
17	40/100GbE node 17	17	40/100GbE node 17
18	40/100GbE node 18	18	40/100GbE node 18
19	40/100GbE node 19	19	40/100GbE node 19
20	40/100GbE node 20	20	40/100GbE node 20
21	40/100GbE node 21	21	40/100GbE node 21

Cluster switch A		Cluster switch B	
22	40/100GbE node 22	22	40/100GbE node 22
23	40/100GbE node 23	23	40/100GbE node 23
24	40/100GbE node 24	24	40/100GbE node 24
25 through 34	Reserved	25 through 34	Reserved
35	100GbE ISL to switch B port 35	35	100GbE ISL to switch A port 35
36	100GbE ISL to switch B port 36	36	100GbE ISL to switch A port 36

9336C-FX2 blank cabling worksheet

You can use the blank cabling worksheet to document the platforms that are supported as nodes in a cluster. The *Supported Cluster Connections* section of the [Hardware Universe](#) defines the cluster ports used by the platform.

Cluster switch A		Cluster switch B	
1		1	
2		2	
3		3	
4		4	
5		5	
6		6	
7		7	
8		8	
9		9	
10		10	
11		11	

Cluster switch A		Cluster switch B	
12		12	
13		13	
14		14	
15		15	
16		16	
17		17	
18		18	
19		19	
20		20	
21		21	
22		22	
23		23	
24		24	
25 through 34	Reserved	25 through 34	Reserved
35	100GbE ISL to switch B port 35	35	100GbE ISL to switch A port 35
36	100GbE ISL to switch B port 36	36	100GbE ISL to switch A port 36

9336C-FX2-T sample cabling worksheet (12-port)

The sample port definition on each pair of switches is as follows:

Cluster switch A		Cluster switch B	
Switch port	Node and port usage	Switch port	Node and port usage
1	4x10GbE node 1	1	4x10GbE node 1
2	4x10GbE node 2	2	4x10GbE node 2

Cluster switch A		Cluster switch B	
3	4x10GbE node 3	3	4x10GbE node 3
4	4x25GbE node 4	4	4x25GbE node 4
5	4x25GbE node 5	5	4x25GbE node 5
6	4x25GbE node 6	6	4x25GbE node 6
7	40/100GbE node 7	7	40/100GbE node 7
8	40/100GbE node 8	8	40/100GbE node 8
9	40/100GbE node 9	9	40/100GbE node 9
10	40/100GbE node 10	10	40/100GbE node 10
11 through 34	Requires license	11 through 34	Requires license
35	100GbE ISL to switch B port 35	35	100GbE ISL to switch A port 35
36	100GbE ISL to switch B port 36	36	100GbE ISL to switch A port 36

9336C-FX2-T blank cabling worksheet (12-port)

You can use the blank cabling worksheet to document the platforms that are supported as nodes in a cluster. The *Supported Cluster Connections* section of the [Hardware Universe](#) defines the cluster ports used by the platform.

Cluster switch A		Cluster switch B	
1		1	
2		2	
3		3	
4		4	
5		5	
6		6	

Cluster switch A		Cluster switch B	
7		7	
8		8	
9		9	
10		10	
11 through 34	Requires license	11 through 34	Requires license
35	100GbE ISL to switch B port 35	35	100GbE ISL to switch A port 35
36	100GbE ISL to switch B port 36	36	100GbE ISL to switch A port 36

See the [Hardware Universe](#) for more information on switch ports.

Install the 9336C-FX2 cluster switch

Follow this procedure to set up and configure the Cisco Nexus 9336C-FX2 switch.

What you'll need

- Access to an HTTP, FTP, or TFTP server at the installation site to download the applicable NX-OS and Reference Configuration File (RCF) releases.
- Applicable NX-OS version, downloaded from the [Cisco Software Download](#) page.
- Applicable licenses, network and configuration information, and cables.
- Completed [cabling worksheets](#).
- Applicable NetApp cluster network and management network RCFs downloaded from the NetApp Support Site at mysupport.netapp.com. All Cisco cluster network and management network switches arrive with the standard Cisco factory-default configuration. These switches also have the current version of the NX-OS software but do not have the RCFs loaded.
- [Required switch and ONTAP documentation](#).

Steps

1. Rack the cluster network and management network switches and controllers.

If you are installing the...	Then...
Cisco Nexus 9336C-FX2 in a NetApp system cabinet	See the <i>Installing a Cisco Nexus 9336C-FX2 cluster switch and pass-through panel in a NetApp cabinet</i> guide for instructions to install the switch in a NetApp cabinet.
Equipment in a Telco rack	See the procedures provided in the switch hardware installation guides and the NetApp installation and setup instructions.

2. Cable the cluster network and management network switches to the controllers using the completed cabling worksheets.
3. Power on the cluster network and management network switches and controllers.

What's next?

Go to [Configure the Cisco Nexus 9336C-FX2 switch](#).

Configure the 9336C-FX2 cluster switch

Follow this procedure to configure the Cisco Nexus 9336C-FX2 switch.

What you'll need


- Access to an HTTP, FTP, or TFTP server at the installation site to download the applicable NX-OS and Reference Configuration File (RCF) releases.
- Applicable NX-OS version, downloaded from the [Cisco software download](#) page.
- Applicable licenses, network and configuration information, and cables.
- Completed [cabling worksheets](#).
- Applicable NetApp cluster network and management network RCFs downloaded from the NetApp Support Site at mysupport.netapp.com. All Cisco cluster network and management network switches arrive with the standard Cisco factory-default configuration. These switches also have the current version of the NX-OS software but do not have the RCFs loaded.
- [Required switch and ONTAP documentation](#).


Steps

1. Perform an initial configuration of the cluster network switches.

Provide applicable responses to the following initial setup questions when you first boot the switch. Your site's security policy defines the responses and services to enable.

Prompt	Response
Abort Auto Provisioning and continue with normal setup? (yes/no)	Respond with yes . The default is no.
Do you want to enforce secure password standard? (yes/no)	Respond with yes . The default is yes.
Enter the password for admin.	The default password is "admin"; you must create a new, strong password. A weak password can be rejected.
Would you like to enter the basic configuration dialog? (yes/no)	Respond with yes at the initial configuration of the switch.
Create another login account? (yes/no)	Your answer depends on your site's policies on alternate administrators. The default is no .

Prompt	Response
Configure read-only SNMP community string? (yes/no)	Respond with no . The default is no.
Configure read-write SNMP community string? (yes/no)	Respond with no . The default is no.
Enter the switch name.	Enter the switch name, which is limited to 63 alphanumeric characters.
Continue with Out-of-band (mgmt0) management configuration? (yes/no)	Respond with yes (the default) at that prompt. At the mgmt0 IPv4 address: prompt, enter your IP address: ip_address.
Configure the default-gateway? (yes/no)	Respond with yes . At the IPv4 address of the default-gateway: prompt, enter your default_gateway.
Configure advanced IP options? (yes/no)	Respond with no . The default is no.
Enable the telnet service? (yes/no)	Respond with no . The default is no.
Enabled SSH service? (yes/no)	Respond with yes . The default is yes. <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-left: 20px;">  SSH is recommended when using Ethernet Switch Health Monitor (CSHM) for its log collection features. SSHv2 is also recommended for enhanced security. </div>
Enter the type of SSH key you want to generate (dsa/rsa/rsa1).	The default is rsa .
Enter the number of key bits (1024-2048).	Enter the number of key bits from 1024 to 2048.
Configure the NTP server? (yes/no)	Respond with no . The default is no.
Configure default interface layer (L3/L2)	Respond with L2 . The default is L2.
Configure default switch port interface state (shut/noshut)	Respond with noshut . The default is noshut.
Configure CoPP system profile (strict/moderate/lenient/dense)	Respond with strict . The default is strict.

Prompt	Response
Would you like to edit the configuration? (yes/no)	You should see the new configuration at this point. Review and make any necessary changes to the configuration you just entered. Respond with no at the prompt if you are satisfied with the configuration. Respond with yes if you want to edit your configuration settings.
Use this configuration and save it? (yes/no)	<p>Respond with yes to save the configuration. This automatically updates the kickstart and system images.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  If you do not save the configuration at this stage, none of the changes will be in effect the next time you reboot the switch. </div>

2. Verify the configuration choices you made in the display that appears at the end of the setup, and make sure that you save the configuration.
3. Check the version on the cluster network switches, and if necessary, download the NetApp-supported version of the software to the switches from the [Cisco software download](#) page.

What's next?

Optionally, you can [install a Cisco Nexus 9336C-FX2 switch in a NetApp cabinet](#). Otherwise, go to [Prepare to install NX-OS and RCF](#).

Install a Cisco Nexus 9336C-FX2 switch in a NetApp cabinet

Depending on your configuration, you might need to install the Cisco Nexus 9336C-FX2 switch and pass-through panel in a NetApp cabinet. Standard brackets are included with the switch.

What you'll need

- The pass-through panel kit, which is available from NetApp (part number X8784-R6).

The NetApp pass-through panel kit contains the following hardware:

- One pass-through blanking panel
- Four 10-32 x .75 screws
- Four 10-32 clip nuts
- For each switch, eight 10-32 or 12-24 screws and clip nuts to mount the brackets and slider rails to the front and rear cabinet posts.
- The Cisco standard rail kit to install the switch in a NetApp cabinet.



The jumper cords are not included with the pass-through kit and should be included with your switches. If they were not shipped with the switches, you can order them from NetApp (part number X1558A-R6).


- For initial preparation requirements, kit contents, and safety precautions, see [Cisco Nexus 9000 Series Hardware Installation Guide](#).

Steps

1. Install the pass-through blanking panel in the NetApp cabinet.
 - a. Determine the vertical location of the switches and blanking panel in the cabinet.


In this procedure, the blanking panel is installed in U40.

- b. Install two clip nuts on each side in the appropriate square holes for front cabinet rails.
- c. Center the panel vertically to prevent intrusion into adjacent rack space, and then tighten the screws.
- d. Insert the female connectors of both 48-inch jumper cords from the rear of the panel and through the brush assembly.




(1) Female connector of the jumper cord.

2. Install the rack-mount brackets on the Nexus 9336C-FX2 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.
3. Install the clip nuts in the square hole locations for all four IEA posts.



The two 9336C-FX2 switches are always mounted in the top 2U of the cabinet RU41 and 42.

4. Install the slider rails in the cabinet.

- 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 cabinet posts.

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

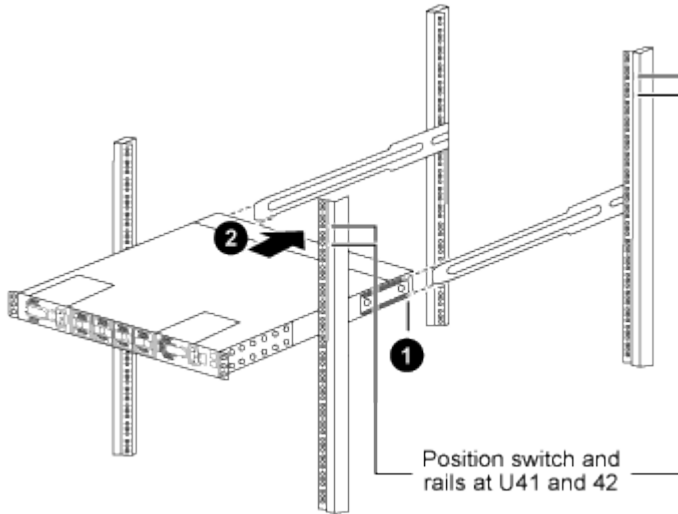
c. Repeat steps 4a and 4b at the RU41 locations on the cabinet.

5. Install the switch in the cabinet.



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) As the chassis is pushed toward the rear posts, align the two rear rack-mount guides with the slider rails.

(2) Gently slide the switch until the front rack-mount brackets are flush with the front posts.

b. Attach the switch to the cabinet.



(1) With one person holding the front of the chassis level, the other person should fully tighten the four rear screws to the cabinet 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.



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.

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

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



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

8. Connect the management port on each 9336C-FX2 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.

What's next?

[Configure the Cisco Nexus 9336C-FX2 switch.](#)

Review cabling and configuration considerations

Before configuring your Cisco 9336C-FX2 switch, review the following considerations.

Support for NVIDIA CX6, CX6-DX, and CX7 Ethernet ports

If connecting a switch port to an ONTAP controller using NVIDIA ConnectX-6 (CX6), ConnectX-6 Dx (CX6-DX), or ConnectX-7 (CX7) NIC ports, you must hard-code the switch port speed.

```
(cs1)(config)# interface Ethernet1/19
For 100GbE speed:
(cs1)(config-if)# speed 100000
For 40GbE speed:
(cs1)(config-if)# speed 40000
(cs1)(config-if)# no negotiate auto
(cs1)(config-if)# exit
(cs1)(config)# exit
Save the changes:
(cs1)# copy running-config startup-config
```

See the [Hardware Universe](#) for more information on switch ports.

25GbE FEC requirements

FAS2820 e0a/e0b ports

FAS2820 e0a and e0b ports require FEC configuration changes to link up with 9336C-FX2 switch ports. For switch ports e0a and e0b, the fec setting is set to `rs-cons16`.


```
(cs1) (config)# interface Ethernet1/8-9
(cs1) (config-if-range)# fec rs-cons16
(cs1) (config-if-range)# exit
(cs1) (config)# exit
Save the changes:
(cs1)# copy running-config startup-config
```

Ports do not link up due to TCAM resources

On the 9336C-FX2 switch, the Ternary Content Addressable Memory (TCAM) resources configured in the configuration utilized by the switch are exhausted.

See the Knowledge Base article [Ports do not link up on Cisco Nexus 9336C-FX2 due to TCAM resources](#) for details on how to resolve this issue.

Configure software

Software install workflow for Cisco Nexus 9336C-FX2 cluster switches

To install and configure the software for a Cisco Nexus 9336C-FX2 switch and to install or upgrade the Reference Configuration File (RCF), follow these steps:

1. [Prepare to install the NX-OS software and RCF.](#)

The Cisco NX-OS software and reference configuration files (RCFs) must be installed on Cisco Nexus 9336C-FX2 cluster switches.

2. [Install or upgrade the NX-OS software.](#)

Download and install or upgrade the NX-OS software on the Cisco Nexus 9336C-FX2 cluster switch.

3. [Install or upgrade the RCF.](#)

Install or upgrade the RCF after setting up the Cisco Nexus 9336C-FX2 switch for the first time. You can also use this procedure to upgrade your RCF version.

4. [Verify SSH configuration on the Nexus 9336C-FX2 cluster switches.](#)

If you use the Ethernet Switch Health Monitor (CSHM) and log collection features, verify SSH is enabled on the switches.

Prepare to install NX-OS software and RCF

Before you install the NX-OS software and the Reference Configuration File (RCF), follow this procedure.

About the examples

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.

About this task

The procedure requires the use of both ONTAP commands and Cisco Nexus 9000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

Steps

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=x h`

where x is the duration of the maintenance window in hours.



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

2. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (`*>`) appears.

3. Display how many cluster interconnect interfaces are configured in each node for each cluster interconnect switch:

```
network device-discovery show -protocol cdp
```

Show example

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

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	
cluster1-02/cdp	e0a	cs1	Eth1/2	N9K-
C9336C	e0b	cs2	Eth1/2	N9K-
C9336C				
cluster1-01/cdp	e0a	cs1	Eth1/1	N9K-
C9336C	e0b	cs2	Eth1/1	N9K-
C9336C				

```
4 entries were displayed.
```

4. Check the administrative or operational status of each cluster interface.
 - a. Display the network port attributes:

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

Show example

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

Node: cluster1-02

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

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
```

Show example

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

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

4 entries were displayed.

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

Packet	Source	Destination
Node	Date	LIF
Loss		
-----	-----	-----
node1		
3/5/2022 19:21:18 -06:00	cluster1-01_clus2	cluster1-02-
clus1 none		
3/5/2022 19:21:20 -06:00	cluster1-01_clus2	cluster1-
02_clus2 none		
node2		
3/5/2022 19:21:18 -06:00	cluster1-02_clus2	cluster1-
01_clus1 none		
3/5/2022 19:21:20 -06:00	cluster1-02_clus2	cluster1-
01_clus2 none		

All ONTAP releases

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

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

```

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

```

6. Verify that the auto-revert command is enabled on all cluster LIFs:

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

Show example

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

Vserver	Logical Interface	Auto-revert
Cluster	cluster1-01_clus1	true
	cluster1-01_clus2	true
	cluster1-02_clus1	true
	cluster1-02_clus2	true

4 entries were displayed.

What's next?

[Install or upgrade the NX-OS software](#)

Install or upgrade the NX-OS software

Follow this procedure to install or upgrade the NX-OS software on the Nexus 9336C-FX2 cluster switch.

Before you begin, complete the procedure in [Prepare to install NX-OS and RCF](#).

Review requirements

What you'll need

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).

Suggested documentation

- [Cisco Ethernet switch page](#)

Consult the switch compatibility table for the supported ONTAP and NX-OS versions.

- [Software Upgrade and downgrade guides](#)

Refer to the appropriate software and upgrade guides available on the Cisco website for complete documentation on the Cisco switch upgrade and downgrade procedures.

- [Cisco Nexus 9000 and 3000 Upgrade and ISSU Matrix](#)

Provides information on Disruptive Upgrade/Downgrade for Cisco NX-OS software on Nexus 9000 Series Switches based on your current and target releases.

On the page, select **Disruptive Upgrade** and select your current release and target release from the

dropdown list.

About the examples

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.

Install the software

The procedure requires the use of both ONTAP commands and Cisco Nexus 9000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

Steps

1. Connect the cluster switch to the management network.
2. Use the ping command to verify connectivity to the server hosting the NX-OS software and the RCF.

Show example

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.
```

3. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show
Node/          Local  Discovered
Protocol       Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
              e0a    cs1                      Ethernet1/7      N9K-
C9336C-FX2
              e0d    cs2                      Ethernet1/7      N9K-
C9336C-FX2
cluster1-02/cdp
              e0a    cs1                      Ethernet1/8      N9K-
C9336C-FX2
              e0d    cs2                      Ethernet1/8      N9K-
C9336C-FX2
cluster1-03/cdp
              e0a    cs1                      Ethernet1/1/1    N9K-
C9336C-FX2
              e0b    cs2                      Ethernet1/1/1    N9K-
C9336C-FX2
cluster1-04/cdp
              e0a    cs1                      Ethernet1/1/2    N9K-
C9336C-FX2
              e0b    cs2                      Ethernet1/1/2    N9K-
C9336C-FX2
cluster1::*>
```

4. 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
```

Show example

```
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
```

```
Health Health Speed (Mbps)
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
```

Show example

```
cluster1::*> network interface show -role cluster
          Logical          Status      Network
Current   Current Is
Vserver   Interface          Admin/Oper Address/Mask      Node
Port      Home
-----
-----
Cluster
cluster1-01  cluster1-01_clus1  up/up      169.254.3.4/23
              e0a          true
cluster1-01  cluster1-01_clus2  up/up      169.254.3.5/23
              e0d          true
cluster1-02  cluster1-02_clus1  up/up      169.254.3.8/23
              e0a          true
cluster1-02  cluster1-02_clus2  up/up      169.254.3.9/23
              e0d          true
cluster1-03  cluster1-03_clus1  up/up      169.254.1.3/23
              e0a          true
cluster1-03  cluster1-03_clus2  up/up      169.254.1.1/23
              e0b          true
cluster1-04  cluster1-04_clus1  up/up      169.254.1.6/23
              e0a          true
cluster1-04  cluster1-04_clus2  up/up      169.254.1.7/23
              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
```

Show example

```
cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                                     Type                                     Address
Model
-----
cs1                                         cluster-network                         10.233.205.90
N9K-C9336C-FX2
  Serial Number: FOCXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                                     9.3(5)
  Version Source: CDP

cs2                                         cluster-network                         10.233.205.91
N9K-C9336C-FX2
  Serial Number: FOCXXXXXXGS
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                                     9.3(5)
  Version Source: CDP
cluster1::*>
```

5. Disable auto-revert on the cluster LIFs. The cluster LIFs fail over to the partner cluster switch and remain there as you perform the upgrade procedure on the targeted switch:

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

6. Copy the NX-OS software and EPLD images to the Nexus 9336C-FX2 switch.

Show example

```
cs2# copy sftp: bootflash: vrf management
Enter source filename: /code/nxos.9.3.5.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.9.3.5.bin /bootflash/nxos.9.3.5.bin
/code/nxos.9.3.5.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.9.3.5.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.9.3.5.img /bootflash/n9000-
epld.9.3.5.img
/code/n9000-epld.9.3.5.img 100% 161MB 9.5MB/s 00:16
sftp> exit
Copy complete, now saving to disk (please wait)...
Copy complete.
```

7. Verify the running version of the NX-OS software:

```
show version
```

Show example

```
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 08.38
  NXOS: version 9.3(4)
  BIOS compile time: 05/29/2020
  NXOS image file is: bootflash:///nxos.9.3.4.bin
  NXOS compile time: 4/28/2020 21:00:00 [04/29/2020 02:28:31]

Hardware
  cisco Nexus9000 C9336C-FX2 Chassis
  Intel(R) Xeon(R) CPU E5-2403 v2 @ 1.80GHz with 8154432 kB of
memory.
  Processor Board ID FOC20291J6K

  Device name: cs2
  bootflash: 53298520 kB
  Kernel uptime is 0 day(s), 0 hour(s), 3 minute(s), 42 second(s)
```



```
Last reset at 157524 usecs after Mon Nov  2 18:32:06 2020
Reason: Reset Requested by CLI command reload
System version: 9.3(4)
Service:

plugin
  Core Plugin, Ethernet Plugin

Active Package(s) :

cs2#
```

8. Install the NX-OS image.

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

Show example

```
cs2# install all nxos bootflash:nxos.9.3.5.bin
```

```
Installer will perform compatibility check first. Please wait.  
Installer is forced disruptive
```

```
Verifying image bootflash:/nxos.9.3.5.bin for boot variable "nxos".  
[] 100% -- SUCCESS
```

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

```
Preparing "nxos" version info using image bootflash:/nxos.9.3.5.bin.  
[] 100% -- SUCCESS
```

```
Preparing "bios" version info using image bootflash:/nxos.9.3.5.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(4)	9.3(5)
yes			
1	bios	v08.37(01/28/2020):v08.23(09/23/2015)	
v08.38(05/29/2020)		yes	

```
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.
```

9. Verify the new version of NX-OS software after the switch has rebooted:

```
show version
```

Show example

```
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 05.33
NXOS: version 9.3(5)
BIOS compile time: 09/08/2018
NXOS image file is: bootflash:///nxos.9.3.5.bin
NXOS compile time: 11/4/2018 21:00:00 [11/05/2018 06:11:06]
```

Hardware

```
cisco Nexus9000 C9336C-FX2 Chassis
Intel(R) Xeon(R) CPU E5-2403 v2 @ 1.80GHz with 8154432 kB of
memory.
Processor Board ID FOC20291J6K
```

```
Device name: cs2
bootflash: 53298520 kB
Kernel uptime is 0 day(s), 0 hour(s), 3 minute(s), 42 second(s)
```

```
Last reset at 277524 usecs after Mon Nov  2 22:45:12 2020
```

```
Reason: Reset due to upgrade
```

```
System version: 9.3(4)
```

```
Service:
```

```
plugin
```

```
Core Plugin, Ethernet Plugin
```

```
Active Package(s):
```

10. Upgrade the EPLD image and reboot the switch.

Show example



```
cs2# show version module 1 epld
```

EPLD	Device	Version
MI	FPGA	0x7
IO	FPGA	0x17
MI	FPGA2	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2

```
cs2# install epld bootflash:n9000-epld.9.3.5.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	0x07	0x07	No
1	SUP	IO FPGA	0x17	0x19	Yes
1	SUP	MI FPGA2	0x02	0x02	No

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

EPLDs upgraded.

Module 1 EPLD upgrade is successful.

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

Show example

```
cs2# show version module 1 epld
```

EPLD	Device	Version
MI	FPGA	0x7
IO	FPGA	0x19
MI	FPGA2	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2

12. Verify the health of cluster ports on the cluster.

a. Verify that cluster ports are up and healthy across all nodes in the cluster:

```
network port show -role cluster
```


Show example

```
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
```

```
Health Health Speed (Mbps)
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.

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol       Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
          e0a    cs1                        Ethernet1/7
N9K-C9336C-FX2
          e0d    cs2                        Ethernet1/7
N9K-C9336C-FX2
cluster01-2/cdp
          e0a    cs1                        Ethernet1/8
N9K-C9336C-FX2
          e0d    cs2                        Ethernet1/8
N9K-C9336C-FX2
cluster01-3/cdp
          e0a    cs1                        Ethernet1/1/1
N9K-C9336C-FX2
          e0b    cs2                        Ethernet1/1/1
N9K-C9336C-FX2
cluster1-04/cdp
          e0a    cs1                        Ethernet1/1/2
N9K-C9336C-FX2
          e0b    cs2                        Ethernet1/1/2
N9K-C9336C-FX2

cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                               Type                Address
Model
-----
-----
cs1                                   cluster-network     10.233.205.90
N9K-C9336C-FX2
  Serial Number: FOCXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                               9.3(5)
  Version Source: CDP

cs2                                   cluster-network     10.233.205.91
```

```

N9K-C9336C-FX2
  Serial Number: FOCXXXXXXGS
    Is Monitored: true
      Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                    9.3(5)
  Version Source: CDP

2 entries were displayed.

```

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.

```

13. Verify that the cluster is healthy:

```
cluster show
```

Show example

```

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::*>

```

14. Repeat steps 6 to 13 to install the NX-OS software on switch cs1.

15. Enable auto-revert on the cluster LIFs.

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

16. Verify that the cluster LIFs have reverted to their home port:

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

Show example

```
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::*>
```

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

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

What's next?

[Install or upgrade the Reference Configuration File \(RCF\)](#)

Install or upgrade the RCF

[Install or upgrade the Reference Configuration File \(RCF\) overview](#)

You install the Reference Configuration File (RCF) after setting up the Nexus 9336C-FX2 switch for the first time. You upgrade your RCF version when you have an existing

version of the RCF file installed on your switch.

See the Knowledge Base article [How to clear configuration on a Cisco interconnect switch while retaining remote connectivity](#) for further information when installing or upgrading your RCF.

Available RCF configurations

The following table describes the RCFs available for different configurations. Choose the RCF applicable to your configuration.

For specific port and VLAN usage details, refer to the banner and important notes section in your RCF.

RCF name	Description
2-Cluster-HA-Breakout	Supports two ONTAP clusters with at least eight nodes, including nodes that use shared Cluster+HA ports.
4-Cluster-HA-Breakout	Supports four ONTAP clusters with at least four nodes, including nodes that use shared Cluster+HA ports.
1-Cluster-HA	All ports are configured for 40/100GbE. Supports shared cluster/HA traffic on ports. Required for AFF A320, AFF A250, and FAS500f systems. Additionally, all ports can be used as dedicated cluster ports.
1-Cluster-HA-Breakout	Ports are configured for 4x10GbE breakout, 4x25GbE breakout (RCF 1.6+ on 100GbE switches), and 40/100GbE. Supports shared cluster/HA traffic on ports for nodes that use shared cluster/HA ports: AFF A320, AFF A250, and FAS500f systems. Additionally, all ports can be used as dedicated cluster ports.
Cluster-HA-Storage	Ports are configured for 40/100GbE for Cluster+HA, 4x10GbE breakout for Cluster and 4x25GbE breakout for Cluster+HA, and 100GbE for each Storage HA Pair.
Cluster	Two flavors of RCF with different allocations of 4x10GbE ports (breakout) and 40/100GbE ports. All FAS/AFF nodes are supported, except for AFF A320, AFF A250, and FAS500f systems.
Storage	All ports are configured for 100GbE NVMe storage connections.

Suggested documentation

- [Cisco Ethernet Switches \(NSS\)](#)

Consult the switch compatibility table for the supported ONTAP and RCF versions on the NetApp Support Site. Note that there can be command dependencies between the command syntax in the RCF and the syntax found in specific versions of NX-OS.

- [Cisco Nexus 3000 Series Switches](#)

Refer to the appropriate software and upgrade guides available on the Cisco website for complete documentation on the Cisco switch upgrade and downgrade procedures.

About the examples

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.

The examples in this procedure use four nodes. These nodes use two 10GbE cluster interconnect ports **e0a** and **e0b**. See the [Hardware Universe](#) to verify the correct cluster ports on your platforms.



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

For details of the available RCF configurations, see [Software install workflow](#).

Commands used

The procedure requires the use of both ONTAP commands and Cisco Nexus 9000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

What's next?

[Install the RCF](#) or [Upgrade your RCF](#).

Install the Reference Configuration File (RCF)

You install the Reference Configuration File (RCF) after setting up the Nexus 9336C-FX2 switch for the first time.

Before you begin

Verify the following installations and connections:

- A console connection to the switch. The console connection is optional if you have remote access to the switch.
- Switch cs1 and switch cs2 are powered up and the initial switch setup is complete (the Management IP address and SSH is setup).
- The desired NX-OS version has been installed.
- ISL connections between switches are connected.
- ONTAP node cluster ports are not connected.

Step 1: Install the RCF on the switches

1. Login to switch cs1 using SSH or by using a serial console.
2. Copy the RCF to the bootflash of switch cs1 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 9000 Series NX-OS Command Reference](#) guides.

Show example

This example shows TFTP being used to copy an RCF to the bootflash on switch cs1:

```
cs1# copy tftp: bootflash: vrf management
Enter source filename: Nexus_9336C_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)...
```

3. Apply the RCF previously downloaded to the bootflash.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 9000 Series NX-OS Command Reference](#) guides.

Show example

This example shows the RCF file `Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt` being installed on switch cs1:

```
cs1# copy Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt running-
config echo-commands
```

4. Examine the banner output from the `show banner motd` command. You must read and follow these instructions to ensure the proper configuration and operation of the switch.

Show example

```
cs1# show banner motd

*****
*****
* NetApp Reference Configuration File (RCF)
*
* Switch   : Nexus N9K-C9336C-FX2
* Filename : Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt
* Date     : 10-23-2020
* Version  : v1.6
*
* Port Usage:
* Ports 1- 3: Breakout mode (4x10G) Intra-Cluster Ports, int
e1/1/1-4, e1/2/1-4
, e1/3/1-4
* Ports 4- 6: Breakout mode (4x25G) Intra-Cluster/HA Ports, int
e1/4/1-4, e1/5/
1-4, e1/6/1-4
* Ports 7-34: 40/100GbE Intra-Cluster/HA Ports, int e1/7-34
* Ports 35-36: Intra-Cluster ISL Ports, int e1/35-36
*
* Dynamic breakout commands:
* 10G: interface breakout module 1 port <range> map 10g-4x
* 25G: interface breakout module 1 port <range> map 25g-4x
*
* Undo breakout commands and return interfaces to 40/100G
configuration in confi
g mode:
* no interface breakout module 1 port <range> map 10g-4x
* no interface breakout module 1 port <range> map 25g-4x
* interface Ethernet <interfaces taken out of breakout mode>
* inherit port-profile 40-100G
* priority-flow-control mode auto
* service-policy input HA
* exit
*
*****
*****
```

5. 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.

6. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
7. After you verify the RCF versions and switch settings are correct, copy the running-config file to the startup-config file.

```
copy running-config startup-config
```

Show example

```
cs1# copy running-config startup-config
[#####] 100% Copy complete
```

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 9000 Series NX-OS Command Reference](#).

8. Reboot switch cs1.

```
cs1# reload
```

This command will reboot the system. (y/n)? [n] **y**

9. Repeat steps 1 through 7 on switch cs2.
10. Connect the cluster ports of all nodes in the ONTAP cluster to switches cs1 and cs2.

Step 2: Verify the switch connections

1. Verify that the switch ports connected to the cluster ports are **up**.

```
show interface brief
```

Show example

```
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) --
.
.
```

2. Verify that the cluster nodes are in their correct cluster VLANs using the following commands:

```
show vlan brief
```

```
show interface trunk
```

Show example

```
cs1# show vlan brief
```

VLAN Name	Status	Ports
1 default	active	Pol, Eth1/1, Eth1/2, Eth1/3 Eth1/4, Eth1/5, Eth1/6, Eth1/7 Eth1/8, Eth1/35, Eth1/36 Eth1/9/1, Eth1/9/2, Eth1/9/3 Eth1/9/4, Eth1/10/1, Eth1/10/2 Eth1/10/3, Eth1/10/4
17 VLAN0017	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4 Eth1/5, Eth1/6, Eth1/7, Eth1/8 Eth1/9/1, Eth1/9/2, Eth1/9/3 Eth1/9/4, Eth1/10/1, Eth1/10/2 Eth1/10/3, Eth1/10/4
18 VLAN0018	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4 Eth1/5, Eth1/6, Eth1/7, Eth1/8 Eth1/9/1, Eth1/9/2, Eth1/9/3 Eth1/9/4, Eth1/10/1, Eth1/10/2 Eth1/10/3, Eth1/10/4
31 VLAN0031	active	Eth1/11, Eth1/12, Eth1/13 Eth1/14, Eth1/15, Eth1/16 Eth1/17, Eth1/18, Eth1/19 Eth1/20, Eth1/21, Eth1/22
32 VLAN0032	active	Eth1/23, Eth1/24, Eth1/25

```

Eth1/28                               Eth1/26, Eth1/27,
Eth1/31                               Eth1/29, Eth1/30,
Eth1/34                               Eth1/32, Eth1/33,
33   VLAN0033                         active   Eth1/11, Eth1/12,
Eth1/13                               Eth1/14, Eth1/15,
Eth1/16                               Eth1/17, Eth1/18,
Eth1/19                               Eth1/20, Eth1/21,
Eth1/22                               Eth1/23, Eth1/24,
34   VLAN0034                         active   Eth1/26, Eth1/27,
Eth1/25                               Eth1/29, Eth1/30,
Eth1/28                               Eth1/32, Eth1/33,
Eth1/31
Eth1/34

```

```
cs1# show interface trunk
```

```

-----
Port          Native  Status      Port
              Vlan               Channel
-----
Eth1/1        1       trunking    --
Eth1/2        1       trunking    --
Eth1/3        1       trunking    --
Eth1/4        1       trunking    --
Eth1/5        1       trunking    --
Eth1/6        1       trunking    --
Eth1/7        1       trunking    --
Eth1/8        1       trunking    --
Eth1/9/1      1       trunking    --
Eth1/9/2      1       trunking    --
Eth1/9/3      1       trunking    --
Eth1/9/4      1       trunking    --
Eth1/10/1     1       trunking    --
Eth1/10/2     1       trunking    --
Eth1/10/3     1       trunking    --
Eth1/10/4     1       trunking    --
Eth1/11       33      trunking    --

```

Eth1/12	33	trunking	--
Eth1/13	33	trunking	--
Eth1/14	33	trunking	--
Eth1/15	33	trunking	--
Eth1/16	33	trunking	--
Eth1/17	33	trunking	--
Eth1/18	33	trunking	--
Eth1/19	33	trunking	--
Eth1/20	33	trunking	--
Eth1/21	33	trunking	--
Eth1/22	33	trunking	--
Eth1/23	34	trunking	--
Eth1/24	34	trunking	--
Eth1/25	34	trunking	--
Eth1/26	34	trunking	--
Eth1/27	34	trunking	--
Eth1/28	34	trunking	--
Eth1/29	34	trunking	--
Eth1/30	34	trunking	--
Eth1/31	34	trunking	--
Eth1/32	34	trunking	--
Eth1/33	34	trunking	--
Eth1/34	34	trunking	--
Eth1/35	1	trnk-bndl	Pol
Eth1/36	1	trnk-bndl	Pol
Pol	1	trunking	--

Port	Vlans Allowed on Trunk
------	------------------------

Eth1/1	1,17-18
Eth1/2	1,17-18
Eth1/3	1,17-18
Eth1/4	1,17-18
Eth1/5	1,17-18
Eth1/6	1,17-18
Eth1/7	1,17-18
Eth1/8	1,17-18
Eth1/9/1	1,17-18
Eth1/9/2	1,17-18
Eth1/9/3	1,17-18
Eth1/9/4	1,17-18
Eth1/10/1	1,17-18
Eth1/10/2	1,17-18
Eth1/10/3	1,17-18
Eth1/10/4	1,17-18

```
Eth1/11      31, 33
Eth1/12      31, 33
Eth1/13      31, 33
Eth1/14      31, 33
Eth1/15      31, 33
Eth1/16      31, 33
Eth1/17      31, 33
Eth1/18      31, 33
Eth1/19      31, 33
Eth1/20      31, 33
Eth1/21      31, 33
Eth1/22      31, 33
Eth1/23      32, 34
Eth1/24      32, 34
Eth1/25      32, 34
Eth1/26      32, 34
Eth1/27      32, 34
Eth1/28      32, 34
Eth1/29      32, 34
Eth1/30      32, 34
Eth1/31      32, 34
Eth1/32      32, 34
Eth1/33      32, 34
Eth1/34      32, 34
Eth1/35      1
Eth1/36      1
Po1         1
..
..
..
..
..
```



For specific port and VLAN usage details, refer to the banner and important notes section in your RCF.

3. Verify that the ISL between cs1 and cs2 is functional:

```
show port-channel summary
```

Show example

```
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         Eth1/35 (P)       Eth1/36 (P)
cs1#
```

Step 3: Set up your ONTAP cluster

NetApp recommends that you use System Manager to set up new clusters.

System Manager provides a simple and easy workflow for cluster set up and configuration including assigning a node management IP address, initializing the cluster, creating a local tier, configuring protocols and provisioning initial storage.

Go to [Configure ONTAP on a new cluster with System Manager](#) for setup instructions.

What's next?

[Verify SSH configuration.](#)

Upgrade your Reference Configuration File (RCF)

You upgrade your RCF version when you have an existing version of the RCF file installed on your operational switches.

Before you begin

Make sure you have the following:

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).
- The current RCF.
- If you are updating your RCF version, you need a boot configuration in the RCF that reflects the desired boot images.

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.



No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure migrates all of the cluster LIFs to the operational partner switch while performing the steps on the target switch.



Before installing a new switch software version and RCFs, you must erase the switch settings and perform basic configuration. You must be connected to the switch using the serial console or have preserved basic configuration information prior to erasing the switch settings.

Step 1: Prepare for the upgrade

1. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show
Node/          Local  Discovered
Protocol       Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
              e0a    cs1                      Ethernet1/7      N9K-
C9336C
              e0d    cs2                      Ethernet1/7      N9K-
C9336C
cluster1-02/cdp
              e0a    cs1                      Ethernet1/8      N9K-
C9336C
              e0d    cs2                      Ethernet1/8      N9K-
C9336C
cluster1-03/cdp
              e0a    cs1                      Ethernet1/1/1    N9K-
C9336C
              e0b    cs2                      Ethernet1/1/1    N9K-
C9336C
cluster1-04/cdp
              e0a    cs1                      Ethernet1/1/2    N9K-
C9336C
              e0b    cs2                      Ethernet1/1/2    N9K-
C9336C
cluster1::*>
```

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
```

Show example

```
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
```

```
Health Health Speed (Mbps)
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
```

Show example

```
cluster1::*> network interface show -role cluster
          Logical          Status      Network
Current   Current Is
Vserver   Interface          Admin/Oper Address/Mask      Node
Port      Home
-----
-----
Cluster
cluster1-01  cluster1-01_clus1  up/up      169.254.3.4/23
              e0a      true
cluster1-01  cluster1-01_clus2  up/up      169.254.3.5/23
              e0d      true
cluster1-02  cluster1-02_clus1  up/up      169.254.3.8/23
              e0a      true
cluster1-02  cluster1-02_clus2  up/up      169.254.3.9/23
              e0d      true
cluster1-03  cluster1-03_clus1  up/up      169.254.1.3/23
              e0a      true
cluster1-03  cluster1-03_clus2  up/up      169.254.1.1/23
              e0b      true
cluster1-04  cluster1-04_clus1  up/up      169.254.1.6/23
              e0a      true
cluster1-04  cluster1-04_clus2  up/up      169.254.1.7/23
              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
```

Show example

```
cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                               Type                Address
Model
-----
cs1                                   cluster-network    10.233.205.90
N9K-C9336C
  Serial Number: FOCXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                               9.3(5)
  Version Source: CDP

cs2                                   cluster-network    10.233.205.91
N9K-C9336C
  Serial Number: FOCXXXXXXGS
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                               9.3(5)
  Version Source: CDP
cluster1::*>
```

3. Disable auto-revert on the cluster LIFs.

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

Step 2: Configure ports

1. On cluster switch cs1, shut down the ports connected to the cluster ports of the nodes.

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



Make sure to shutdown **all** connected cluster ports to avoid any network connection issues. See the Knowledge Base article [Node out of quorum when migrating cluster LIF during switch OS upgrade](#) for further details.

2. Verify that the cluster LIFs have failed over to the ports hosted on cluster switch cs1. This might take a few seconds.

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

Show example

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

Current Is	Logical	Status	Network	Current
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::*>
```

3. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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::*>
```

4. If you have not already done so, save a copy of the current switch configuration by copying the output of the following command to a text file:

```
show running-config
```

- a. Record any custom additions between the current running-config and the RCF file in use (such as an SNMP configuration for your organization).
- b. For NX-OS 10.2 and newer use the `show diff running-config` command to compare with the saved RCF file in the bootflash. Otherwise, use a third part diff/compare tool.

5. Save basic configuration details to the `write_erase.cfg` file on the bootflash.

```
switch# show run | i "username admin password" > bootflash:write_erase.cfg
```

```
switch# show run | section "vrf context management" >>
bootflash:write_erase.cfg
```

```
switch# show run | section "interface mgmt0" >> bootflash:write_erase.cfg
```

```
switch# show run | section "switchname" >> bootflash:write_erase.cfg
```

6. Issue the write erase command to erase the current saved configuration:

```
switch# write erase
```

Warning: This command will erase the startup-configuration.

```
Do you wish to proceed anyway? (y/n) [n] y
```

7. Copy the previously saved basic configuration into the startup configuration.

```
switch# copy write_erase.cfg startup-config
```

8. Perform a reboot of the switch:

```
switch# reload
```


This command will reboot the system. (y/n)? [n] y

9. After the management IP address is reachable again, log in to the switch through SSH.

You may need to update host file entries related to the SSH keys.

10. Copy the RCF to the bootflash of switch cs1 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 9000 Series NX-OS Command Reference](#) guides.

Show example

This example shows TFTP being used to copy an RCF to the bootflash on switch cs1:

```
cs1# copy tftp: bootflash: vrf management
Enter source filename: Nexus_9336C_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)...
```

11. Apply the RCF previously downloaded to the bootflash.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 9000 Series NX-OS Command Reference](#) guides.

Show example

This example shows the RCF file `Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt` being installed on switch cs1:

```
cs1# copy Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt running-
config echo-commands
```

12. Examine the banner output from the `show banner motd` command. You must read and follow these instructions to ensure the proper configuration and operation of the switch.

Show example

```
cs1# show banner motd

*****
*****
* NetApp Reference Configuration File (RCF)
*
* Switch      : Nexus N9K-C9336C-FX2
* Filename    : Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt
* Date       : 10-23-2020
* Version    : v1.6
*
* Port Usage:
* Ports 1- 3: Breakout mode (4x10G) Intra-Cluster Ports, int
e1/1/1-4, e1/2/1-4
, e1/3/1-4
* Ports 4- 6: Breakout mode (4x25G) Intra-Cluster/HA Ports, int
e1/4/1-4, e1/5/
1-4, e1/6/1-4
* Ports 7-34: 40/100GbE Intra-Cluster/HA Ports, int e1/7-34
* Ports 35-36: Intra-Cluster ISL Ports, int e1/35-36
*
* Dynamic breakout commands:
* 10G: interface breakout module 1 port <range> map 10g-4x
* 25G: interface breakout module 1 port <range> map 25g-4x
*
* Undo breakout commands and return interfaces to 40/100G
configuration in confi
g mode:
* no interface breakout module 1 port <range> map 10g-4x
* no interface breakout module 1 port <range> map 25g-4x
* interface Ethernet <interfaces taken out of breakout mode>
* inherit port-profile 40-100G
* priority-flow-control mode auto
* service-policy input HA
* exit
*
*****
*****
```

13. 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.

14. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
15. After you verify the RCF versions, custom additions, 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 9000 Series NX-OS Command Reference](#) guides.

```
cs1# copy running-config startup-config
```

```
[ ] 100% Copy complete
```

16. Reboot switch cs1. You can ignore the “cluster switch health monitor” alerts and “cluster ports down” events reported on the nodes while the switch reboots.

```
cs1# reload
```

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

17. Verify the health of cluster ports on the cluster.
 - a. Verify that cluster ports are up and healthy across all nodes in the cluster:

```
network port show -role cluster
```

Show example

```
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
```

```
Health Health Speed (Mbps)
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.

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol       Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
          e0a   cs1                      Ethernet1/7
N9K-C9336C
          e0d   cs2                      Ethernet1/7
N9K-C9336C
cluster01-2/cdp
          e0a   cs1                      Ethernet1/8
N9K-C9336C
          e0d   cs2                      Ethernet1/8
N9K-C9336C
cluster01-3/cdp
          e0a   cs1                      Ethernet1/1/1
N9K-C9336C
          e0b   cs2                      Ethernet1/1/1
N9K-C9336C
cluster1-04/cdp
          e0a   cs1                      Ethernet1/1/2
N9K-C9336C
          e0b   cs2                      Ethernet1/1/2
N9K-C9336C

cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                               Type                Address
Model
-----
-----
cs1                                   cluster-network    10.233.205.90
NX9-C9336C
    Serial Number: FOCXXXXXXGD
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                                9.3(5)
    Version Source: CDP

cs2                                   cluster-network    10.233.205.91
```

```

NX9-C9336C
  Serial Number: FOCXXXXXXGS
    Is Monitored: true
      Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                9.3(5)
  Version Source: CDP

2 entries were displayed.

```

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.

```

18. Verify that the cluster is healthy:

```
cluster show
```

Show example

```

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::*>

```

19. Repeat steps 1 to 18 on switch cs2.

20. Enable auto-revert on the cluster LIFs.

```

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

```

Step 3: Verify the cluster network configuration and cluster health

1. Verify that the switch ports connected to the cluster ports are **up**.

```
show interface brief
```

Show example

```
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)  --
.
.
```

2. Verify that the expected nodes are still connected:

```
show cdp neighbors
```


Show example

```
cs1# 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
node1             Eth1/1        133     H           FAS2980
e0a
node2             Eth1/2        133     H           FAS2980
e0a
cs1               Eth1/35       175     R S I s     N9K-C9336C
Eth1/35
cs1               Eth1/36       175     R S I s     N9K-C9336C
Eth1/36

Total entries displayed: 4
```

3. Verify that the cluster nodes are in their correct cluster VLANs using the following commands:

```
show vlan brief
```

```
show interface trunk
```

Show example

```
cs1# show vlan brief
```

VLAN Name	Status	Ports
1 default	active	Pol, Eth1/1, Eth1/2, Eth1/3 Eth1/4, Eth1/5, Eth1/6, Eth1/7 Eth1/8, Eth1/35, Eth1/36 Eth1/9/1, Eth1/9/2, Eth1/9/3 Eth1/9/4, Eth1/10/1, Eth1/10/2 Eth1/10/3, Eth1/10/4
17 VLAN0017	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4 Eth1/5, Eth1/6, Eth1/7, Eth1/8 Eth1/9/1, Eth1/9/2, Eth1/9/3 Eth1/9/4, Eth1/10/1, Eth1/10/2 Eth1/10/3, Eth1/10/4
18 VLAN0018	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4 Eth1/5, Eth1/6, Eth1/7, Eth1/8 Eth1/9/1, Eth1/9/2, Eth1/9/3 Eth1/9/4, Eth1/10/1, Eth1/10/2 Eth1/10/3, Eth1/10/4
31 VLAN0031	active	Eth1/11, Eth1/12, Eth1/13 Eth1/14, Eth1/15, Eth1/16 Eth1/17, Eth1/18, Eth1/19 Eth1/20, Eth1/21, Eth1/22
32 VLAN0032	active	Eth1/23, Eth1/24, Eth1/25

```

Eth1/28                               Eth1/26, Eth1/27,
Eth1/31                               Eth1/29, Eth1/30,
Eth1/34                               Eth1/32, Eth1/33,
33   VLAN0033                         active   Eth1/11, Eth1/12,
Eth1/13                               Eth1/14, Eth1/15,
Eth1/16                               Eth1/17, Eth1/18,
Eth1/19                               Eth1/20, Eth1/21,
Eth1/22                               Eth1/23, Eth1/24,
34   VLAN0034                         active   Eth1/26, Eth1/27,
Eth1/25                               Eth1/29, Eth1/30,
Eth1/28                               Eth1/32, Eth1/33,
Eth1/31
Eth1/34

```

```
cs1# show interface trunk
```

```

-----
Port          Native  Status      Port
              Vlan               Channel
-----
Eth1/1        1       trunking    --
Eth1/2        1       trunking    --
Eth1/3        1       trunking    --
Eth1/4        1       trunking    --
Eth1/5        1       trunking    --
Eth1/6        1       trunking    --
Eth1/7        1       trunking    --
Eth1/8        1       trunking    --
Eth1/9/1      1       trunking    --
Eth1/9/2      1       trunking    --
Eth1/9/3      1       trunking    --
Eth1/9/4      1       trunking    --
Eth1/10/1     1       trunking    --
Eth1/10/2     1       trunking    --
Eth1/10/3     1       trunking    --
Eth1/10/4     1       trunking    --
Eth1/11       33      trunking    --

```

Eth1/12	33	trunking	--
Eth1/13	33	trunking	--
Eth1/14	33	trunking	--
Eth1/15	33	trunking	--
Eth1/16	33	trunking	--
Eth1/17	33	trunking	--
Eth1/18	33	trunking	--
Eth1/19	33	trunking	--
Eth1/20	33	trunking	--
Eth1/21	33	trunking	--
Eth1/22	33	trunking	--
Eth1/23	34	trunking	--
Eth1/24	34	trunking	--
Eth1/25	34	trunking	--
Eth1/26	34	trunking	--
Eth1/27	34	trunking	--
Eth1/28	34	trunking	--
Eth1/29	34	trunking	--
Eth1/30	34	trunking	--
Eth1/31	34	trunking	--
Eth1/32	34	trunking	--
Eth1/33	34	trunking	--
Eth1/34	34	trunking	--
Eth1/35	1	trnk-bndl	Pol
Eth1/36	1	trnk-bndl	Pol
Pol	1	trunking	--

Port	Vlans Allowed on Trunk
------	------------------------

Eth1/1	1,17-18
Eth1/2	1,17-18
Eth1/3	1,17-18
Eth1/4	1,17-18
Eth1/5	1,17-18
Eth1/6	1,17-18
Eth1/7	1,17-18
Eth1/8	1,17-18
Eth1/9/1	1,17-18
Eth1/9/2	1,17-18
Eth1/9/3	1,17-18
Eth1/9/4	1,17-18
Eth1/10/1	1,17-18
Eth1/10/2	1,17-18
Eth1/10/3	1,17-18
Eth1/10/4	1,17-18

```
Eth1/11      31, 33
Eth1/12      31, 33
Eth1/13      31, 33
Eth1/14      31, 33
Eth1/15      31, 33
Eth1/16      31, 33
Eth1/17      31, 33
Eth1/18      31, 33
Eth1/19      31, 33
Eth1/20      31, 33
Eth1/21      31, 33
Eth1/22      31, 33
Eth1/23      32, 34
Eth1/24      32, 34
Eth1/25      32, 34
Eth1/26      32, 34
Eth1/27      32, 34
Eth1/28      32, 34
Eth1/29      32, 34
Eth1/30      32, 34
Eth1/31      32, 34
Eth1/32      32, 34
Eth1/33      32, 34
Eth1/34      32, 34
Eth1/35      1
Eth1/36      1
Po1          1
..
..
..
..
..
```



For specific port and VLAN usage details, refer to the banner and important notes section in your RCF.

4. Verify that the ISL between cs1 and cs2 is functional:

```
show port-channel summary
```

Show example

```
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         Eth1/35 (P)       Eth1/36 (P)
cs1#
```

5. Verify that the cluster LIFs have reverted to their home port:

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

Show example

```
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::*>
```

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

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

6. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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::*>
```

7. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

Packet	Source	Destination
Node	Date	LIF
Loss		
-----	-----	-----
node1		
3/5/2022 19:21:18 -06:00	cluster1-01_clus2	cluster1-02-
clus1 none		
3/5/2022 19:21:20 -06:00	cluster1-01_clus2	cluster1-
02_clus2 none		
node2		
3/5/2022 19:21:18 -06:00	cluster1-02_clus2	cluster1-
01_clus1 none		
3/5/2022 19:21:20 -06:00	cluster1-02_clus2	cluster1-
01_clus2 none		

All ONTAP releases

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

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

```

cluster1::~*> cluster ping-cluster -node local
Host is 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)

```

What's next?

[Verify SSH configuration.](#)

Verify your SSH configuration

If you are using the Ethernet Switch Health Monitor (CSHM) and log collection features,

verify that SSH and SSH keys are enabled on the cluster switches.

Steps

1. Verify that SSH is enabled:

```
(switch) show ssh server  
ssh version 2 is enabled
```

2. Verify that the SSH keys are enabled:

```
show ssh key
```

Show example

```
(switch)# show ssh key

rsa Keys generated:Fri Jun 28 02:16:00 2024

ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQGDINrD52Q586wTGJjFABjBlFaA23EpDrZ2sDCew
l7nwlIoC6HBejxluIObAH8hrW8kR+gj0ZAfPpNeLGTg3APj/yIPTBoIZZxbWRShywAM5
PqyxWwRb7kp9Zt1YHzVuHYpSO82KUDowKrL6lox/YtpKoZUDZjrZjAp8hTv3JZsPgQ==

bitcount:1024
fingerprint:
SHA256:aHwhpzo7+YCDsrp3isJv2uVGz+mjMMokqdMeXVVXfdo

could not retrieve dsa key information

ecdsa Keys generated:Fri Jun 28 02:30:56 2024

ecdsa-sha2-nistp521
AAAAE2VjZHNhLXNoYTItbmlzdHA1MjEAAAABmlzdHA1MjEAAACFBABJ+ZX5SFKhS57e
vkE273e0VoqZi4/32dt+f14fBuKv80MjMsmLfjKtCWylwgVt1Zi+C5TIBbugpzez529z
kFSF0ADb8JaGCoaAYe2HvWR/f6QLbKbqVIewCdqWgxzrIY5BPP5GBdxQJMBiOwEdnHg1
u/9Pzh/Vz9cHDcCW9qGE780QHA==

bitcount:521
fingerprint:
SHA256:TFGe2hXn6QIpcs/vyHzftHJ7Dceg0vQaULYRALZeHwQ

(switch)# show feature | include scpServer
scpServer          1          enabled
(switch)# show feature | include ssh
sshServer          1          enabled
(switch)#
```



When enabling FIPS, you must change the bitcount to 256 on the switch using the command `ssh key ecdsa 256 force`. See [Configure network security using FIPS](#) for more details.

What's next?

[Configure switch health monitoring.](#)

Migrate switches

Migrate from a NetApp CN1610 cluster switch to a Cisco 9336C-FX2 cluster switch

You can migrate NetApp CN1610 cluster switches for an ONTAP cluster to Cisco 9336C-FX2 cluster switches. This is a nondisruptive procedure.

Review requirements

You must be aware of certain configuration information, port connections and cabling requirements when you are replacing NetApp CN1610 cluster switches with Cisco 9336C-FX2 cluster switches. Also, verify the switch serial number to ensure that the correct switch is migrated.

Supported switches

The following cluster switches are supported:

- NetApp CN1610
- Cisco 9336C-FX2

For details of supported ports and their configurations, see the [Hardware Universe](#).

What you'll need

Verify that your configuration meets the following requirements:

- The existing cluster is correctly set up and functioning.
- All cluster ports are in the **up** state to ensure nondisruptive operations.
- The Cisco 9336C-FX2 cluster switches are configured and operating under the correct version of NX-OS installed with the reference configuration file (RCF) applied.
- The existing cluster network configuration has the following:
 - A redundant and fully functional NetApp cluster using NetApp CN1610 switches.
 - Management connectivity and console access to both the NetApp CN1610 switches and the new switches.
 - All cluster LIFs in the up state with the cluster LIFs are on their home ports.
- Some of the ports are configured on Cisco 9336C-FX2 switches to run at 40GbE or 100GbE.
- You have planned, migrated, and documented 40GbE and 100GbE connectivity from nodes to Cisco 9336C-FX2 cluster switches.

Migrate the switches

About the examples

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

- The existing CN1610 cluster switches are *C1* and *C2*.
- The new 9336C-FX2 cluster switches are *cs1* and *cs2*.
- The nodes are *node1* and *node2*.
- The cluster LIFs are *node1_clus1* and *node1_clus2* on node 1, and *node2_clus1* and *node2_clus2* on node 2 respectively.
- The `cluster1::*>` prompt indicates the name of the cluster.
- The cluster ports used in this procedure are *e3a* and *e3b*.

About this task

This procedure covers the following scenario:

- Switch C2 is replaced by switch cs2 first.
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - All cluster LIFs fail over to the new switch cs2.
 - The cabling between the nodes and C2 is then disconnected from C2 and reconnected to cs2.
- Switch C1 is replaced by switch cs1.
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - All cluster LIFs fail over to the new switch cs1.
 - The cabling between the nodes and C1 is then disconnected from C1 and reconnected to cs1.



No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure fails over all of the cluster LIFs to the operational partner switch while performing the steps on the target switch.

Step 1: Prepare for migration

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
```

where *x* is the duration of the maintenance window in hours.

2. Change the privilege level to advanced, entering *y* when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (*>) appears.

3. Disable auto-revert on the cluster LIFs.

By disabling auto-revert for this procedure, the cluster LIFs will not automatically move back to their home port. They remain on the current port while it continues to be up and operational.

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

Step 2: Configure ports and cabling

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

Each port should display up for `Link` and `healthy` for `Health Status`.

- a. Display the network port attributes:

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

Show example

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

Node: node1

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
-----
e3a         Cluster   Cluster           up   9000  auto/100000
healthy    false
e3b         Cluster   Cluster           up   9000  auto/100000
healthy    false

Node: node2

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
-----
e3a         Cluster   Cluster           up   9000  auto/100000
healthy    false
e3b         Cluster   Cluster           up   9000  auto/100000
healthy    false
```

- b. Display information about the LIFs and their designated home nodes:

```
network interface show -vserver Cluster
```

Each LIF should display up/up for Status Admin/Oper and true for Is Home.

Show example

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

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

Cluster				
e3a	node1_clus1	up/up	169.254.209.69/16	node1
e3b	true			
e3a	node1_clus2	up/up	169.254.49.125/16	node1
e3b	true			
e3a	node2_clus1	up/up	169.254.47.194/16	node2
e3a	true			
e3a	node2_clus2	up/up	169.254.19.183/16	node2
e3b	true			

- The cluster ports on each node are connected to existing cluster switches in the following way (from the nodes' perspective) using the command:

```
network device-discovery show -protocol
```

Show example

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


Node/	Local	Discovered		
Protocol	Port	Device (LLDP: ChassisID)	Interface	
Platform				

node1	/cdp			
	e3a	C1 (6a:ad:4f:98:3b:3f)	0/1	-
	e3b	C2 (6a:ad:4f:98:4c:a4)	0/1	-
node2	/cdp			
	e3a	C1 (6a:ad:4f:98:3b:3f)	0/2	-
	e3b	C2 (6a:ad:4f:98:4c:a4)	0/2	-

- The cluster ports and switches are connected in the following way (from the switches' perspective) using the command:

```
show cdp neighbors
```


Show example



```
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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e3a	Eth1/1	124	H	AFF-A400
node2 e3a	Eth1/2	124	H	AFF-A400
C2 0/13	0/13	179	S I s	CN1610
C2 0/14	0/14	175	S I s	CN1610
C2 0/15	0/15	179	S I s	CN1610
C2 0/16	0/16	175	S I s	CN1610

```
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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e3b	Eth1/1	124	H	AFF-A400
node2 e3b	Eth1/2	124	H	AFF-A400
C1 0/13	0/13	175	S I s	CN1610
C1 0/14	0/14	175	S I s	CN1610
C1 0/15	0/15	175	S I s	CN1610
C1 0/16	0/16	175	S I s	CN1610

4. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e3a
Cluster node1_clus2 169.254.49.125 node1 e3b
Cluster node2_clus1 169.254.47.194 node2 e3a
Cluster node2_clus2 169.254.19.183 node2 e3b
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)

```

5. On switch C2, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs.



Do not attempt to manually migrate the cluster LIFs.

```

(C2)# configure
(C2)(Config)# interface 0/1-0/12
(C2)(Interface 0/1-0/12)# shutdown
(C2)(Interface 0/1-0/12)# exit
(C2)(Config)# exit

```

6. Move the node cluster ports from the old switch C2 to the new switch cs2, using appropriate cabling supported by Cisco 9336C-FX2.
7. Display the network port attributes:

```
network port show -ipspace Cluster
```

Show example

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

```
Node: node1
```

```
Ignore
```

						Speed (Mbps)	Health
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/100000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

```
Node: node2
```

```
Ignore
```

						Speed (Mbps)	Health
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
-----	-----						
e3a	Cluster	Cluster		up	9000	auto/100000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

8. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

```
network device-discovery show -protocol
```

Show example

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

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	
node1 CN1610	/cdp e3a	C1 (6a:ad:4f:98:3b:3f)	0/1	
C9336C-FX2	e3b	cs2 (b8:ce:f6:19:1a:7e)	Ethernet1/1/1	N9K-
node2 CN1610	/cdp e3a	C1 (6a:ad:4f:98:3b:3f)	0/2	
C9336C-FX2	e3b	cs2 (b8:ce:f6:19:1b:96)	Ethernet1/1/2	N9K-

9. On switch cs2, verify that all node cluster ports are up:

```
network interface show -vserver Cluster
```

Show example

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

Current Is Vserver Port	Logical Interfac Home	Status Admin/Oper	Network Address/Mask	Current Node
Cluster	node1_clus1	up/up	169.254.3.4/16	node1
e0b	false			
	node1_clus2	up/up	169.254.3.5/16	node1
e0b	true			
	node2_clus1	up/up	169.254.3.8/16	node2
e0b	false			
	node2_clus2	up/up	169.254.3.9/16	node2
e0b	true			

10. On switch C1, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs.

```
(C1) # configure
(C1) (Config) # interface 0/1-0/12
(C1) (Interface 0/1-0/12) # shutdown
(C1) (Interface 0/1-0/12) # exit
(C1) (Config) # exit
```

11. Move the node cluster ports from the old switch C1 to the new switch cs1, using appropriate cabling supported by Cisco 9336C-FX2.
12. Verify the final configuration of the cluster:

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

Each port should display `up` for Link and `healthy` for Health Status.

Show example

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

```
Node: node1
```

```
Ignore
```

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

e3a	Cluster	Cluster		up	9000	auto/100000		healthy	false
e3b	Cluster	Cluster		up	9000	auto/100000		healthy	false

```
Node: node2
```

```
Ignore
```

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

e3a	Cluster	Cluster		up	9000	auto/100000		healthy	false
e3b	Cluster	Cluster		up	9000	auto/100000		healthy	false

13. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

```
network device-discovery show -protocol
```

Show example

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

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	
node1	/cdp			
C9336C-FX2	e3a	cs1 (b8:ce:f6:19:1a:7e)	Ethernet1/1/1	N9K-
C9336C-FX2	e3b	cs2 (b8:ce:f6:19:1b:96)	Ethernet1/1/2	N9K-
node2	/cdp			
C9336C-FX2	e3a	cs1 (b8:ce:f6:19:1a:7e)	Ethernet1/1/1	N9K-
C9336C-FX2	e3b	cs2 (b8:ce:f6:19:1b:96)	Ethernet1/1/2	N9K-

14. On switches cs1 and cs2, verify that all node cluster ports are up:

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

Show example

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

```
Node: node1
```

```
Ignore
```

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

```
-----  
-----  
e0a      Cluster      Cluster      up    9000  auto/10000  
healthy  false  
e0b      Cluster      Cluster      up    9000  auto/10000  
healthy  false
```

```
Node: node2
```

```
Ignore
```

						Speed(Mbps)	Health
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
```

15. Verify that both nodes each have one connection to each switch:

```
network device-discovery show -protocol
```

Show example

The following example shows the appropriate results for both switches:

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
node1         /cdp
              e0a    cs1 (b8:ce:f6:19:1b:42)   Ethernet1/1/1   N9K-
C9336C-FX2
              e0b    cs2 (b8:ce:f6:19:1b:96)   Ethernet1/1/2   N9K-
C9336C-FX2
node2         /cdp
              e0a    cs1 (b8:ce:f6:19:1b:42)   Ethernet1/1/1   N9K-
C9336C-FX2
              e0b    cs2 (b8:ce:f6:19:1b:96)   Ethernet1/1/2   N9K-
C9336C-FX2
```

Step 3: Verify the configuration

1. Enable auto-revert on the cluster LIFs:

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

2. Verify that all cluster network LIFs are back on their home ports:

```
network interface show
```

Show example

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

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

Cluster				
e3a	node1_clus1	up/up	169.254.209.69/16	node1
e3b	true			
e3a	node1_clus2	up/up	169.254.49.125/16	node1
e3b	true			
e3a	node2_clus1	up/up	169.254.47.194/16	node2
e3b	true			
e3a	node2_clus2	up/up	169.254.19.183/16	node2
e3b	true			

3. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

What's next?

[Configure switch health monitoring.](#)

Migrate from an older Cisco switch to a Cisco Nexus 9336C-FX2 switch

You can perform a nondisruptive migration from an older Cisco cluster switch to a Cisco Nexus 9336C-FX2 cluster network switch.

Review requirements

Ensure that:

- You have verified the switch serial number to ensure that the correct switch is migrated.
- Some of the ports on Nexus 9336C-FX2 switches are configured to run at 10GbE or 40GbE.
- The 10GbE and 40GbE connectivity from nodes to Nexus 9336C-FX2 cluster switches have been planned, migrated, and documented.
- The cluster is fully functioning (there should be no errors in the logs or similar issues).

- Initial customization of the Cisco Nexus 9336C-FX2 switches is complete, so that:
 - 9336C-FX2 switches are running the latest recommended version of software.
 - Confirm that Reference Configuration Files (RCFs) have been fully applied to any new switches before migrating the LIFs to the new switches.
 - Check the running and startup configurations on both switches prior to shifting network traffic.
 - Any site customization, such as DNS, NTP, SMTP, SNMP, and SSH, are configured on the new switches.
- You have access to the switch compatibility table on the [Cisco Ethernet Switches](#) page for the supported ONTAP, NX-OS, and RCF versions.
- You have reviewed the appropriate software and upgrade guides available on the Cisco web site for the Cisco switch upgrade and downgrade procedures at [Cisco Nexus 9000 Series Switches Support](#) page.



If you are changing the port speed of the e0a and e1a cluster ports on AFF A800 or AFF C800 systems, you might observe malformed packets being received after the speed conversion. See [Bug 1570339](#) and the Knowledge Base article [CRC errors on T6 ports after converting from 40GbE to 100GbE](#) for guidance.

Migrate the switches

About the examples

The examples in this procedure use two nodes. These nodes use two 10GbE cluster interconnect ports e0a and e0b. See the [Hardware Universe](#) to verify the correct cluster ports on your platforms.




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

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

- The names of the existing two Cisco switches are **cs1** and **cs2**
- The new Nexus 9336C-FX2 cluster switches are **cs1-new** and **cs2-new**.
- The node names are **node1** and **node2**.
- The cluster LIF names are **node1_clus1** and **node1_clus2** for node 1, and **node2_clus1** and **node2_clus2** for node 2.
- The **cluster1::>** prompt indicates the name of the cluster.

During this procedure, refer to the following example:



About this task

The procedure requires the use of both ONTAP commands and [Nexus 9000 Series Switches](#) commands; ONTAP commands are used, unless otherwise indicated.

This procedure covers the following scenario:

- Switch cs2 is replaced by switch cs2-new first.
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - All cluster LIFs fail over to the new switch cs2-new.
 - Cabling between the nodes and cs2 are then disconnected from cs2 and reconnected to cs2-new.
- Switch cs1 is replaced by switch cs1-new.
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - All cluster LIFs fail over to the new switch cs1-new.
 - Cabling between the nodes and cs1 are then disconnected from cs1 and reconnected to cs1-new.



No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure fails over all of the cluster LIFs to the operational partner switch while performing the steps on the target switch.

Step 1: Prepare for migration

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`

where *x* is the duration of the maintenance window in hours.



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

2. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (***>**) appears.

Step 2: Configure ports and cabling

1. On the new switches, confirm that the ISL is cabled and healthy between the switches cs1-new and cs2-new:

```
show port-channel summary
```

Show example

```
cs1-new# 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      Eth1/35(P)  Eth1/36(P)

cs2-new# 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      Eth1/35(P)  Eth1/36(P)
```

2. Display the cluster ports on each node that are connected to the existing cluster switches:

```
network device-discovery show
```


Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
node1         /cdp
              e0a    cs1                       Ethernet1/1      N5K-
C5596UP
              e0b    cs2                       Ethernet1/2      N5K-
C5596UP
node2         /cdp
              e0a    cs1                       Ethernet1/1      N5K-
C5596UP
              e0b    cs2                       Ethernet1/2      N5K-
C5596UP
```

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

a. Verify that all the cluster ports are up with a healthy status:

```
network port show -ipspace Cluster
```

Show example

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

Node: node1

Ignore

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

Node: node2

Ignore

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

- b. Verify that all the cluster interfaces (LIFs) are on their home ports:

```
network interface show -vserver Cluster
```

Show example

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

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			
Cluster				
e0a	node1_clus1	up/up	169.254.209.69/16	node1
e0b	true			
e0a	node1_clus2	up/up	169.254.49.125/16	node1
e0b	true			
e0a	node2_clus1	up/up	169.254.47.194/16	node2
e0b	true			
e0a	node2_clus2	up/up	169.254.19.183/16	node2
e0b	true			

c. Verify that the cluster displays information for both cluster switches:

```
system cluster-switch show -is-monitoring-enabled-operational true
```

Show example

```
cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                                     Type                                     Address
Model
-----
cs1                                         cluster-network                         10.233.205.92
N5K-C5596UP
  Serial Number: FOXXXXXXXXGS
  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
N5K-C5596UP
  Serial Number: FOXXXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                                     9.3(4)
  Version Source: CDP
```

4. Disable auto-revert on the cluster LIFs.

By disabling auto-revert for this procedure, the cluster LIFs will not automatically move back to their home port. They remain on the current port while it continues to be up and operational.

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



Disabling auto-revert ensures ONTAP only fails over the cluster LIFs when the switch ports are shutdown later.

5. On cluster switch cs2, shut down the ports connected to the cluster ports of **all** the nodes in order to fail over the cluster LIFs:

```

cs2# configure
cs2(config)# interface eth1/1-1/2
cs2(config-if-range)# shutdown
cs2(config-if-range)# exit
cs2(config)# exit
cs2#

```

- Verify that the cluster LIFs have failed over to the ports hosted on cluster switch cs1. This might take a few seconds.

```
network interface show -vserver Cluster
```

Show example

```

cluster1::*> network interface show -vserver Cluster

```

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

Cluster				
e0a	true	up/up	169.254.3.4/16	node1
e0a	false	up/up	169.254.3.5/16	node1
e0a	true	up/up	169.254.3.8/16	node2
e0a	false	up/up	169.254.3.9/16	node2

- Verify that the cluster is healthy:

```
cluster show
```

Show example

```

cluster1::*> cluster show

```

Node	Health	Eligibility	Epsilon
node1	true	true	false
node2	true	true	false

8. If the cluster LIFs have failed over to switch cs1 and the cluster is healthy, go to [Step 10](#). If some cluster LIFs are not healthy or the cluster is unhealthy, you can roll back the connectivity to the switch cs2, as follows:
 - a. Bring up the ports connected to the cluster ports of **all** the nodes:

```
cs2# configure
cs2(config)# interface eth1/1-1/2
cs2(config-if-range)# no shutdown
cs2(config-if-range)# exit
cs2(config)# exit
cs2#
```

- b. Verify that the cluster LIFs have failed over to the ports hosted on cluster switch cs1. This might take a few seconds.

```
network interface show -vserver Cluster
```

Show example

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

Current	Is	Logical	Status	Network	Current
Vserver	Port	Interface	Admin/Oper	Address/Mask	Node
Home					
Cluster		node1_clus1	up/up	169.254.3.4/16	node1
e0a	true	node1_clus2	up/up	169.254.3.5/16	node1
e0a	false	node2_clus1	up/up	169.254.3.8/16	node2
e0a	true	node2_clus2	up/up	169.254.3.9/16	node2
e0a	false				

- c. Verify that the cluster is healthy:


```
cluster show
```

Show example

```
cluster1::*> cluster show
Node      Health  Eligibility  Epsilon
-----  -
node1     true    true         false
node2     true    true         false
```

- 9. Once you have restored LIF and cluster health, restart the process from [Step. 4](#).
- 10. Move all cluster node connection cables from the old cs2 switch to the new cs2-new switch.

Cluster node connection cables moved to the cs2-new switch



- 11. Confirm the health of the network connections moved to cs2-new:

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

Show example

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

```
Node: node1
```

```
Ignore
```

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

```
-----  
-----  
e0a      Cluster      Cluster      up    9000  auto/10000  
healthy  false  
e0b      Cluster      Cluster      up    9000  auto/10000  
healthy  false
```

```
Node: node2
```

```
Ignore
```

						Speed(Mbps)	Health
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
```

All cluster ports that were moved should be up.

12. Check neighbor information on the cluster ports:

```
network device-discovery show -protocol cdp
```


Show example

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

Node/	Local	Discovered		
Protocol	Port	Device (LLDP: ChassisID)	Interface	Platform

node1	/cdp			
	e0a	cs1	Ethernet1/1	N5K-
C5596UP				
	e0b	cs2-new	Ethernet1/1/1	N9K-
C9336C-FX2				
node2	/cdp			
	e0a	cs1	Ethernet1/2	N5K-
C5596UP				
	e0b	cs2-new	Ethernet1/1/2	N9K-
C9336C-FX2				

Verify that the moved cluster ports see the cs2-new switch as the neighbor.

13. Confirm the switch port connections from switch cs2-new's perspective:

```
cs2-new# show interface brief
cs2-new# show cdp neighbors
```

14. On cluster switch cs1, shut down the ports connected to the cluster ports of **all** the nodes in order to fail over the cluster LIFs.

```
cs1# configure
cs1(config)# interface eth1/1-1/2
cs1(config-if-range)# shutdown
cs1(config-if-range)# exit
cs1(config)# exit
cs1#
```

All cluster LIFs fail over to the cs2-new switch.

15. Verify that the cluster LIFs have failed over to the ports hosted on switch cs2-new. This might take a few seconds:

```
network interface show -vserver Cluster
```

Show example

```
cluster1::*> network interface show -vserver Cluster
          Logical      Status      Network      Current
Current Is
Vserver   Interfac   Admin/Oper  Address/Mask  Node
Port      Home
-----
Cluster
          node1_clus1  up/up      169.254.3.4/16  node1
e0b       false
          node1_clus2  up/up      169.254.3.5/16  node1
e0b       true
          node2_clus1  up/up      169.254.3.8/16  node2
e0b       false
          node2_clus2  up/up      169.254.3.9/16  node2
e0b       true
```

16. Verify that the cluster is healthy:


```
cluster show
```

Show example

```
cluster1::*> cluster show
Node      Health  Eligibility  Epsilon
-----
node1     true    true         false
node2     true    true         false
```

17. Move the cluster node connection cables from cs1 to the new cs1-new switch.

Cluster node connection cables moved to the cs1-new switch



18. Confirm the health of the network connections moved to cs1-new:

```
network port show -ipspace Cluster
```

Show example

```

cluster1::*> network port show -ipspace Cluster

Node: node1

Ignore

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

Node: node2

Ignore

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
  
```

All cluster ports that were moved should be up.

19. Check neighbor information on the cluster ports:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
node1         /cdp
              e0a    cs1-new                   Ethernet1/1/1  N9K-
C9336C-FX2
              e0b    cs2-new                   Ethernet1/1/2  N9K-
C9336C-FX2
node2         /cdp
              e0a    cs1-new                   Ethernet1/1/1  N9K-
C9336C-FX2
              e0b    cs2-new                   Ethernet1/1/2  N9K-
C9336C-FX2
```

Verify that the moved cluster ports see the cs1-new switch as the neighbor.

20. Confirm the switch port connections from switch cs1-new's perspective:

```
cs1-new# show interface brief
cs1-new# show cdp neighbors
```

21. Verify that the ISL between cs1-new and cs2-new is still operational:

```
show port-channel summary
```

Show example

```
cs1-new# 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      Eth1/35(P)  Eth1/36(P)

cs2-new# 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      Eth1/35(P)  Eth1/36(P)
```

Step 3: Verify the configuration

1. Enable auto-revert on the cluster LIFs.

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

2. Verify that the cluster LIFs have reverted to their home ports (this might take a minute):

```
network interface show -vserver Cluster
```

If the cluster LIFs have not reverted to their home port, manually revert them:

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

3. Verify that the cluster is healthy:

```
cluster show
```

4. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

Packet	Source	Destination
Node	Date	LIF
Loss		LIF
node1	3/5/2022 19:21:18 -06:00	node1_clus2
node1	3/5/2022 19:21:20 -06:00	node2_clus1
node2	3/5/2022 19:21:18 -06:00	node1_clus1
node2	3/5/2022 19:21:20 -06:00	node1_clus2

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node node2
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
    Local 169.254.19.183 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)

```

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

What's next?

[Configure switch health monitoring.](#)

Migrate to two-node switched cluster

If you have an existing two-node *switchless* cluster environment, you can migrate to a two-node *switched* cluster environment using Cisco Nexus 9336C-FX2 switches.

The migration process works for all nodes using optical or Twinax ports, but is not supported on this switch if nodes are using onboard 10Gb BASE-T RJ45 ports for the cluster-network ports.

Review requirements

What you'll need

- For the two-node switchless configuration:
 - The two-node switchless configuration is properly set up and functioning.
 - All cluster ports are in the **up** state.

- All cluster logical interfaces (LIFs) are in the **up** state and on their home ports.
- See [Hardware Universe](#) for all supported ONTAP versions.
- For the Cisco Nexus 9336C-FX2 switch configuration:
 - Both switches have management network connectivity.
 - There is console access to the cluster switches.
 - Nexus 9336C-FX2 node-to-node switch and switch-to-switch connections use Twinax or fiber cables.

See [Hardware Universe](#) for more information about cabling.

- Inter-Switch Link (ISL) cables are connected to ports 1/35 and 1/36 on both 9336C-FX2 switches.
- Initial customization of both the 9336C-FX2 switches are completed, so that:
 - 9336C-FX2 switches are running the latest version of software.
 - Reference Configuration Files (RCFs) are applied to the switches.
 Any site customization, such as SMTP, SNMP, and SSH, is configured on the new switches.

About the examples

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

- The names of the 9336C-FX2 switches are cs1 and cs2.
- The names of the cluster SVMs are node1 and node2.
- The names of the LIFs are node1_clus1 and node1_clus2 on node 1, and node2_clus1 and node2_clus2 on node 2 respectively.
- The `cluster1::*>` prompt indicates the name of the cluster.
- The cluster ports used in this procedure are e0a and e0b.

See [Hardware Universe](#) for information about the cluster ports for your platforms.

Migrate the switches

Step 1: Prepare for migration

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
```

where x is the duration of the maintenance window in hours.



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

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

```
set -privilege advanced
```

The advanced prompt (`*>`) appears.

Step 2: Configure ports and cabling

1. Disable all node-facing ports (not ISL ports) on both the new cluster switches cs1 and cs2.

Do not disable the ISL ports.

Show example

The following example shows that node-facing ports 1 through 34 are disabled on switch cs1:

```
cs1# config
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# interface e1/1/1-4, e1/2/1-4, e1/3/1-4, e1/4/1-4,
e1/5/1-4, e1/6/1-4, e1/7-34
cs1(config-if-range)# shutdown
```

2. Verify that the ISL and the physical ports on the ISL between the two 9336C-FX2 switches cs1 and cs2 are up on ports 1/35 and 1/36:

```
show port-channel summary
```

Show example

The following example shows that the ISL ports are up on switch cs1:

```
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      Eth1/35 (P)  Eth1/36 (P)
```

The following example shows that the ISL ports are up on switch cs2:

```
(cs2)# 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      Eth1/35 (P)  Eth1/36 (P)
```

3. Display the list of neighboring devices:

```
show cdp neighbors
```

This command provides information about the devices that are connected to the system.

Show example

The following example lists the neighboring devices on switch cs1:

```
cs1# 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
cs2                 Eth1/35        175     R S I s     N9K-C9336C
Eth1/35
cs2                 Eth1/36        175     R S I s     N9K-C9336C
Eth1/36

Total entries displayed: 2
```

The following example lists the neighboring devices on switch cs2:

```
cs2# 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
cs1                 Eth1/35        177     R S I s     N9K-C9336C
Eth1/35
cs1                 Eth1/36        177     R S I s     N9K-C9336C
Eth1/36

Total entries displayed: 2
```

4. Verify that all cluster ports are up:

```
network port show -ipspace Cluster
```

Each port should display up for Link and healthy for Health Status.

Show example

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

Node: node1

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

Node: node2

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

4 entries were displayed.
```

5. Verify that all cluster LIFs are up and operational:

```
network interface show -vserver Cluster
```

Each cluster LIF should display true for Is Home and have a Status Admin/Oper of up/up.

Show example

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

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

Cluster				
e0a	node1_clus1	up/up	169.254.209.69/16	node1
	true			
e0b	node1_clus2	up/up	169.254.49.125/16	node1
	true			
e0a	node2_clus1	up/up	169.254.47.194/16	node2
	true			
e0b	node2_clus2	up/up	169.254.19.183/16	node2
	true			

4 entries were displayed.

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

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

Show example

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

Vserver	Logical	Auto-revert
Interface		

Cluster		
	node1_clus1	true
	node1_clus2	true
	node2_clus1	true
	node2_clus2	true

4 entries were displayed.

7. Disconnect the cable from cluster port e0a on node1, and then connect e0a to port 1 on cluster switch cs1, using the appropriate cabling supported by the 9336C-FX2 switches.

The [Hardware Universe - Switches](#) contains more information about cabling.

Hardware Universe - Switches

8. Disconnect the cable from cluster port e0a on node2, and then connect e0a to port 2 on cluster switch cs1, using the appropriate cabling supported by the 9336C-FX2 switches.
9. Enable all node-facing ports on cluster switch cs1.

Show example

The following example shows that ports 1/1 through 1/34 are enabled on switch cs1:

```
cs1# config
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# interface e1/1/1-4, e1/2/1-4, e1/3/1-4, e1/4/1-4,
e1/5/1-4, e1/6/1-4, e1/7-34
cs1(config-if-range)# no shutdown
```

10. Verify that all cluster LIFs are up, operational, and display as `true` for `Is Home`:

```
network interface show -vserver Cluster
```

Show example

The following example shows that all of the LIFs are up on node1 and node2 and that Is Home results are true:

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

Current Is Home	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Port
Cluster	node1_clus1	up/up	169.254.209.69/16	node1	e0a
true	node1_clus2	up/up	169.254.49.125/16	node1	e0b
true	node2_clus1	up/up	169.254.47.194/16	node2	e0a
true	node2_clus2	up/up	169.254.19.183/16	node2	e0b
true					

4 entries were displayed.

11. Display information about the status of the nodes in the cluster:

```
cluster show
```

Show example

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

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
node1	true	true	false
node2	true	true	false

2 entries were displayed.

12. Disconnect the cable from cluster port e0b on node1, and then connect e0b to port 1 on cluster switch cs2, using the appropriate cabling supported by the 9336C-FX2 switches.

13. Disconnect the cable from cluster port e0b on node2, and then connect e0b to port 2 on cluster switch cs2, using the appropriate cabling supported by the 9336C-FX2 switches.
14. Enable all node-facing ports on cluster switch cs2.

Show example

The following example shows that ports 1/1 through 1/34 are enabled on switch cs2:

```
cs2# config
Enter configuration commands, one per line. End with CNTL/Z.
cs2(config)# interface e1/1/1-4, e1/2/1-4, e1/3/1-4, e1/4/1-4,
e1/5/1-4, e1/6/1-4, e1/7-34
cs2(config-if-range)# no shutdown
```

15. Verify that all cluster ports are up:

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

Show example

The following example shows that all of the cluster ports are up on node1 and node2:

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

Node: node1

Ignore

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

Node: node2

Ignore

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

4 entries were displayed.
```

Step 3: Verify the configuration

1. Verify that all interfaces display true for Is Home:

```
network interface show -vserver Cluster
```



This might take several minutes to complete.

Show example

The following example shows that all LIFs are up on node1 and node2 and that Is Home results are true:

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

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	-----
-----	-----				
Cluster					
true	node1_clus1	up/up	169.254.209.69/16	node1	e0a
true	node1_clus2	up/up	169.254.49.125/16	node1	e0b
true	node2_clus1	up/up	169.254.47.194/16	node2	e0a
true	node2_clus2	up/up	169.254.19.183/16	node2	e0b
true					

4 entries were displayed.

2. Verify that both nodes each have one connection to each switch:

```
show cdp neighbors
```

Show example

The following example shows the appropriate results for both switches:

```
(cs1)# 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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0a	Eth1/1	133	H	FAS2980
node2 e0a	Eth1/2	133	H	FAS2980
cs2 Eth1/35	Eth1/35	175	R S I s	N9K-C9336C
cs2 Eth1/36	Eth1/36	175	R S I s	N9K-C9336C

```
Total entries displayed: 4
```

```
(cs2)# 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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0b	Eth1/1	133	H	FAS2980
node2 e0b	Eth1/2	133	H	FAS2980
cs1 Eth1/35	Eth1/35	175	R S I s	N9K-C9336C
cs1 Eth1/36	Eth1/36	175	R S I s	N9K-C9336C

```
Total entries displayed: 4
```

3. Display information about the discovered network devices in your cluster:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol       Port   Device (LLDP: ChassisID)  Interface
Platform
-----
node2          /cdp
              e0a   cs1                       0/2          N9K-
C9336C
              e0b   cs2                       0/2          N9K-
C9336C
node1          /cdp
              e0a   cs1                       0/1          N9K-
C9336C
              e0b   cs2                       0/1          N9K-
C9336C

4 entries were displayed.
```

4. Verify that the settings are disabled:

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



It might take several minutes for the command to complete. Wait for the '3 minute lifetime to expire' announcement.

Show example

The false output in the following example shows that the configuration settings are disabled:

```
cluster1::*> network options switchless-cluster show
Enable Switchless Cluster: false
```

5. Verify the status of the node members in the cluster:

```
cluster show
```

Show example

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

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
node1	true	true	false
node2	true	true	false

6. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

7. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

What's next?

[Configure switch health monitoring.](#)

Replace switches

Replace a Cisco Nexus 9336C-FX2 cluster switch

Follow these steps to replace a defective Nexus 9336C-FX2 switch in a cluster network. This is a nondisruptive procedure (NDU).

Review requirements

Before performing the switch replacement, make sure that:

- You have verified the switch serial number to ensure that the correct switch is replaced.
- On the existing cluster and network infrastructure:
 - The existing cluster is verified as completely functional, with at least one fully connected cluster switch.
 - All cluster ports are **up**.
 - All cluster logical interfaces (LIFs) are **up** and on their home ports.
 - The ONTAP `cluster ping-cluster -node node1` command must indicate that basic connectivity and larger than PMTU communication are successful on all paths.
- On the Nexus 9336C-FX2 replacement switch:
 - Management network connectivity on the replacement switch is functional.
 - Console access to the replacement switch is in place.
 - The node connections are ports 1/1 through 1/34.
 - All Inter-Switch Link (ISL) ports is disabled on ports 1/35 and 1/36.
 - The desired reference configuration file (RCF) and NX-OS operating system image switch is loaded onto the switch.
 - Initial customization of the switch is complete, as detailed in [Configure the 9336C-FX2 cluster switch](#).

Any previous site customizations, such as STP, SNMP, and SSH, are copied to the new switch.
- You have executed the command for migrating a cluster LIF from the node where the cluster LIF is hosted.

Enable console logging

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

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

Replace the switch

About the examples

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

- The names of the existing Nexus 9336C-FX2 switches are cs1 and cs2.
- The name of the new Nexus 9336C-FX2 switch is newcs2.
- The node names are node1 and node2.
- The cluster ports on each node are named e0a and e0b.
- The cluster LIF names are node1_clus1 and node1_clus2 for node1, and node2_clus1 and node2_clus2 for node2.
- The prompt for changes to all cluster nodes is cluster1::<*>

About this task

The following procedure is based on the following cluster network topology:

Show example

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

```
Node: node1
```

```
Ignore
```

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

e0a	Cluster	Cluster		up	9000	auto/10000	healthy
false							
e0b	Cluster	Cluster		up	9000	auto/10000	healthy
false							

```
Node: node2
```

```
Ignore
```

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

e0a	Cluster	Cluster		up	9000	auto/10000	healthy
false							
e0b	Cluster	Cluster		up	9000	auto/10000	healthy
false							

```
4 entries were displayed.
```

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

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

Cluster					
	node1_clus1	up/up	169.254.209.69/16	node1	e0a
true					
	node1_clus2	up/up	169.254.49.125/16	node1	e0b

```

true
      node2_clus1  up/up    169.254.47.194/16  node2      e0a
true
      node2_clus2  up/up    169.254.19.183/16  node2      e0b
true
4 entries were displayed.

```

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

Node/ Protocol	Local Port	Discovered Device (LLDP: ChassisID)	Interface	Platform
node2 C9336C	/cdp e0a	cs1	Eth1/2	N9K-
C9336C	e0b	cs2	Eth1/2	N9K-
node1 C9336C	/cdp e0a	cs1	Eth1/1	N9K-
C9336C	e0b	cs2	Eth1/1	N9K-

```
4 entries were displayed.
```

```
cs1# 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 ID	Local Intrfce	Hldtme	Capability	Platform	Port
node1	Eth1/1	144	H	FAS2980	e0a
node2	Eth1/2	145	H	FAS2980	e0a
cs2 Eth1/35	Eth1/35	176	R S I s	N9K-C9336C	
cs2 (FDO220329V5) Eth1/36	Eth1/36	176	R S I s	N9K-C9336C	

```
Total entries displayed: 4
```

```
cs2# 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
node1	Eth1/1	139	H	FAS2980	e0b
node2	Eth1/2	124	H	FAS2980	e0b
cs1	Eth1/35	178	R S I s	N9K-C9336C	
Eth1/35					
cs1	Eth1/36	178	R S I s	N9K-C9336C	
Eth1/36					

```
Total entries displayed: 4
```

Step 1: Prepare for replacement

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

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

where x is the duration of the maintenance window in hours.



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

2. Install the appropriate RCF and image on the switch, newcs2, and make any necessary site preparations.

If necessary, verify, download, and install the appropriate versions of the RCF and NX-OS software for the new switch. If you have verified that the new switch is correctly set up and does not need updates to the RCF and NX-OS software, continue to step 2.

- a. Go to the *NetApp Cluster and Management Network Switches Reference Configuration File Description Page* on the NetApp Support Site.
 - b. Click the link for the *Cluster Network and Management Network Compatibility Matrix*, and then note the required switch software version.
 - c. Click your browser's back arrow to return to the Description page, click **CONTINUE**, accept the license agreement, and then go to the Download page.
 - d. Follow the steps on the Download page to download the correct RCF and NX-OS files for the version of ONTAP software you are installing.
3. On the new switch, log in as admin and shut down all of the ports that will be connected to the node cluster interfaces (ports 1/1 to 1/34).

If the switch that you are replacing is not functional and is powered down, go to Step 4. The LIFs on the

cluster nodes should have already failed over to the other cluster port for each node.

Show example

```
newcs2# config
Enter configuration commands, one per line. End with CNTL/Z.
newcs2(config)# interface e1/1-34
newcs2(config-if-range)# shutdown
```

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

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

Show example

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

Vserver	Logical Interface	Auto-revert
Cluster	node1_clus1	true
Cluster	node1_clus2	true
Cluster	node2_clus1	true
Cluster	node2_clus2	true

4 entries were displayed.

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

Step 2: Configure cables and ports

1. Shut down the ISL ports 1/35 and 1/36 on the Nexus 9336C-FX2 switch cs1.

Show example

```

cs1# configure
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# interface e1/35-36
cs1(config-if-range)# shutdown
cs1(config-if-range)#

```

2. Remove all of the cables from the Nexus 9336C-FX2 cs2 switch, and then connect them to the same ports on the Nexus C9336C-FX2 newcs2 switch.
3. Bring up the ISLs ports 1/35 and 1/36 between the cs1 and newcs2 switches, and then verify the port channel operation status.

Port-Channel should indicate Po1(SU) and Member Ports should indicate Eth1/35(P) and Eth1/36(P).

Show example

This example enables ISL ports 1/35 and 1/36 and displays the port channel summary on switch cs1:

```
cs1# configure
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# int e1/35-36
cs1(config-if-range)# no shutdown

cs1(config-if-range)# 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       Eth1/35(P)  Eth1/36(P)

cs1(config-if-range)#
```

4. Verify that port e0b is up on all nodes:

```
network port show ipspace Cluster
```

Show example

The output should be similar to the following:

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

Node: node1

Ignore

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

Node: node2

Ignore

Health      Health      Speed (Mbps)
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/auto  -
false

4 entries were displayed.
```

5. On the same node you used in the previous step, revert the cluster LIF associated with the port in the previous step by using the network interface revert command.

Show example

In this example, LIF node1_clus2 on node1 is successfully reverted if the Home value is true and the port is e0b.

The following commands return LIF node1_clus2 on node1 to home port e0a and displays information about the LIFs on both nodes. Bringing up the first node is successful if the Is Home column is true for both cluster interfaces and they show the correct port assignments, in this example e0a and e0b on node1.

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

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

Cluster				
e0a	node1_clus1	up/up	169.254.209.69/16	node1
e0b	node1_clus2	up/up	169.254.49.125/16	node1
e0a	node2_clus1	up/up	169.254.47.194/16	node2
e0a	node2_clus2	up/up	169.254.19.183/16	node2
e0a	false			

4 entries were displayed.

6. Display information about the nodes in a cluster:

```
cluster show
```

Show example

This example shows that the node health for node1 and node2 in this cluster is true:

```
cluster1::*> cluster show
```

Node	Health	Eligibility

node1	false	true
node2	true	true

7. Verify that all physical cluster ports are up:

```
network port show ipspace Cluster
```

Show example

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

Node node1
Ignore

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

-----
e0a         Cluster   Cluster          up   9000 auto/10000
healthy    false
e0b         Cluster   Cluster          up   9000 auto/10000
healthy    false

Node: node2

Ignore

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

-----
e0a         Cluster   Cluster          up   9000 auto/10000
healthy    false
e0b         Cluster   Cluster          up   9000 auto/10000
healthy    false

4 entries were displayed.
```

8. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

Step 3: Verify the configuration

1. Confirm the following cluster network configuration:

```
network port show
```

Show example

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

```
Node: node1
```

```
Ignore
```

				Speed (Mbps)		Health	
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: node2
```

```
Ignore
```

				Speed (Mbps)		Health	
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.
```

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

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

	Cluster				
	e0a	node1_clus1	up/up	169.254.209.69/16	node1
		node1_clus2	up/up	169.254.49.125/16	node1

```

e0b      true
          node2_clus1  up/up    169.254.47.194/16  node2
e0a      true
          node2_clus2  up/up    169.254.19.183/16  node2
e0b      true

```

4 entries were displayed.

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

```

Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface
Platform
-----
node2      /cdp
          e0a    cs1                        0/2          N9K-
C9336C
          e0b    newcs2                    0/2          N9K-
C9336C
node1      /cdp
          e0a    cs1                        0/1          N9K-
C9336C
          e0b    newcs2                    0/1          N9K-
C9336C

```

4 entries were displayed.

```
cs1# 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
node1          Eth1/1        144     H           FAS2980
e0a
node2          Eth1/2        145     H           FAS2980
e0a
newcs2        Eth1/35       176     R S I s     N9K-C9336C
Eth1/35
newcs2        Eth1/36       176     R S I s     N9K-C9336C

```



```
Eth1/36
```

```
Total entries displayed: 4
```

```
cs2# 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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0b	Eth1/1	139	H	FAS2980
node2 e0b	Eth1/2	124	H	FAS2980
cs1 Eth1/35	Eth1/35	178	R S I s	N9K-C9336C
cs1 Eth1/36	Eth1/36	178	R S I s	N9K-C9336C

```
Total entries displayed: 4
```

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

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

What's next?

[Configure switch health monitoring.](#)

Replace Cisco Nexus 9336C-FX2 cluster switches with switchless connections

You can migrate from a cluster with a switched cluster network to one where two nodes are directly connected for ONTAP 9.3 and later.

Review requirements

Guidelines

Review the following guidelines:

- Migrating to a two-node switchless cluster configuration is a nondisruptive operation. Most systems have two dedicated cluster interconnect ports on each node, but you can also use this procedure for systems with a larger number of dedicated cluster interconnect ports on each node, such as four, six or eight.

- You cannot use the switchless cluster interconnect feature with more than two nodes.
- If you have an existing two-node cluster that uses cluster interconnect switches and is running ONTAP 9.3 or later, you can replace the switches with direct, back-to-back connections between the nodes.


What you'll need

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

Migrate the switches

About this task

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



About the examples

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

Step 1: Prepare for migration

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

```
set -privilege advanced
```

The advanced prompt `*>` appears.

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

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

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

Show example

The following example output shows if the option is enabled.

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

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

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

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

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

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

Show example


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

Step 2: Configure ports and cabling

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

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

In the following example for nodes with cluster ports "e0a" and "e0b", one group is identified as "node1:e0a" and "node2:e0a" and the other group as "node1:e0b" and "node2:e0b". Your nodes might be using different cluster ports because they vary by system.



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

Show example

```

cluster::> network port show -ipspace Cluster
Node: node1

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

Node: node2

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

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

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

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

Show example

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

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

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

4. Disable auto-revert for the cluster LIFs:

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

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

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

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

Show example

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

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

6. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

7. Verify that the cluster is healthy:

```
cluster ring show
```

All units must be either master or secondary.


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



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


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

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



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

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



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

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

The following example shows that the switchless cluster is enabled:

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

10. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```



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


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



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


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

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



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

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



Step 3: Verify the configuration

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

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

Show example

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

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

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

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

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

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

Show example

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

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

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

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

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

```
cluster show
```

Show example

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

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

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

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

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

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

7. Change the privilege level back to admin:

```
set -privilege admin
```

NVIDIA SN2100

Overview

Overview of installation and configuration for NVIDIA SN2100 switches

The NVIDIA SN2100 is a cluster switch that allows you to build ONTAP clusters with more than two nodes.

Initial configuration overview

To configure a NVIDIA SN2100 switch on systems running ONTAP, follow these steps:

1. [Install the hardware for the NVIDIA SN2100 switch.](#)

Instructions are available in the *NVIDIA Switch Installation Guide*.

2. [Configure the switch.](#)

Instructions are available in NVIDIA's documentation.

3. [Review cabling and configuration considerations.](#)

Review requirements for optical connections, the QSA adapter, and the switchport speed.

4. [Cable the NS224 shelves as switch-attached storage.](#)

Follow the cabling procedures if you have a system in which the NS224 drive shelves need to be cabled as switch-attached storage (not direct-attached storage).

5. [Install Cumulus Linux in Cumulus mode](#) or [install Cumulus Linux in ONIE mode](#).

You can install Cumulus Linux (CL) OS when the switch is running either Cumulus Linux or ONIE.

6. [Install the Reference Configuration File \(RCF\) script.](#)

There are two RCF scripts available for Clustering and Storage applications. The procedure for each is the same.

7. [Install the CSHM file.](#)

You can install the applicable configuration file for Ethernet switch health monitoring of NVIDIA cluster switches.

Additional information

Before you begin installation or maintenance, be sure to review the following:

- [Configuration requirements](#)
- [Components and part numbers](#)
- [Required documentation](#)
- [Hardware Universe](#) for all supported ONTAP versions.

Configuration requirements for NVIDIA SN2100 switches

For NVIDIA SN2100 switch installation and maintenance, be sure to review all configuration requirements.

Installation requirements

If you want to build ONTAP clusters with more than two nodes, you need two supported cluster network switches. You can use additional management switches, which are optional.

You install the NVIDIA SN2100 switch (X190006) in the NVIDIA dual/single switch cabinet with the standard brackets that are included with the switch.

For cabling guidelines, see [Review cabling and configuration considerations](#).

ONTAP and Linux support

The NVIDIA SN2100 switch is a 10/25/40/100GbE switch running Cumulus Linux. The switch supports the following:

- ONTAP 9.10.1P3 and later

The SN2100 switch serves Cluster and Storage applications in ONTAP 9.10.1P3 and later over different switch-pairs.

- Cumulus Linux (CL) OS versions
 - Specific CL versions are qualified and supported by NetApp. For current compatibility information, see the [NVIDIA Ethernet Switches information](#) page or the [NetApp Hardware Universe](#).
 - In order to download the SN2100 Cumulus software from NVIDIA, you must have login credentials to access NVIDIA's Enterprise Support Portal. See the Knowledge Base article [How to register with NVIDIA for Enterprise Support Portal Access](#).
- You can install Cumulus Linux when the switch is running Cumulus Linux or ONIE.

Components and part numbers for NVIDIA SN2100 switches

For NVIDIA SN2100 switch installation and maintenance, be sure to review the list of components and part numbers for the cabinet and rail kit.

Cabinet details

You install the NVIDIA SN2100 switch (X190006) in the NVIDIA dual/single switch cabinet with the standard brackets that are included with the switch.

Rail kit details

The following table lists the part number and description for the SN2100 switches and rail kits:

Part number	Description
X190006-PE	Cluster Switch, NVIDIA SN2100, 16PT 100GbE, PTSX
X190006-PI	Cluster Switch, NVIDIA SN2100, 16PT 100GbE, PSIN
X-MTEF-KIT-D	Rail Kit, NVIDIA Dual switch side by side
X-MTEF-KIT-E	Rail Kit, NVIDIA Single switch short depth



See NVIDIA documentation for details on [installing your SN2100 switch and rail kit](#).

Documentation requirements for NVIDIA SN2100 switches

For NVIDIA SN2100 switch installation and maintenance, be sure to review all the recommended documentation.

Title	Description
NVIDIA Switch Installation Guide	Describes how to install your NVIDIA SN2100 switches.
NS224 NVMe Drive Shelf Cabling Guide	Overview and illustrations showing how to configure cabling for drive shelves.
NetApp Hardware Universe	Allows you to confirm supported hardware, such as storage switches and cables, for your platform model.

Install hardware

Install the hardware for the NVIDIA SN2100 switch

To install the SN2100 hardware, refer to NVIDIA's documentation.

Steps

1. Review the [configuration requirements](#).
2. Follow the instructions in [NVIDIA Switch Installation Guide](#).

What's next?

[Configure the switch](#).

Configure the NVIDIA SN2100 switch

To configure the SN2100 switch, refer to NVIDIA's documentation.

Steps

1. Review the [configuration requirements](#).
2. Follow the instructions in [NVIDIA System Bring-Up..](#)

What's next?

[Review cabling and configuration considerations](#).

Review cabling and configuration considerations

Before configuring your NVIDIA SN2100 switch, review the following considerations.

NVIDIA port details

Switch ports	Ports usage
swp1s0-3	4x10GbE breakout cluster port nodes

swp2s0-3	4x25GbE breakout cluster port nodes
swp3-14	40/100GbE cluster port nodes
swp15-16	100GbE Inter-Switch Link (ISL) ports

See the [Hardware Universe](#) for more information on switch ports.

Link-up delays with optical connections

If you are experiencing link-up delays of more than five seconds, Cumulus Linux 5.4 and later includes support for fast link-up. You can configure the links by using the `nv set` command as follows:

```
nv set interface <interface-id> link fast-linkup on  
nv config apply  
reload the switchd
```

Show example

```
cumulus@cumulus-cs13:mgmt:~$ nv set interface swp5 link fast-linkup on  
cumulus@cumulus-cs13:mgmt:~$ nv config apply  
switchd need to reload on this config change  
  
Are you sure? [y/N] y  
applied [rev_id: 22]  
  
Only switchd reload required
```

Support for copper connections

The following configuration changes are required to fix this issue.

Cumulus Linux 4.4.3

1. Identify the name for each interface using 40GbE/100GbE copper cables:

```
cumulus@cumulus:mgmt:~$ net show interface pluggables
```

Interface	Identifier	Vendor Name	Vendor PN	Vendor SN
Vendor Rev				
-----	-----	-----	-----	-----

swp3	0x11 (QSFP28)	Molex	112-00576	93A2229911111
B0				
swp4	0x11 (QSFP28)	Molex	112-00576	93A2229922222
B0				

2. Add the following two lines to the `/etc/cumulus/switchd.conf` file for every port (swp<n>) that is using 40GbE/100GbE copper cables:

- `interface.swp<n>.enable_media_depended_linkup_flow=TRUE`
- `interface.swp<n>.enable_short_tuning=TRUE`

For example:

```
cumulus@cumulus:mgmt:~$ sudo nano /etc/cumulus/switchd.conf
.
.
interface.swp3.enable_media_depended_linkup_flow=TRUE
interface.swp3.enable_short_tuning=TRUE
interface.swp4.enable_media_depended_linkup_flow=TRUE
interface.swp4.enable_short_tuning=TRUE
```

3. Restart the `switchd` service:

```
cumulus@cumulus:mgmt:~$ sudo systemctl restart switchd.service
```

4. Confirm that the ports are up:

```
cumulus@cumulus:mgmt:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp3	100G	9216	Trunk/L2		Master: bridge (UP)
UP	swp4	100G	9216	Trunk/L2		Master: bridge (UP)

Cumulus Linux 5.x

1. Identify the name for each interface using 40GbE/100GbE copper cables:

```
cumulus@cumulus:mgmt:~$ nv show interface --view=pluggables
```

Interface	Identifier	Vendor Name	Vendor PN	Vendor SN
swp3	0x11 (QSFP28)	Molex	112-00576	93A2229911111
B0				
swp4	0x11 (QSFP28)	Molex	112-00576	93A2229922222
B0				

2. Configure the links using the `nv set` command as follows:

- `nv set interface <interface-id> link fast-linkup on`
- `nv config apply`
- Reload the `switchd` service

For example:

```
cumulus@cumulus:mgmt:~$ nv set interface swp5 link fast-linkup on
cumulus@cumulus:mgmt:~$ nv config apply
switchd need to reload on this config change
```

```
Are you sure? [y/N] y
applied [rev_id: 22]
```

```
Only switchd reload required
```

3. Confirm that the ports are up:

```
cumulus@cumulus:mgmt:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp3	100G	9216	Trunk/L2		Master: bridge (UP)
UP	swp4	100G	9216	Trunk/L2		Master: bridge (UP)

See the Knowledge Base article [SN2100 switch fails to connect using 40/100GbE copper cables](#) for further details.

On Cumulus Linux 4.4.2, copper connections are not supported on SN2100 switches with X1151A NIC, X1146A NIC, or onboard 100GbE ports.

For example:

- AFF A800 on ports e0a and e0b
- AFF A320 on ports e0g and e0h

QSA adapter

When a QSA adapter is used to connect to the 10GbE/25GbE cluster ports on a platform, the link might not come up.

To resolve this issue, do the following:

- For 10GbE, manually set the swp1s0-3 link speed to 10000 and set auto-negotiation to off.
- For 25GbE, manually set the swp2s0-3 link speed to 25000 and set auto-negotiation to off.



When using 10GbE/25GbE QSA adapters, insert them in non-breakout 40GbE/100GbE ports (swp3-swp14). Do not insert the QSA adapter in a port that is configured for breakout.

Setting interface speed on breakout ports

Depending on the transceiver in the switch port, you might need to set the speed on the switch interface to a fixed speed. If using 10GbE and 25GbE breakout ports, verify that auto-negotiation is off and set the interface speed on the switch.

Cumulus Linux 4.4.3

For example:

```
cumulus@cumulus:mgmt:~$ net add int swpls3 link autoneg off && net com
--- /etc/network/interfaces      2019-11-17 00:17:13.470687027 +0000
+++ /run/nclu/ifupdown2/interfaces.tmp  2019-11-24 00:09:19.435226258
+0000
@@ -37,21 +37,21 @@
     alias 10G Intra-Cluster Node
     link-autoneg off
     link-speed 10000 <---- port speed set
     mstpctl-bpduguard yes
     mstpctl-portadminedge yes
     mtu 9216

auto swpls3
iface swpls3
    alias 10G Intra-Cluster Node
-   link-autoneg off
+   link-autoneg on
    link-speed 10000 <---- port speed set
    mstpctl-bpduguard yes
    mstpctl-portadminedge yes
    mtu 9216

auto swp2s0
iface swp2s0
    alias 25G Intra-Cluster Node
    link-autoneg off
    link-speed 25000 <---- port speed set
```

Check the interface and port status to verify that the settings are applied:


```
cumulus@cumulus:mgmt:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp1s0	10G	9216	Trunk/L2	cs07 (e4c)	Master: br_default(UP)
UP	swp1s1	10G	9216	Trunk/L2	cs07 (e4d)	Master: br_default(UP)
UP	swp1s2	10G	9216	Trunk/L2	cs08 (e4c)	Master: br_default(UP)
UP	swp1s3	10G	9216	Trunk/L2	cs08 (e4d)	Master: br_default(UP)
UP	swp3	40G	9216	Trunk/L2	cs03 (e4e)	Master: br_default(UP)
UP	swp4	40G	9216	Trunk/L2	cs04 (e4e)	Master: br_default(UP)
DN	swp5	N/A	9216	Trunk/L2		Master: br_default(UP)
DN	swp6	N/A	9216	Trunk/L2		Master: br_default(UP)
DN	swp7	N/A	9216	Trunk/L2		Master: br_default(UP)
UP	swp15	100G	9216	BondMember	cs01 (swp15)	Master: cluster_isl(UP)
UP	swp16	100G	9216	BondMember	cs01 (swp16)	Master: cluster_isl(UP)

Cumulus Linux 5.x

For example:

```
cumulus@cumulus:mgmt:~$ nv set interface swp1s3 link auto-negotiate off
cumulus@cumulus:mgmt:~$ nv set interface swp1s3 link speed 10G
cumulus@cumulus:mgmt:~$ nv show interface swp1s3

link

  auto-negotiate      off          off
off
  duplex              full         full
full
  speed               10G         10G
10G
  fec                 auto         auto
auto
  mtu                 9216        9216
9216
[breakout]

  state               up           up
up
```

Check the interface and port status to verify that the settings are applied:

```
cumulus@cumulus:mgmt:~$ nv show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp1s0	10G	9216	Trunk/L2	cs07 (e4c)	Master: br_default(UP)
UP	swp1s1	10G	9216	Trunk/L2	cs07 (e4d)	Master: br_default(UP)
UP	swp1s2	10G	9216	Trunk/L2	cs08 (e4c)	Master: br_default(UP)
UP	swp1s3	10G	9216	Trunk/L2	cs08 (e4d)	Master: br_default(UP)
UP	swp3	40G	9216	Trunk/L2	cs03 (e4e)	Master: br_default(UP)
UP	swp4	40G	9216	Trunk/L2	cs04 (e4e)	Master: br_default(UP)
DN	swp5	N/A	9216	Trunk/L2		Master: br_default(UP)
DN	swp6	N/A	9216	Trunk/L2		Master: br_default(UP)
DN	swp7	N/A	9216	Trunk/L2		Master: br_default(UP)
UP	swp15	100G	9216	BondMember	cs01 (swp15)	Master: cluster_isl(UP)
UP	swp16	100G	9216	BondMember	cs01 (swp16)	Master: cluster_isl(UP)

What's next?

[Cable NS224 shelves as switch-attached storage.](#)

Cable the NS224 shelves as switch-attached storage

If you have a system in which the NS224 drive shelves need to be cabled as switch-attached storage (not direct-attached storage), use the information provided here.

- Cable NS224 drive shelves through storage switches:

Cabling switch-attached NS224 drive shelves

- Confirm supported hardware, such as storage switches and cables, for your platform model:

[NetApp Hardware Universe](#)

What's next?

[Install Cumulus Linux in Cumulus mode](#) or [Install Cumulus Linux in ONIE mode](#).

Configure software

Software install workflow for NVIDIA SN2100 switches

To install and configure software for a NVIDIA SN2100 switch, follow these steps:

1. [Install Cumulus Linux in Cumulus mode](#) or [install Cumulus Linux in ONIE mode](#).

You can install Cumulus Linux (CL) OS when the switch is running either Cumulus Linux or ONIE.

2. [Install the Reference Configuration File \(RCF\) script](#).

There are two RCF scripts available for Clustering and Storage applications. The procedure for each is the same.

3. [Install the CSHM file](#).

You can install the applicable configuration file for Ethernet switch health monitoring of NVIDIA cluster switches.

Install Cumulus Linux in Cumulus mode

Follow this procedure to install Cumulus Linux (CL) OS when the switch is running in Cumulus mode.



Cumulus Linux (CL) OS can be installed either when the switch is running Cumulus Linux or ONIE (see [Install in ONIE mode](#)).

What you'll need

- Intermediate-level Linux knowledge.
- Familiarity with basic text editing, UNIX file permissions, and process monitoring. A variety of text editors are pre-installed, including `vi` and `nano`.
- Access to a Linux or UNIX shell. If you are running Windows, use a Linux environment as your command line tool for interacting with Cumulus Linux.
- The baud rate requirement is set to 115200 on the serial console switch for NVIDIA SN2100 switch console access, as follows:
 - 115200 baud
 - 8 data bits
 - 1 stop bit

- parity: none
- flow control: none

About this task

Be aware of the following:



Each time Cumulus Linux is installed, the entire file system structure is erased and rebuilt.



The default password for the cumulus user account is **cumulus**. The first time you log into Cumulus Linux, you must change this default password. Be sure to update any automation scripts before installing a new image. Cumulus Linux provides command line options to change the default password automatically during the installation process.

Example 1. Steps

Cumulus Linux 4.4.3

1. Log in to the switch.

First time log in to the switch requires username/password of **cumulus/cumulus** with `sudo` privileges.

```
cumulus login: cumulus
Password: cumulus
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>
```

2. Check the Cumulus Linux version: `net show system`

```
cumulus@cumulus:mgmt:~$ net show system
Hostname..... cumulus
Build..... Cumulus Linux 4.4.3
Uptime..... 0:08:20.860000
Model..... Mlnx X86
CPU..... x86_64 Intel Atom C2558 2.40GHz
Memory..... 8GB
Disk..... 14.7GB
ASIC..... Mellanox Spectrum MT52132
Ports..... 16 x 100G-QSFP28
Part Number..... MSN2100-CB2FC
Serial Number.... MT2105T05177
Platform Name.... x86_64-mlnx_x86-r0
Product Name..... MSN2100
ONIE Version..... 2019.11-5.2.0020-115200
Base MAC Address. 04:3F:72:43:92:80
Manufacturer..... Mellanox
```

3. Configure the hostname, IP address, subnet mask, and default gateway. The new hostname only becomes effective after restarting the console/SSH session.



A Cumulus Linux switch provides at least one dedicated Ethernet management port called `eth0`. This interface is specifically for out-of-band management use. By default, the management interface uses DHCPv4 for addressing.



Do not use an underscore (`_`), apostrophe (`'`), or non-ASCII characters in the hostname.

```
cumulus@cumulus:mgmt:~$ net add hostname sw1
cumulus@cumulus:mgmt:~$ net add interface eth0 ip address
10.233.204.71
cumulus@cumulus:mgmt:~$ net add interface eth0 ip gateway
10.233.204.1
cumulus@cumulus:mgmt:~$ net pending
cumulus@cumulus:mgmt:~$ net commit
```

This command modifies both the `/etc/hostname` and `/etc/hosts` files.

4. Confirm that the hostname, IP address, subnet mask, and default gateway have been updated.

```
cumulus@sw1:mgmt:~$ hostname sw1
cumulus@sw1:mgmt:~$ ifconfig eth0
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 10.233.204.71 netmask 255.255.254.0 broadcast 10.233.205.255
inet6 fe80::bace:f6ff:fe19:1df6 prefixlen 64 scopeid 0x20<link>
ether b8:ce:f6:19:1d:f6 txqueuelen 1000 (Ethernet)
RX packets 75364 bytes 23013528 (21.9 MiB)
RX errors 0 dropped 7 overruns 0 frame 0
TX packets 4053 bytes 827280 (807.8 KiB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0 device
memory 0xdfc00000-dfc1ffff

cumulus@sw1::mgmt:~$ ip route show vrf mgmt
default via 10.233.204.1 dev eth0
unreachable default metric 4278198272
10.233.204.0/23 dev eth0 proto kernel scope link src 10.233.204.71
127.0.0.0/8 dev mgmt proto kernel scope link src 127.0.0.1
```

5. Set the date, time, time zone, and NTP server on the switch.
 - a. Verify the current time zone:

```
cumulus@sw1:~$ cat /etc/timezone
```

- b. Update to the new time zone:

```
cumulus@sw1:~$ sudo dpkg-reconfigure --frontend noninteractive
tzdata
```

c. Verify your current time zone:

```
cumulus@switch:~$ date +%Z
```

d. To set the time zone using the guided wizard, run the following command:

```
cumulus@sw1:~$ sudo dpkg-reconfigure tzdata
```

e. Set the software clock according to the configured time zone:

```
cumulus@switch:~$ sudo date -s "Tue Oct 28 00:37:13 2023"
```

f. Set the current value of the software clock to the hardware clock:

```
cumulus@switch:~$ sudo hwclock -w
```

g. Add an NTP server if required:

```
cumulus@sw1:~$ net add time ntp server <cumulus.network.ntp.org>  
iburst  
cumulus@sw1:~$ net pending  
cumulus@sw1:~$ net commit
```

h. Verify that ntpd is running on the system:

```
cumulus@sw1:~$ ps -ef | grep ntp  
ntp          4074      1  0 Jun20 ?           00:00:33 /usr/sbin/ntpd -p  
/var/run/ntpd.pid -g -u 101:102
```

i. Specify the NTP source interface. By default, the source interface that NTP uses is eth0. You can configure a different NTP source interface as follows:

```
cumulus@sw1:~$ net add time ntp source <src_int>  
cumulus@sw1:~$ net pending  
cumulus@sw1:~$ net commit
```

6. Install Cumulus Linux 4.4.3:


```
cumulus@sw1:mgmt:~$ sudo onie-install -a -i http://<web-server>/<path>/cumulus-linux-4.4.3-mlx-amd64.bin
```

The installer starts the download. Type **y** when prompted.

7. Reboot the NVIDIA SN2100 switch:

```
cumulus@sw1:mgmt:~$ sudo reboot
```

8. The installation starts automatically, and the following GRUB screen choices appear. Do **not** make any selections.

- Cumulus-Linux GNU/Linux
- ONIE: Install OS
- CUMULUS-INSTALL
- Cumulus-Linux GNU/Linux

9. Repeat steps 1 to 4 to log in.

10. Verify that the Cumulus Linux version is 4.4.3: `net show version`

```
cumulus@sw1:mgmt:~$ net show version  
NCLU_VERSION=1.0-cl4.4.3u0  
DISTRIB_ID="Cumulus Linux"  
DISTRIB_RELEASE=4.4.3  
DISTRIB_DESCRIPTION="Cumulus Linux 4.4.3"
```

11. Create a new user and add this user to the `sudo` group. This user only becomes effective after the console/SSH session is restarted.

```
sudo adduser --ingroup netedit admin
```

```

cumulus@sw1:mgmt:~$ sudo adduser --ingroup netedit admin
[sudo] password for cumulus:
Adding user 'admin' ...
Adding new user 'admin' (1001) with group `netedit' ...
Creating home directory '/home/admin' ...
Copying files from '/etc/skel' ...
New password:
Retype new password:
passwd: password updated successfully
Changing the user information for admin
Enter the new value, or press ENTER for the default
Full Name []:
Room Number []:
Work Phone []:
Home Phone []:
Other []:
Is the information correct? [Y/n] y

cumulus@sw1:mgmt:~$ sudo adduser admin sudo
[sudo] password for cumulus:
Adding user `admin' to group `sudo' ...
Adding user admin to group sudo
Done.
cumulus@sw1:mgmt:~$ exit
logout
Connection to 10.233.204.71 closed.

[admin@cycrh6svl01 ~]$ ssh admin@10.233.204.71
admin@10.233.204.71's password:
Linux sw1 4.19.0-cl-1-amd64 #1 SMP Cumulus 4.19.206-1+cl4.4.1u1
(2021-09-09) x86_64
Welcome to NVIDIA Cumulus (R) Linux (R)

For support and online technical documentation, visit
http://www.cumulusnetworks.com/support

The registered trademark Linux (R) is used pursuant to a sublicense
from LMI, the exclusive licensee of Linus Torvalds, owner of the
mark on a world-wide basis.
admin@sw1:mgmt:~$

```

Cumulus Linux 5.4.0

1. Log in to the switch.

First time log in to the switch requires username/password of **cumulus/cumulus** with **sudo**

privileges.

```
cumulus login: cumulus
Password: cumulus
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>
```

2. Check the Cumulus Linux version: `nv show system`

```
cumulus@cumulus:mgmt:~$ nv show system
operational          applied              description
-----
hostname             cumulus             cumulus
build                Cumulus Linux 5.3.0 system build version
uptime              6 days, 8:37:36    system uptime
timezone            Etc/UTC             system time zone
```

3. Configure the hostname, IP address, subnet mask, and default gateway. The new hostname only becomes effective after restarting the console/SSH session.



A Cumulus Linux switch provides at least one dedicated Ethernet management port called `eth0`. This interface is specifically for out-of-band management use. By default, the management interface uses DHCPv4 for addressing.



Do not use an underscore (`_`), apostrophe (`'`), or non-ASCII characters in the hostname.

```
cumulus@cumulus:mgmt:~$ nv set system hostname sw1
cumulus@cumulus:mgmt:~$ nv set interface eth0 ip address
10.233.204.71/24
cumulus@cumulus:mgmt:~$ nv set interface eth0 ip gateway
10.233.204.1
cumulus@cumulus:mgmt:~$ nv config apply
cumulus@cumulus:mgmt:~$ nv config save
```

This command modifies both the `/etc/hostname` and `/etc/hosts` files.

4. Confirm that the hostname, IP address, subnet mask, and default gateway have been updated.

```

cumulus@sw1:mgmt:~$ hostname sw1
cumulus@sw1:mgmt:~$ ifconfig eth0
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 10.233.204.71 netmask 255.255.254.0 broadcast 10.233.205.255
inet6 fe80::bace:f6ff:fe19:1df6 prefixlen 64 scopeid 0x20<link>
ether b8:ce:f6:19:1d:f6 txqueuelen 1000 (Ethernet)
RX packets 75364 bytes 23013528 (21.9 MiB)
RX errors 0 dropped 7 overruns 0 frame 0
TX packets 4053 bytes 827280 (807.8 KiB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0 device
memory 0xdfc00000-dfc1ffff

cumulus@sw1::mgmt:~$ ip route show vrf mgmt
default via 10.233.204.1 dev eth0
unreachable default metric 4278198272
10.233.204.0/23 dev eth0 proto kernel scope link src 10.233.204.71
127.0.0.0/8 dev mgmt proto kernel scope link src 127.0.0.1

```

5. Set the time zone, date, time, and NTP server on the switch.

a. Set the time zone:

```

cumulus@sw1:~$ nv set system timezone US/Eastern
cumulus@sw1:~$ nv config apply

```

b. Verify your current time zone:

```

cumulus@switch:~$ date +%Z

```

c. To set the time zone using the guided wizard, run the following command:

```

cumulus@sw1:~$ sudo dpkg-reconfigure tzdata

```

d. Set the software clock according to the configured time zone:

```

cumulus@sw1:~$ sudo date -s "Tue Oct 28 00:37:13 2023"

```

e. Set the current value of the software clock to the hardware clock:

```

cumulus@sw1:~$ sudo hwclock -w

```

f. Add an NTP server if required:

```
cumulus@sw1:~$ nv set service ntp default server  
<cumulus.network.ntp.org> iburst on  
cumulus@sw1:~$ nv config apply  
cumulus@sw1:~$ nv config save
```

g. Verify that ntpd is running on the system:

```
cumulus@sw1:~$ ps -ef | grep ntp  
ntp          4074      1  0 Jun20 ?           00:00:33 /usr/sbin/ntpd -p  
/var/run/ntpd.pid -g -u 101:102
```

h. Specify the NTP source interface. By default, the source interface that NTP uses is eth0. You can configure a different NTP source interface as follows:

```
cumulus@sw1:~$ nv set service ntp default listen <src_int>  
cumulus@sw1:~$ nv config apply
```

6. Install Cumulus Linux 5.4.0:

```
cumulus@sw1:mgmt:~$ sudo onie-install -a -i http://<web-  
server>/<path>/cumulus-linux-5.4-mlx-amd64.bin
```

The installer starts the download. Type **y** when prompted.

7. Reboot the NVIDIA SN2100 switch:

```
cumulus@sw1:mgmt:~$ sudo reboot
```

8. The installation starts automatically, and the following GRUB screen choices appear. Do **not** make any selections.

- Cumulus-Linux GNU/Linux
- ONIE: Install OS
- CUMULUS-INSTALL
- Cumulus-Linux GNU/Linux

9. Repeat steps 1 to 4 to log in.

10. Verify that the Cumulus Linux version is 5.4.0: `nv show system`

```
cumulus@cumulus:mgmt:~$ nv show system
operational      applied          description
-----
hostname         cumulus         cumulus
build            Cumulus Linux 5.4.0  system build version
uptime           6 days, 13:37:36  system uptime
timezone         Etc/UTC         system time zone
```

11. Verify that the nodes each have a connection to each switch:

```
cumulus@sw1:mgmt:~$ net show lldp

LocalPort  Speed  Mode          RemoteHost
RemotePort
-----
-----
eth0       100M   Mgmt          mgmt-sw1
Eth110/1/29
swp2s1     25G    Trunk/L2      node1
e0a
swp15      100G   BondMember    sw2
swp15
swp16      100G   BondMember    sw2
swp16
```

12. Create a new user and add this user to the `sudo` group. This user only becomes effective after the console/SSH session is restarted.

```
sudo adduser --ingroup netedit admin
```

```

cumulus@sw1:mgmt:~$ sudo adduser --ingroup netedit admin
[sudo] password for cumulus:
Adding user 'admin' ...
Adding new user 'admin' (1001) with group `netedit' ...
Creating home directory '/home/admin' ...
Copying files from '/etc/skel' ...
New password:
Retype new password:
passwd: password updated successfully
Changing the user information for admin
Enter the new value, or press ENTER for the default
Full Name []:
Room Number []:
Work Phone []:
Home Phone []:
Other []:
Is the information correct? [Y/n] y

cumulus@sw1:mgmt:~$ sudo adduser admin sudo
[sudo] password for cumulus:
Adding user `admin' to group `sudo' ...
Adding user admin to group sudo
Done.
cumulus@sw1:mgmt:~$ exit
logout
Connection to 10.233.204.71 closed.

[admin@cycrh6sv101 ~]$ ssh admin@10.233.204.71
admin@10.233.204.71's password:
Linux sw1 4.19.0-cl-1-amd64 #1 SMP Cumulus 4.19.206-1+cl4.4.1u1
(2021-09-09) x86_64
Welcome to NVIDIA Cumulus (R) Linux (R)

For support and online technical documentation, visit
http://www.cumulusnetworks.com/support

The registered trademark Linux (R) is used pursuant to a sublicense
from LMI, the exclusive licensee of Linus Torvalds, owner of the
mark on a world-wide basis.
admin@sw1:mgmt:~$

```

13. Add additional user groups for the admin user to access `nv` commands:

```
cumulus@sw1:mgmt:~$ sudo adduser admin nvshow
[sudo] password for cumulus:
Adding user 'admin' to group 'nvshow' ...
Adding user admin to group nvshow
Done.
```

See [NVIDIA User Accounts](#) for more information.

Cumulus Linux 5.11.0

1. Log in to the switch.

When you log in to the switch for the first time, it requires the username/password of **cumulus** /**cumulus** with sudo privileges.

```
cumulus login: cumulus
Password: cumulus
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>
```

2. Check the Cumulus Linux version: `nv show system`

```
cumulus@cumulus:mgmt:~$ nv show system
operational      applied          description
-----
hostname         cumulus         cumulus
build            Cumulus Linux 5.4.0  system build version
uptime           6 days, 8:37:36  system uptime
timezone         Etc/UTC         system time zone
```

3. Configure the hostname, IP address, subnet mask, and default gateway. The new hostname only becomes effective after restarting the console/SSH session.



A Cumulus Linux switch provides at least one dedicated Ethernet management port called `eth0`. This interface is specifically for out-of-band management use. By default, the management interface uses DHCPv4 for addressing.



Do not use an underscore (`_`), apostrophe (`'`), or non-ASCII characters in the hostname.


```
cumulus@cumulus:mgmt:~$ nv unset interface eth0 ip address dhcp
cumulus@cumulus:mgmt:~$ nv set interface eth0 ip address
10.233.204.71/24
cumulus@cumulus:mgmt:~$ nv set interface eth0 ip gateway
10.233.204.1
cumulus@cumulus:mgmt:~$ nv config apply
cumulus@cumulus:mgmt:~$ nv config save
```

This command modifies both the `/etc/hostname` and `/etc/hosts` files.

4. Confirm that the hostname, IP address, subnet mask, and default gateway have been updated.

```
cumulus@sw1:mgmt:~$ hostname sw1
cumulus@sw1:mgmt:~$ ifconfig eth0
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 10.233.204.71 netmask 255.255.254.0 broadcast 10.233.205.255
inet6 fe80::bace:f6ff:fe19:1df6 prefixlen 64 scopeid 0x20<link>
ether b8:ce:f6:19:1d:f6 txqueuelen 1000 (Ethernet)
RX packets 75364 bytes 23013528 (21.9 MiB)
RX errors 0 dropped 7 overruns 0 frame 0
TX packets 4053 bytes 827280 (807.8 KiB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0 device
memory 0xdfc00000-dfc1ffff

cumulus@sw1::mgmt:~$ ip route show vrf mgmt
default via 10.233.204.1 dev eth0
unreachable default metric 4278198272
10.233.204.0/23 dev eth0 proto kernel scope link src 10.233.204.71
127.0.0.0/8 dev mgmt proto kernel scope link src 127.0.0.1
```

5. Set the time zone, date, time, and NTP server on the switch.

- a. Set the time zone:

```
cumulus@sw1:~$ nv set system timezone US/Eastern
cumulus@sw1:~$ nv config apply
```

- b. Verify your current time zone:

```
cumulus@switch:~$ date +%Z
```

- c. To set the time zone using the guided wizard, run the following command:

```
cumulus@sw1:~$ sudo dpkg-reconfigure tzdata
```

- d. Set the software clock according to the configured time zone:

```
cumulus@sw1:~$ sudo date -s "Tue Oct 28 00:37:13 2023"
```

- e. Set the current value of the software clock to the hardware clock:

```
cumulus@sw1:~$ sudo hwclock -w
```

- f. Add an NTP server if required:

```
cumulus@sw1:~$ nv set service ntp default server  
<cumulus.network.ntp.org> iburst on  
cumulus@sw1:~$ nv config apply  
cumulus@sw1:~$ nv config save
```

- g. Verify that ntpd is running on the system:

```
cumulus@sw1:~$ ps -ef | grep ntp  
ntp          4074      1  0 Jun20 ?           00:00:33 /usr/sbin/ntpd -p  
/var/run/ntpd.pid -g -u 101:102
```

- h. Specify the NTP source interface. By default, the source interface that NTP uses is eth0. You can configure a different NTP source interface as follows:

```
cumulus@sw1:~$ nv set service ntp default listen <src_int>  
cumulus@sw1:~$ nv config apply
```

6. Install Cumulus Linux 5.11.0:

```
cumulus@sw1:mgmt:~$ sudo onie-install -a -i http://<web-  
server>/<path>/cumulus-linux-5.11.0-mlx-amd64.bin
```

The installer starts the download. Type **y** when prompted.

7. Reboot the NVIDIA SN2100 switch:

```
cumulus@sw1:mgmt:~$ sudo reboot
```

8. The installation starts automatically, and the following GRUB screen choices appear. Do **not** make any selections.

- Cumulus-Linux GNU/Linux
- ONIE: Install OS
- CUMULUS-INSTALL
- Cumulus-Linux GNU/Linux

9. Repeat steps 1 to 4 to log in.

10. Verify that the Cumulus Linux version is 5.11.0:

```
nv show system
```

```
cumulus@cumulus:mgmt:~$ nv show system
operational          applied              description
-----
build                Cumulus Linux 5.11.0
uptime              153 days, 2:44:16
hostname            cumulus              cumulus
product-name        Cumulus Linux
product-release     5.11.0
platform            x86_64-mlnx_x86-r0
system-memory       2.76 GB used / 2.28 GB free / 7.47 GB total
swap-memory         0 Bytes used / 0 Bytes free / 0 Bytes total
health-status       not OK
date-time           2025-04-23 09:55:24
status              N/A
timezone            Etc/UTC
maintenance
  mode              disabled
  ports             enabled
version
  kernel            6.1.0-cl-1-amd64
  build-date        Thu Nov 14 13:06:38 UTC 2024
  image             5.11.0
  onie              2019.11-5.2.0020-115200
```

11. Verify that each node has a connection to each switch:

```
cumulus@sw1:mgmt:~$ nv show interface lldp
```

LocalPort	Speed	Mode	RemoteHost
RemotePort			
eth0	100M	eth	mgmt-sw1
Eth110/1/14			
swp2s1	25G	Trunk/L2	node1
e0a			
swp1s1	10G	swp	sw2
e0a			
swp9	100G	swp	sw3
e4a			
swp10	100G	swp	sw4
e4a			
swp15	100G	swp	sw5
swp15			
swp16	100G	swp	sw6
swp16			

See [NVIDIA User Accounts](#) for more information.

What's next?

[Install the Reference Configuration File \(RCF\) script.](#)

Install Cumulus Linux in ONIE mode

Follow this procedure to install Cumulus Linux (CL) OS when the switch is running in ONIE mode.



Cumulus Linux (CL) OS can be installed either when the switch is running ONIE or Cumulus Linux (see [Install in Cumulus mode](#)).

About this task

You can install Cumulus Linux using Open Network Install Environment (ONIE) that allows for automatic discovery of a network installer image. This facilitates the system model of securing switches with an operating system choice, such as Cumulus Linux. The easiest way to install Cumulus Linux with ONIE is with local HTTP discovery.



If your host is IPv6-enabled, make sure it is running a web server. If your host is IPv4-enabled, make sure it is running DHCP in addition to a web server.

This procedure demonstrates how to upgrade Cumulus Linux after the admin has booted in ONIE.

Example 2. Steps

Cumulus Linux 4.4.3

1. Download the Cumulus Linux installation file to the root directory of the web server. Rename this file to: `onie-installer`.
2. Connect your host to the management Ethernet port of the switch using an Ethernet cable.
3. Power on the switch.

The switch downloads the ONIE image installer and boots. After the installation completes, the Cumulus Linux login prompt appears in the terminal window.



Each time Cumulus Linux is installed, the entire file system structure is erased and rebuilt.

4. Reboot the SN2100 switch:

```
cumulus@cumulus:mgmt:~$ sudo reboot
```

5. Press the **Esc** key at the GNU GRUB screen to interrupt the normal boot process, select **ONIE**, and press **Enter**.
6. On the next screen, select **ONIE: Install OS**.
7. The ONIE installer discovery process runs searching for the automatic installation. Press **Enter** to temporarily stop the process.
8. When the discovery process has stopped:

```
ONIE:/ # onie-stop  
discover: installer mode detected.  
Stopping: discover...start-stop-daemon: warning: killing process  
427:  
No such process done.
```

9. If the DHCP service is running on your network, verify that the IP address, subnet mask, and the default gateway are correctly assigned:

```
ifconfig eth0
```

```

ONIE:/ # ifconfig eth0
eth0  Link encap:Ethernet  HWaddr B8:CE:F6:19:1D:F6
      inet addr:10.233.204.71  Bcast:10.233.205.255
Mask:255.255.254.0
      inet6 addr: fe80::bace:f6ff:fe19:1df6/64 Scope:Link
UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
RX packets:21344 errors:0 dropped:2135 overruns:0 frame:0
TX packets:3500 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:6119398 (5.8 MiB)  TX bytes:472975 (461.8 KiB)
Memory:dfc00000-dfc1ffff

```

```

ONIE:/ # route
Kernel IP routing table
Destination      Gateway          Genmask         Flags Metric Ref
Use Iface
default          10.233.204.1    0.0.0.0         UG    0     0
0 eth0
10.233.204.0    *                255.255.254.0   U     0     0
0 eth0

```

10. If the IP addressing scheme is manually defined, do the following:

```

ONIE:/ # ifconfig eth0 10.233.204.71 netmask 255.255.254.0
ONIE:/ # route add default gw 10.233.204.1

```

11. Repeat step 9 to verify that the static information is correctly entered.

12. Install Cumulus Linux:

```

# onie-nos-install http://<web-server>/<path>/cumulus-linux-4.4.3-
mlx-amd64.bin

```

```

ONIE:/ # route

Kernel IP routing table

ONIE:/ # onie-nos-install http://<web-server>/<path>/cumulus-
linux-4.4.3-mlx-amd64.bin

Stopping: discover... done.
Info: Attempting
http://10.60.132.97/x/eng/testbedN,svl/nic/files/cumulus-linux-
4.4.3-mlx-amd64.bin ...
Connecting to 10.60.132.97 (10.60.132.97:80)
installer          100% |*|    552M  0:00:00 ETA
...
...

```

13. After the installation has completed, log in to the switch.

```

cumulus login: cumulus
Password: cumulus
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>

```

14. Verify the Cumulus Linux version: net show version

```

cumulus@cumulus:mgmt:~$ net show version
NCLU_VERSION=1.0-cl4.4.3u4
DISTRIB_ID="Cumulus Linux"
DISTRIB_RELEASE=4.4.3
DISTRIB_DESCRIPTION="Cumulus Linux 4.4.3"

```

Cumulus Linux 5.x

1. Download the Cumulus Linux installation file to the root directory of the web server. Rename this file to: onie-installer.
2. Connect your host to the management Ethernet port of the switch using an Ethernet cable.
3. Power on the switch.

The switch downloads the ONIE image installer and boots. After the installation completes, the Cumulus Linux login prompt appears in the terminal window.



Each time Cumulus Linux is installed, the entire file system structure is erased and rebuilt.

4. Reboot the SN2100 switch:

```
cumulus@cumulus:mgmt:~$ sudo reboot
.
.
GNU GRUB version 2.06-3
+-----+
-----+
| Cumulus-Linux GNU/Linux
|
| Advanced options for Cumulus-Linux GNU/Linux
|
| ONIE
|
|
|
|
|
|
|
|
|
|
|
|
|
|
|
|
|
+-----+
-----+
```

5. Press the Esc key at the GNU GRUB screen to interrupt the normal boot process, select ONIE, and press Enter.


```

.
.
Loading ONIE ...

GNU GRUB version 2.02
+-----+
-----+
| ONIE: Install OS
|
| ONIE: Rescue
|
| ONIE: Uninstall OS
|
| ONIE: Update ONIE
|
| ONIE: Embed ONIE
|
|
|
|
|
|
|
|
|
|
|
|
|
+-----+
-----+

```

Select ONIE: **Install OS**.

6. The ONIE installer discovery process runs searching for the automatic installation. Press **Enter** to temporarily stop the process.
7. When the discovery process has stopped:

```

ONIE:/ # onie-stop
discover: installer mode detected.
Stopping: discover...start-stop-daemon: warning: killing process
427:
No such process done.

```

8. Configure the IP address, subnet mask, and the default gateway:

```
ifconfig eth0
```

```

ONIE:/ # ifconfig eth0
eth0  Link encap:Ethernet  HWaddr B8:CE:F6:19:1D:F6
      inet addr:10.233.204.71  Bcast:10.233.205.255
Mask:255.255.254.0
      inet6 addr: fe80::bace:f6ff:fe19:1df6/64 Scope:Link
      UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
      RX packets:21344 errors:0 dropped:2135 overruns:0 frame:0
      TX packets:3500 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:6119398 (5.8 MiB)  TX bytes:472975 (461.8 KiB)
      Memory:dfc00000-dfc1ffff

ONIE:/ #
ONIE:/ # ifconfig eth0 10.228.140.27 netmask 255.255.248.0
ONIE:/ # ifconfig eth0
eth0  Link encap:Ethernet HWaddr B8:CE:F6:5E:05:E6
      inet addr:10.228.140.27 Bcast:10.228.143.255
Mask:255.255.248.0
      inet6 addr: fd20:8b1e:b255:822b:bace:f6ff:fe5e:5e6/64
Scope:Global
      inet6 addr: fe80::bace:f6ff:fe5e:5e6/64 Scope:Link
      UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
      RX packets:18813 errors:0 dropped:1418 overruns:0 frame:0
      TX packets:491 errors:0 dropped:0 overruns:0 carrier:0
      collisions:0 txqueuelen:1000
      RX bytes:1339596 (1.2 MiB) TX bytes:49379 (48.2 KiB)
      Memory:dfc00000-dfc1ffff

ONIE:/ # route add default gw 10.228.136.1
ONIE:/ # route
Kernel IP routing table
Destination      Gateway          Genmask          Flags Metric Ref
Use Iface

default          10.228.136.1    0.0.0.0          UG    0      0
0 eth0
10.228.136.1    *                255.255.248.0    U    0      0
0 eth0

```

9. Install Cumulus Linux 5.4:

```
# onie-nos-install http://<web-server>/<path>/cumulus-linux-5.4-mlx-amd64.bin
```

```

ONIE:/ # route

Kernel IP routing table

ONIE:/ # onie-nos-install http://<web-server>/<path>/cumulus-
linux-5.4-mlx-amd64.bin

Stopping: discover... done.
Info: Attempting
http://10.60.132.97/x/eng/testbedN,svl/nic/files/cumulus-linux-5.4-
mlx-amd64.bin ...
Connecting to 10.60.132.97 (10.60.132.97:80)
installer          100% |*|    552M  0:00:00 ETA
...
...

```

10. After the installation has completed, log in to the switch.

```

cumulus login: cumulus
Password: cumulus
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>

```

11. Verify the Cumulus Linux version: `nv show system`

```

cumulus@cumulus:mgmt:~$ nv show system
operational      applied          description
-----
hostname         cumulus         cumulus
build            Cumulus Linux 5.4.0  system build version
uptime          6 days, 13:37:36  system uptime
timezone        Etc/UTC         system time zone

```

12. Create a new user and add this user to the `sudo` group. This user only becomes effective after the console/SSH session is restarted.

```

sudo adduser --ingroup netedit admin

```

```

cumulus@sw1:mgmt:~$ sudo adduser --ingroup netedit admin
[sudo] password for cumulus:
Adding user 'admin' ...
Adding new user 'admin' (1001) with group `netedit' ...
Creating home directory '/home/admin' ...
Copying files from '/etc/skel' ...
New password:
Retype new password:
passwd: password updated successfully
Changing the user information for admin
Enter the new value, or press ENTER for the default
Full Name []:
Room Number []:
Work Phone []:
Home Phone []:
Other []:
Is the information correct? [Y/n] y

cumulus@sw1:mgmt:~$ sudo adduser admin sudo
[sudo] password for cumulus:
Adding user `admin' to group `sudo' ...
Adding user admin to group sudo
Done.
cumulus@sw1:mgmt:~$ exit
logout
Connection to 10.233.204.71 closed.

[admin@cycrh6svl01 ~]$ ssh admin@10.233.204.71
admin@10.233.204.71's password:
Linux sw1 4.19.0-cl-1-amd64 #1 SMP Cumulus 4.19.206-1+cl4.4.1u1
(2021-09-09) x86_64
Welcome to NVIDIA Cumulus (R) Linux (R)

For support and online technical documentation, visit
http://www.cumulusnetworks.com/support

The registered trademark Linux (R) is used pursuant to a sublicense
from LMI, the exclusive licensee of Linus Torvalds, owner of the
mark on a world-wide basis.
admin@sw1:mgmt:~$

```

13. Add additional user groups for the admin user to access `nv` commands:

```
cumulus@cumulus:mgmt:~$ sudo adduser admin nvshow
[sudo] password for cumulus:
Adding user `admin' to group `nvshow' ...
Adding user admin to group nvshow
Done.
```

See [NVIDIA User Accounts](#) for more information.

What's next?

[Install the Reference Configuration File \(RCF\) script.](#)

Upgrade Cumulus Linux versions

Complete the following procedure to upgrade your Cumulus Linux version as required.

What you'll need

- Intermediate-level Linux knowledge.
- Familiarity with basic text editing, UNIX file permissions, and process monitoring. A variety of text editors are pre-installed, including `vi` and `nano`.
- Access to a Linux or UNIX shell. If you are running Windows, use a Linux environment as your command line tool for interacting with Cumulus Linux.
- The baud rate requirement is set to 115200 on the serial console switch for NVIDIA SN2100 switch console access, as follows:
 - 115200 baud
 - 8 data bits
 - 1 stop bit
 - parity: none
 - flow control: none

About this task

Be aware of the following:



Each time Cumulus Linux is upgraded, the entire file system structure is erased and rebuilt. Your existing configuration will be erased. You must save and record your switch configuration before updating Cumulus Linux.



The default password for the cumulus user account is **cumulus**. The first time you log into Cumulus Linux, you must change this default password. You must update any automation scripts before installing a new image. Cumulus Linux provides command line options to change the default password automatically during the installation process.

See [Installing a New Cumulus Linux Image](#) for further information.

Example 3. Steps

From Cumulus Linux 4.4.x to Cumulus Linux 5.x

1. Check the current Cumulus Linux version and connected ports:

```
admin@sw1:mgmt:~$ net show system
Hostname..... cumulus
Build..... Cumulus Linux 4.4.3
Uptime..... 0:08:20.860000
Model..... Mlnx X86
CPU..... x86_64 Intel Atom C2558 2.40GHz
Memory..... 8GB
Disk..... 14.7GB
ASIC..... Mellanox Spectrum MT52132
Ports..... 16 x 100G-QSFP28
Part Number..... MSN2100-CB2FC
Serial Number.... MT2105T05177
Platform Name.... x86_64-mlnx_x86-r0
Product Name..... MSN2100
ONIE Version..... 2019.11-5.2.0020-115200
Base MAC Address. 04:3F:72:43:92:80
Manufacturer..... Mellanox

admin@sw1:mgmt:~$ net show interface

State Name      Spd   MTU   Mode      LLDP
Summary
-----
.
.
UP      swp1      100G  9216   Trunk/L2   node1 (e5b)
Master: bridge(UP)
UP      swp2      100G  9216   Trunk/L2   node2 (e5b)
Master: bridge(UP)
UP      swp3      100G  9216   Trunk/L2   SHFFG1826000112 (e0b)
Master: bridge(UP)
UP      swp4      100G  9216   Trunk/L2   SHFFG1826000112 (e0b)
Master: bridge(UP)
UP      swp5      100G  9216   Trunk/L2   SHFFG1826000102 (e0b)
Master: bridge(UP)
UP      swp6      100G  9216   Trunk/L2   SHFFG1826000102 (e0b)
Master: bridge(UP)
.
.
```

2. Download the Cumulux Linux 5.x image:tmp

```
admin@sw1:mgmt:~$ sudo onie-install -a -i http://<ip-to-  
webserver>/path/to/cumulus-linux-5.4.0-mlx-amd64.bin  
[sudo] password for cumulus:  
Fetching installer: http://<ip-to-webserver>/path/to/cumulus-linux-  
5.4.0-mlx-amd64.bin  
Downloading URL: http://<ip-to-webserver>/path/to/cumulus-linux-  
5.4.0-mlx-amd64.bin  
# 100.0%  
Success: HTTP download complete.  
EFI variables are not supported on this system  
Warning: SecureBoot is not available.  
Image is signed.  
. . .  
Staging installer image...done.  
WARNING:  
WARNING: Activating staged installer requested.  
WARNING: This action will wipe out all system data.  
WARNING: Make sure to back up your data.  
WARNING:  
Are you sure (y/N)? y  
Activating staged installer...done.  
Reboot required to take effect.
```

3. Reboot the switch:

```
admin@sw1:mgmt:~$ sudo reboot
```

4. Change the password:

```
cumulus login: cumulus
Password:
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>
Linux cumulus 5.10.0-cl-1-amd64 #1 SMP Debian 5.10.162-1+cl5.4.0u1
(2023-01-20) x86_64

Welcome to NVIDIA Cumulus (R) Linux (R)

ZTP in progress. To disable, do 'ztp -d'
```

5. Check the Cumulus Linux version: `nv show system`

```
cumulus@cumulus:mgmt:~$ nv show system
           operational    applied
-----
hostname   cumulus        cumulus
build      Cumulus Linux 5.4.0
uptime     14:07:08
timezone   Etc/UTC
```

6. Change the hostname:

```
cumulus@cumulus:mgmt:~$ nv set system hostname sw1
cumulus@cumulus:mgmt:~$ nv config apply
Warning: The following files have been changed since the last save,
and they WILL be overwritten.
- /etc/nsswitch.conf
- /etc/synced/synced.conf
.
.
```

7. Logout and log in to the switch again to see the updated switch name at the prompt:


```
cumulus@cumulus:mgmt:~$ exit
logout

Debian GNU/Linux 10 cumulus ttyS0

cumulus login: cumulus
Password:
Last login: Tue Dec 15 21:43:13 UTC 2020 on ttyS0
Linux cumulus 5.10.0-cl-1-amd64 #1 SMP Debian 5.10.162-1+cl5.4.0u1
(2023-01-20) x86_64

Welcome to NVIDIA Cumulus (R) Linux (R)

ZTP in progress. To disable, do 'ztp -d'

cumulus@sw1:mgmt:~$
```

8. Set the IP address:

```
cumulus@sw1:mgmt:~$ nv set interface eth0 ip address
10.231.80.206/22
cumulus@sw1:mgmt:~$ nv set interface eth0 ip gateway 10.231.80.1
cumulus@sw1:mgmt:~$ nv config apply
applied [rev_id: 2]
cumulus@sw1:mgmt:~$ ip route show vrf mgmt
default via 10.231.80.1 dev eth0 proto kernel
unreachable default metric 4278198272
10.231.80.0/22 dev eth0 proto kernel scope link src 10.231.80.206
127.0.0.0/8 dev mgmt proto kernel scope link src 127.0.0.1
```

9. Create a new user and add this user to the `sudo` group. This user only becomes effective after the console/SSH session is restarted.

```
sudo adduser --ingroup netedit admin
```

```

cumulus@sw1:mgmt:~$ sudo adduser --ingroup netedit admin
[sudo] password for cumulus:
Adding user 'admin' ...
Adding new user 'admin' (1001) with group `netedit' ...
Creating home directory '/home/admin' ...
Copying files from '/etc/skel' ...
New password:
Retype new password:
passwd: password updated successfully
Changing the user information for admin
Enter the new value, or press ENTER for the default
Full Name []:
Room Number []:
Work Phone []:
Home Phone []:
Other []:
Is the information correct? [Y/n] y

cumulus@sw1:mgmt:~$ sudo adduser admin sudo
[sudo] password for cumulus:
Adding user `admin' to group `sudo' ...
Adding user admin to group sudo
Done.
cumulus@sw1:mgmt:~$ exit
logout
Connection to 10.233.204.71 closed.

[admin@cycrh6svl01 ~]$ ssh admin@10.233.204.71
admin@10.233.204.71's password:
Linux sw1 4.19.0-cl-1-amd64 #1 SMP Cumulus 4.19.206-1+cl4.4.1u1
(2021-09-09) x86_64
Welcome to NVIDIA Cumulus (R) Linux (R)

For support and online technical documentation, visit
http://www.cumulusnetworks.com/support

The registered trademark Linux (R) is used pursuant to a sublicense
from LMI, the exclusive licensee of Linus Torvalds, owner of the
mark on a world-wide basis.
admin@sw1:mgmt:~$

```

10. Add additional user groups for the admin user to access `nv` commands:

```
cumulus@sw1:mgmt:~$ sudo adduser admin nvshow
[sudo] password for cumulus:
Adding user `admin' to group `nvshow' ...
Adding user admin to group nvshow
Done.
```

See [NVIDIA User Accounts](#) for more information.

From Cumulus Linux 5.x to Cumulus Linux 5.x

1. Check the current Cumulus Linux version and connected ports:

```
admin@sw1:mgmt:~$ nv show system
              operational          applied
-----
hostname      cumulus              cumulus
build         Cumulus Linux 5.3.0
uptime        6 days, 8:37:36
timezone      Etc/UTC

admin@sw1:mgmt:~$ nv show interface
Interface      MTU   Speed State Remote Host      Remote Port-
Type           Summary
-----
+ cluster_isl 9216  200G  up
bond
+ eth0         1500  100M  up   mgmt-sw1      Eth105/1/14
eth            IP Address: 10.231.80 206/22
  eth0
IP Address: fd20:8b1e:f6ff:fe31:4a0e/64
+ lo           65536          up
loopback      IP Address: 127.0.0.1/8
  lo
IP Address: ::1/128
+ swp1s0       9216  10G   up   cluster01     e0b
swp
.
.
.
+ swp15        9216  100G  up   sw2           swp15
swp
+ swp16        9216  100G  up   sw2           swp16
swp
```

2. Download the Cumulus Linux 5.4.0 image:

```
admin@sw1:mgmt:~$ sudo onie-install -a -i http://<ip-to-  
webserver>/path/to/cumulus-linux-5.4.0-mlx-amd64.bin  
[sudo] password for cumulus:  
Fetching installer: http://<ip-to-webserver>/path/to/cumulus-linux-  
5.4.0-mlx-amd64.bin  
Downloading URL: http://<ip-to-webserver>/path/to/cumulus-linux-  
5.4.0-mlx-amd64.bin  
# 100.0%  
Success: HTTP download complete.  
EFI variables are not supported on this system  
Warning: SecureBoot is not available.  
Image is signed.  
. . .  
Staging installer image...done.  
WARNING:  
WARNING: Activating staged installer requested.  
WARNING: This action will wipe out all system data.  
WARNING: Make sure to back up your data.  
WARNING:  
Are you sure (y/N)? y  
Activating staged installer...done.  
Reboot required to take effect.
```

3. Reboot the switch:

```
admin@sw1:mgmt:~$ sudo reboot
```

4. Change the password:

```
cumulus login: cumulus
Password:
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>
Linux cumulus 5.10.0-cl-1-amd64 #1 SMP Debian 5.10.162-1+cl5.4.0u1
(2023-01-20) x86_64

Welcome to NVIDIA Cumulus (R) Linux (R)

ZTP in progress. To disable, do 'ztp -d'
```

5. Check the Cumulus Linux version: `nv show system`

```
cumulus@cumulus:mgmt:~$ nv show system
operational      applied
-----
hostname         cumulus cumulus
build            Cumulus Linux 5.4.0
uptime           14:07:08
timezone         Etc/UTC
```

6. Change the hostname:

```
cumulus@cumulus:mgmt:~$ nv set system hostname sw1
cumulus@cumulus:mgmt:~$ nv config apply
Warning: The following files have been changed since the last save,
and they WILL be overwritten.
- /etc/nsswitch.conf
- /etc/synced/synced.conf
.
.
```

7. Logout and log in again to the switch to see the updated switch name at the prompt:

```
cumulus@cumulus:mgmt:~$ exit
logout

Debian GNU/Linux 10 cumulus ttyS0

cumulus login: cumulus
Password:
Last login: Tue Dec 15 21:43:13 UTC 2020 on ttyS0
Linux cumulus 5.10.0-cl-1-amd64 #1 SMP Debian 5.10.162-1+cl5.4.0u1
(2023-01-20) x86_64

Welcome to NVIDIA Cumulus (R) Linux (R)

ZTP in progress. To disable, do 'ztp -d'

cumulus@sw1:mgmt:~$
```

8. Set the IP address:

```
cumulus@sw1:mgmt:~$ nv unset interface eth0 ip address dhcp
cumulus@sw1:mgmt:~$ nv set interface eth0 ip address
10.231.80.206/22
cumulus@sw1:mgmt:~$ nv set interface eth0 ip gateway 10.231.80.1
cumulus@sw1:mgmt:~$ nv config apply
applied [rev_id: 2]
cumulus@sw1:mgmt:~$ ip route show vrf mgmt
default via 10.231.80.1 dev eth0 proto kernel
unreachable default metric 4278198272
10.231.80.0/22 dev eth0 proto kernel scope link src 10.231.80.206
127.0.0.0/8 dev mgmt proto kernel scope link src 127.0.0.1
```

9. Create a new user and add this user to the `sudo` group. This user only becomes effective after the console/SSH session is restarted.

```
sudo adduser --ingroup netedit admin
```

```

cumulus@sw1:mgmt:~$ sudo adduser --ingroup netedit admin
[sudo] password for cumulus:
Adding user 'admin' ...
Adding new user 'admin' (1001) with group `netedit' ...
Creating home directory '/home/admin' ...
Copying files from '/etc/skel' ...
New password:
Retype new password:
passwd: password updated successfully
Changing the user information for admin
Enter the new value, or press ENTER for the default
Full Name []:
Room Number []:
Work Phone []:
Home Phone []:
Other []:
Is the information correct? [Y/n] y

cumulus@sw1:mgmt:~$ sudo adduser admin sudo
[sudo] password for cumulus:
Adding user `admin' to group `sudo' ...
Adding user admin to group sudo
Done.
cumulus@sw1:mgmt:~$ exit
logout
Connection to 10.233.204.71 closed.

[admin@cycrh6svl01 ~]$ ssh admin@10.233.204.71
admin@10.233.204.71's password:
Linux sw1 4.19.0-cl-1-amd64 #1 SMP Cumulus 4.19.206-1+cl4.4.1u1
(2021-09-09) x86_64
Welcome to NVIDIA Cumulus (R) Linux (R)

For support and online technical documentation, visit
http://www.cumulusnetworks.com/support

The registered trademark Linux (R) is used pursuant to a sublicense
from LMI, the exclusive licensee of Linus Torvalds, owner of the
mark on a world-wide basis.
admin@sw1:mgmt:~$

```

10. Add additional user groups for the admin user to access `nv` commands:

```
cumulus@sw1:mgmt:~$ sudo adduser admin nvshow
[sudo] password for cumulus:
Adding user `admin' to group `nvshow' ...
Adding user admin to group nvshow
Done.
```

See [NVIDIA User Accounts](#) for more information.

What's next?

[Install the Reference Configuration File \(RCF\) script.](#)

Install or upgrade the Reference Configuration File (RCF) script

Follow this procedure to install or upgrade the RCF script.

What you'll need

Before installing or upgrading the RCF script, make sure that the following are available on the switch:

- Cumulus Linux is installed. See the [Hardware Universe](#) for supported versions.
- IP address, subnet mask, and default gateway defined via DHCP or manually configured.



You must specify a user in the RCF (in addition to the admin user) to be used specifically for log collection.

Current RCF script versions

There are two RCF scripts available for Cluster and Storage applications. Download RCFs from [NVIDIA SN2100 software download](#) page. The procedure for each is the same.

- Cluster: **MSN2100-RCF-v1.x-Cluster-HA-Breakout-LLDP**
- Storage: **MSN2100-RCF-v1.x-Storage**

About the examples

The following example procedure shows how to download and apply the RCF script for Cluster switches.

Example command output uses switch management IP address 10.233.204.71, netmask 255.255.254.0 and default gateway 10.233.204.1.

Example 4. Steps

Cumulus Linux 4.4.3

1. Display the available interfaces on the SN2100 switch:

```
admin@sw1:mgmt:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
-----	-----	---	-----	-----	-----	-----
.....						
ADMDN	swp1	N/A	9216	NotConfigured		
ADMDN	swp2	N/A	9216	NotConfigured		
ADMDN	swp3	N/A	9216	NotConfigured		
ADMDN	swp4	N/A	9216	NotConfigured		
ADMDN	swp5	N/A	9216	NotConfigured		
ADMDN	swp6	N/A	9216	NotConfigured		
ADMDN	swp7	N/A	9216	NotConfigured		
ADMDN	swp8	N/A	9216	NotConfigured		
ADMDN	swp9	N/A	9216	NotConfigured		
ADMDN	swp10	N/A	9216	NotConfigured		
ADMDN	swp11	N/A	9216	NotConfigured		
ADMDN	swp12	N/A	9216	NotConfigured		
ADMDN	swp13	N/A	9216	NotConfigured		
ADMDN	swp14	N/A	9216	NotConfigured		
ADMDN	swp15	N/A	9216	NotConfigured		
ADMDN	swp16	N/A	9216	NotConfigured		

2. Copy the RCF python script to the switch.

```
cumulus@cumulus:mgmt:~$ cd /tmp
cumulus@cumulus:mgmt:/tmp$ scp <user>@<host>:<path>/MSN2100-RCF-v1.x
-Cluster-HA-Breakout-LLDP .
ssologin@10.233.204.71's password:
MSN2100-RCF-v1.x-Cluster-HA-Breakout-LLDP          100% 8607
111.2KB/s          00:00
```



While `scp` is used in the example, you can use your preferred method of file transfer.

3. Apply the RCF python script **MSN2100-RCF-v1.x-Cluster-HA-Breakout-LLDP**.

```

cumulus@cumulus:mgmt:/tmp$ sudo python3 MSN2100-RCF-v1.x-Cluster-HA
-Breakout-LLDP
[sudo] password for cumulus:
...
Step 1: Creating the banner file
Step 2: Registering banner message
Step 3: Updating the MOTD file
Step 4: Ensuring passwordless use of cl-support command by admin
Step 5: Disabling apt-get
Step 6: Creating the interfaces
Step 7: Adding the interface config
Step 8: Disabling cdp
Step 9: Adding the lldp config
Step 10: Adding the RoCE base config
Step 11: Modifying RoCE Config
Step 12: Configure SNMP
Step 13: Reboot the switch

```

The RCF script completes the steps listed in the example above.



In step 3 **Updating the MOTD file** above, the command `cat /etc/motd` is run. This allows you to verify the RCF filename, RCF version, ports to use, and other important information in the RCF banner.



For any RCF python script issues that cannot be corrected, contact [NetApp Support](#) for assistance.

4. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
5. Verify the configuration after the reboot:

```

admin@sw1:mgmt:~$ net show interface all

```

State	Name	Spd	MTU	Mode	LLDP	Summary
...						
...						
DN	swp1s0	N/A	9216	Trunk/L2		Master:
	bridge (UP)					
DN	swp1s1	N/A	9216	Trunk/L2		Master:
	bridge (UP)					
DN	swp1s2	N/A	9216	Trunk/L2		Master:
	bridge (UP)					
DN	swp1s3	N/A	9216	Trunk/L2		Master:
	bridge (UP)					

```

DN      swp2s0    N/A    9216    Trunk/L2    Master:
bridge (UP)
DN      swp2s1    N/A    9216    Trunk/L2    Master:
bridge (UP)
DN      swp2s2    N/A    9216    Trunk/L2    Master:
bridge (UP)
DN      swp2s3    N/A    9216    Trunk/L2    Master:
bridge (UP)
UP      swp3      100G   9216    Trunk/L2    Master:
bridge (UP)
UP      swp4      100G   9216    Trunk/L2    Master:
bridge (UP)
DN      swp5      N/A    9216    Trunk/L2    Master:
bridge (UP)
DN      swp6      N/A    9216    Trunk/L2    Master:
bridge (UP)
DN      swp7      N/A    9216    Trunk/L2    Master:
bridge (UP)
DN      swp8      N/A    9216    Trunk/L2    Master:
bridge (UP)
DN      swp9      N/A    9216    Trunk/L2    Master:
bridge (UP)
DN      swp10     N/A    9216    Trunk/L2    Master:
bridge (UP)
DN      swp11     N/A    9216    Trunk/L2    Master:
bridge (UP)
DN      swp12     N/A    9216    Trunk/L2    Master:
bridge (UP)
DN      swp13     N/A    9216    Trunk/L2    Master:
bridge (UP)
DN      swp14     N/A    9216    Trunk/L2    Master:
bridge (UP)
UP      swp15     N/A    9216    BondMember  Master:
bond_15_16 (UP)
UP      swp16     N/A    9216    BondMember  Master:
bond_15_16 (UP)
...
...

```

```
admin@sw1:mgmt:~$ net show roce config
```

```
RoCE mode..... lossless
```

```
Congestion Control:
```

```
Enabled SPs.... 0 2 5
```

```
Mode..... ECN
```

```
Min Threshold.. 150 KB
```

```
Max Threshold.. 1500 KB
```

PFC:

```
Status..... enabled
Enabled SPs.... 2 5
Interfaces..... swp10-16,swp1s0-3,swp2s0-3,swp3-9
```

DSCP	802.1p	switch-priority
0 1 2 3 4 5 6 7	0	0
8 9 10 11 12 13 14 15	1	1
16 17 18 19 20 21 22 23	2	2
24 25 26 27 28 29 30 31	3	3
32 33 34 35 36 37 38 39	4	4
40 41 42 43 44 45 46 47	5	5
48 49 50 51 52 53 54 55	6	6
56 57 58 59 60 61 62 63	7	7

switch-priority	TC	ETS
0 1 3 4 6 7	0	DWRR 28%
2	2	DWRR 28%
5	5	DWRR 43%

6. Verify information for the transceiver in the interface:

```
admin@sw1:mgmt:~$ net show interface pluggables
Interface Identifier      Vendor Name  Vendor PN      Vendor SN
Vendor Rev
-----
swp3      0x11 (QSFP28)  Amphenol     112-00574
APF20379253516 B0
swp4      0x11 (QSFP28)  AVAGO        332-00440      AF1815GU05Z
A0
swp15     0x11 (QSFP28)  Amphenol     112-00573
APF21109348001 B0
swp16     0x11 (QSFP28)  Amphenol     112-00573
APF21109347895 B0
```

7. Verify that the nodes each have a connection to each switch:

```
admin@sw1:mgmt:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
swp3	100G	Trunk/L2	sw1	e3a
swp4	100G	Trunk/L2	sw2	e3b
swp15	100G	BondMember	sw13	swp15
swp16	100G	BondMember	sw14	swp16

8. Verify the health of cluster ports on the cluster.

a. Verify that cluster ports are up and healthy across all nodes in the cluster:

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

```
Node: node1
```

```
Ignore
```

Health	Health					Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
e3a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

```
Node: node2
```

```
Ignore
```

Health	Health					Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
e3a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

b. Verify the switch health from the cluster (this might not show switch sw2, since LIFs are not homed on e0d).

```

cluster1::*> network device-discovery show -protocol lldp
Node/          Local  Discovered
Protocol       Port   Device (LLDP: ChassisID)  Interface Platform
-----
node1/lldp
              e3a    sw1 (b8:ce:f6:19:1a:7e)   swp3          -
              e3b    sw2 (b8:ce:f6:19:1b:96)   swp3          -

node2/lldp
              e3a    sw1 (b8:ce:f6:19:1a:7e)   swp4          -
              e3b    sw2 (b8:ce:f6:19:1b:96)   swp4          -

cluster1::*> system switch ethernet show -is-monitoring-enabled
-operational true
Switch                               Type                Address
Model
-----
sw1                                    cluster-network     10.233.205.90
MSN2100-CB2RC
    Serial Number: MNXXXXXXGD
    Is Monitored: true
    Reason: None
    Software Version: Cumulus Linux version 4.4.3 running on
Mellanox
                                Technologies Ltd. MSN2100
    Version Source: LLDP

sw2                                    cluster-network     10.233.205.91
MSN2100-CB2RC
    Serial Number: MNCXXXXXXGS
    Is Monitored: true
    Reason: None
    Software Version: Cumulus Linux version 4.4.3 running on
Mellanox
                                Technologies Ltd. MSN2100
    Version Source: LLDP

```

Cumulus Linux 5.4.0

1. Display the available interfaces on the SN2100 switch:

```

admin@sw1:mgmt:~$ nv show interface
Interface      MTU   Speed State Remote Host      Remote Port-
Type           Summary
-----
+ cluster_isl  9216  200G  up
bond
+ eth0         1500  100M  up   mgmt-sw1      Eth105/1/14
eth           IP Address: 10.231.80 206/22
  eth0
IP Address: fd20:8b1e:f6ff:fe31:4a0e/64
+ lo           65536      up
loopback     IP Address: 127.0.0.1/8
  lo
IP Address: ::1/128
+ swp1s0       9216  10G   up cluster01      e0b
swp
.
.
.
+ swp15        9216  100G  up sw2           swp15
swp
+ swp16        9216  100G  up sw2           swp16
swp

```

2. Copy the RCF python script to the switch.

```

cumulus@cumulus:mgmt:~$ cd /tmp
cumulus@cumulus:mgmt:/tmp$ scp <user>@<host:/<path>/MSN2100-RCF-v1.x
-Cluster-HA-Breakout-LLDP .
ssologin@10.233.204.71's password:
MSN2100-RCF-v1.x-Cluster-HA-Breakout-LLDP          100% 8607
111.2KB/s          00:00

```



While `scp` is used in the example, you can use your preferred method of file transfer.

3. Apply the RCF python script **MSN2100-RCF-v1.x-Cluster-HA-Breakout-LLDP**.

```
cumulus@cumulus:mgmt:/tmp$ sudo python3 MSN2100-RCF-v1.x-Cluster-HA  
-Breakout-LLDP  
[sudo] password for cumulus:  
.  
.  
Step 1: Creating the banner file  
Step 2: Registering banner message  
Step 3: Updating the MOTD file  
Step 4: Ensuring passwordless use of cl-support command by admin  
Step 5: Disabling apt-get  
Step 6: Creating the interfaces  
Step 7: Adding the interface config  
Step 8: Disabling cdp  
Step 9: Adding the lldp config  
Step 10: Adding the RoCE base config  
Step 11: Modifying RoCE Config  
Step 12: Configure SNMP  
Step 13: Reboot the switch
```

The RCF script completes the steps listed in the example above.



In step 3 **Updating the MOTD file** above, the command `cat /etc/issue` is run. This allows you to verify the RCF filename, RCF version, ports to use, and other important information in the RCF banner.

For example:


```

admin@sw1:mgmt:~$ cat /etc/issue
*****
*****
*
* NetApp Reference Configuration File (RCF)
* Switch      : Mellanox MSN2100
* Filename    : MSN2100-RCF-1._x_-Cluster-HA-Breakout-LLDP
* Release Date : 13-02-2023
* Version     : 1._x_-Cluster-HA-Breakout-LLDP
*
* Port Usage:
* Port 1      : 4x10G Breakout mode for Cluster+HA Ports, swp1s0-3
* Port 2      : 4x25G Breakout mode for Cluster+HA Ports, swp2s0-3
* Ports 3-14  : 40/100G for Cluster+HA Ports, swp3-14
* Ports 15-16 : 100G Cluster ISL Ports, swp15-16
*
* NOTE:
* RCF manually sets swp1s0-3 link speed to 10000 and
* auto-negotiation to off for Intel 10G
* RCF manually sets swp2s0-3 link speed to 25000 and
* auto-negotiation to off for Chelsio 25G
*
* IMPORTANT: Perform the following steps to ensure proper RCF
installation:
* - Copy the RCF file to /tmp
* - Ensure the file has execute permission
* - From /tmp run the file as sudo python3 <filename>
*
*****
*****

```



For any RCF python script issues that cannot be corrected, contact [NetApp Support](#) for assistance.

4. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
5. Verify the configuration after the reboot:

```

admin@sw1:mgmt:~$ nv show interface
Interface      MTU      Speed State Remote Host      Remote Port-
Type          Summary
-----
-----

```

```

+ cluster_isl 9216 200G up
bond
+ eth0          1500 100M up    mgmt-sw1          Eth105/1/14
eth            IP Address: 10.231.80 206/22
  eth0
IP Address: fd20:8ble:f6ff:fe31:4a0e/64
+ lo            65536          up
loopback      IP Address: 127.0.0.1/8
  lo
IP Address: ::1/128
+ swp1s0       9216 10G      up cluster01      e0b
swp
.
.
.
+ swp15        9216 100G      up sw2            swp15
swp
+ swp16        9216 100G      up sw2            swp16
swp

admin@sw1:mgmt:~$ nv show qos roce
              operational  applied  description
-----
enable        on                    Turn feature 'on' or
'off'. This feature is disabled by default.
mode          lossless  lossless  Roce Mode
congestion-control
  congestion-mode  ECN,RED  Congestion config mode
  enabled-tc      0,2,5   Congestion config enabled
Traffic Class
  max-threshold  195.31 KB  Congestion config max-
threshold
  min-threshold  39.06 KB  Congestion config min-
threshold
  probability    100
lldp-app-tlv
  priority       3          switch-priority of roce
  protocol-id    4791      L4 port number
  selector       UDP        L4 protocol
pfc
  pfc-priority   2, 5      switch-prio on which PFC
is enabled
  rx-enabled     enabled    PFC Rx Enabled status
  tx-enabled     enabled    PFC Tx Enabled status
trust

```

```
trust-mode          pcp,dscp          Trust Setting on the port
for packet classification
```

RoCE PCP/DSCP->SP mapping configurations

```
=====
```

	pcp	dscp	switch-prio
--	---	-----	-----
0	0	0,1,2,3,4,5,6,7	0
1	1	8,9,10,11,12,13,14,15	1
2	2	16,17,18,19,20,21,22,23	2
3	3	24,25,26,27,28,29,30,31	3
4	4	32,33,34,35,36,37,38,39	4
5	5	40,41,42,43,44,45,46,47	5
6	6	48,49,50,51,52,53,54,55	6
7	7	56,57,58,59,60,61,62,63	7

RoCE SP->TC mapping and ETS configurations

```
=====
```

	switch-prio	traffic-class	scheduler-weight
--	-----	-----	-----
0	0	0	DWRR-28%
1	1	0	DWRR-28%
2	2	2	DWRR-28%
3	3	0	DWRR-28%
4	4	0	DWRR-28%
5	5	5	DWRR-43%
6	6	0	DWRR-28%
7	7	0	DWRR-28%

RoCE pool config

```
=====
```

	name	mode	size	switch-priorities	
traffic-class					
--	-----	-----	----	-----	
0	lossy-default-ingress	Dynamic	50%	0,1,3,4,6,7	-
1	roce-reserved-ingress	Dynamic	50%	2,5	-
2	lossy-default-egress	Dynamic	50%	-	0
3	roce-reserved-egress	Dynamic	inf	-	2,5

Exception List

```
=====
```

description
--


```
1  RoCE PFC Priority Mismatch.Expected pfc-priority: 3.
2  Congestion Config TC Mismatch.Expected enabled-tc: 0,3.
3  Congestion Config mode Mismatch.Expected congestion-mode:
ECN.
4  Congestion Config min-threshold Mismatch.Expected min-
threshold: 150000.
5  Congestion Config max-threshold Mismatch.Expected max-
threshold:
1500000.
6  Scheduler config mismatch for traffic-class mapped to
switch-prio0.
Expected scheduler-weight: DWRR-50%.
7  Scheduler config mismatch for traffic-class mapped to
switch-prio1.
Expected scheduler-weight: DWRR-50%.
8  Scheduler config mismatch for traffic-class mapped to
switch-prio2.
Expected scheduler-weight: DWRR-50%.
9  Scheduler config mismatch for traffic-class mapped to
switch-prio3.
Expected scheduler-weight: DWRR-50%.
10 Scheduler config mismatch for traffic-class mapped to
switch-prio4.
Expected scheduler-weight: DWRR-50%.
11 Scheduler config mismatch for traffic-class mapped to
switch-prio5.
Expected scheduler-weight: DWRR-50%.
12 Scheduler config mismatch for traffic-class mapped to
switch-prio6.
Expected scheduler-weight: strict-priority.
13 Scheduler config mismatch for traffic-class mapped to
switch-prio7.
Expected scheduler-weight: DWRR-50%.
14 Invalid reserved config for ePort.TC[2].Expected 0 Got 1024
15 Invalid reserved config for ePort.TC[5].Expected 0 Got 1024
16 Invalid traffic-class mapping for switch-priority 2.Expected
0 Got 2
17 Invalid traffic-class mapping for switch-priority 3.Expected
3 Got 0
18 Invalid traffic-class mapping for switch-priority 5.Expected
0 Got 5
19 Invalid traffic-class mapping for switch-priority 6.Expected
6 Got 0
Incomplete Command: set interface swp3-16 link fast-linkupp3-16 link
fast-linkup
Incomplete Command: set interface swp3-16 link fast-linkupp3-16 link
```

```
fast-linkup
```

```
Incomplete Command: set interface swp3-16 link fast-linkupp3-16 link
fast-linkup
```



The exceptions listed don't affect performance and can be safely ignored.

6. Verify information for the transceiver in the interface:

```
admin@sw1:mgmt:~$ nv show interface --view=pluggables
Interface  Identifier      Vendor Name  Vendor PN      Vendor
SN         Vendor Rev
-----
-----
swp1s0     0x00 None
swp1s1     0x00 None
swp1s2     0x00 None
swp1s3     0x00 None
swp2s0     0x11 (QSFP28)  CISCO-LEONI  L45593-D278-D20
LCC2321GTTJ  00
swp2s1     0x11 (QSFP28)  CISCO-LEONI  L45593-D278-D20
LCC2321GTTJ  00
swp2s2     0x11 (QSFP28)  CISCO-LEONI  L45593-D278-D20
LCC2321GTTJ  00
swp2s3     0x11 (QSFP28)  CISCO-LEONI  L45593-D278-D20
LCC2321GTTJ  00
swp3       0x00 None
swp4       0x00 None
swp5       0x00 None
swp6       0x00 None
.
.
.
swp15      0x11 (QSFP28)  Amphenol    112-00595
APF20279210117  B0
swp16      0x11 (QSFP28)  Amphenol    112-00595
APF20279210166  B0
```

7. Verify that the nodes each have a connection to each switch:

```
admin@sw1:mgmt:~$ nv show interface --view=lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
eth0	100M	Mgmt	mgmt-sw1	Eth110/1/29
swp2s1	25G	Trunk/L2	node1	e0a
swp15	100G	BondMember	sw2	swp15
swp16	100G	BondMember	sw2	swp16

8. Verify the health of cluster ports on the cluster.

a. Verify that cluster ports are up and healthy across all nodes in the cluster:

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

```
Node: node1
```

```
Ignore
```

Health	Health					Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
e3a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

```
Node: node2
```

```
Ignore
```

Health	Health					Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
e3a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

b. Verify the switch health from the cluster (this might not show switch sw2, since LIFs are not homed on e0d).

```

cluster1::*> network device-discovery show -protocol lldp
Node/          Local  Discovered
Protocol       Port   Device (LLDP: ChassisID)  Interface Platform
-----
node1/lldp
              e3a    sw1 (b8:ce:f6:19:1a:7e)   swp3      -
              e3b    sw2 (b8:ce:f6:19:1b:96)   swp3      -

node2/lldp
              e3a    sw1 (b8:ce:f6:19:1a:7e)   swp4      -
              e3b    sw2 (b8:ce:f6:19:1b:96)   swp4      -

cluster1::*> system switch ethernet show -is-monitoring-enabled
-operational true
Switch          Type          Address
Model
-----
sw1              cluster-network  10.233.205.90
MSN2100-CB2RC
  Serial Number: MNXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cumulus Linux version 5.4.0 running on
Mellanox
                  Technologies Ltd. MSN2100
  Version Source: LLDP

sw2              cluster-network  10.233.205.91
MSN2100-CB2RC
  Serial Number: MNCXXXXXXGS
  Is Monitored: true
  Reason: None
  Software Version: Cumulus Linux version 5.4.0 running on
Mellanox
                  Technologies Ltd. MSN2100
  Version Source: LLDP

```

Cumulus Linux 5.11.0

1. Display the available interfaces on the SN2100 switch:

```

admin@sw1:mgmt:~$ nv show interface
Interface      MTU   Speed State Remote Host      Remote Port-
Type           Summary
-----
+ cluster_isl 9216  200G  up
bond
+ eth0         1500  100M  up   mgmt-sw1      Eth105/1/14
eth           IP Address: 10.231.80 206/22
  eth0
IP Address: fd20:8b1e:f6ff:fe31:4a0e/64
+ lo           65536      up
loopback     IP Address: 127.0.0.1/8
  lo
IP Address: ::1/128
+ swp1s0       9216  10G   up cluster01      e0b
swp
.
.
.
+ swp15        9216  100G  up sw2            swp15
swp
+ swp16        9216  100G  up sw2            swp16
swp

```

2. Copy the RCF python script to the switch.

```

cumulus@cumulus:mgmt:~$ cd /tmp
cumulus@cumulus:mgmt:/tmp$ scp <user>@<host:/<path>/MSN2100-RCF-v1.x
-Cluster-HA-Breakout-LLDP .
ssologin@10.233.204.71's password:
MSN2100-RCF-v1.x-Cluster-HA-Breakout-LLDP          100% 8607
111.2KB/s          00:00

```



Although `scp` is used in the example, you can use your preferred method of file transfer.

3. Apply the RCF python script **MSN2100-RCF-v1.x-Cluster-HA-Breakout-LLDP**.


```
cumulus@cumulus:mgmt:/tmp$ sudo python3 MSN2100-RCF-v1.x-Cluster-HA  
-Breakout-LLDP  
[sudo] password for cumulus:  
.  
.  
Step 1: Creating the banner file  
Step 2: Registering banner message  
Step 3: Updating the MOTD file  
Step 4: Ensuring passwordless use of cl-support command by admin  
Step 5: Disabling apt-get  
Step 6: Creating the interfaces  
Step 7: Adding the interface config  
Step 8: Disabling cdp  
Step 9: Adding the lldp config  
Step 10: Adding the RoCE base config  
Step 11: Modifying RoCE Config  
Step 12: Configure SNMP  
Step 13: Reboot the switch
```

The RCF script completes the steps listed in the example above.



In step 3 **Updating the MOTD file** above, the command `cat /etc/issue` is run. This allows you to verify the RCF filename, RCF version, ports to use, and other important information in the RCF banner.

For example:

```

admin@sw1:mgmt:~$ cat /etc/issue
*****
*****
*
* NetApp Reference Configuration File (RCF)
* Switch      : Mellanox MSN2100
* Filename    : MSN2100-RCF-1._x_-Cluster-HA-Breakout-LLDP
* Release Date : 13-02-2023
* Version     : 1._x_-Cluster-HA-Breakout-LLDP
*
* Port Usage:
* Port 1      : 4x10G Breakout mode for Cluster+HA Ports, swp1s0-3
* Port 2      : 4x25G Breakout mode for Cluster+HA Ports, swp2s0-3
* Ports 3-14  : 40/100G for Cluster+HA Ports, swp3-14
* Ports 15-16 : 100G Cluster ISL Ports, swp15-16
*
* NOTE:
*   RCF manually sets swp1s0-3 link speed to 10000 and
*   auto-negotiation to off for Intel 10G
*   RCF manually sets swp2s0-3 link speed to 25000 and
*   auto-negotiation to off for Chelsio 25G
*
* IMPORTANT: Perform the following steps to ensure proper RCF
installation:
* - Copy the RCF file to /tmp
* - Ensure the file has execute permission
* - From /tmp run the file as sudo python3 <filename>
*
*****
*****

```



For any RCF python script issues that cannot be corrected, contact [NetApp Support](#) for assistance.

4. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
5. Verify the configuration after the reboot:

```

admin@sw1:mgmt:~$ nv show interface
Interface      MTU   Speed State Remote Host      Remote Port-
Type          Summary
-----
-----

```

```

+ cluster_isl 9216 200G up
bond
+ eth0          1500 100M up    mgmt-sw1          Eth105/1/14
eth            IP Address: 10.231.80 206/22
  eth0
IP Address: fd20:8ble:f6ff:fe31:4a0e/64
+ lo            65536          up
loopback      IP Address: 127.0.0.1/8
  lo
IP Address: ::1/128
+ swp1s0       9216 10G      up cluster01          e0b
swp
.
.
.
+ swp15        9216 100G      up sw2              swp15
swp
+ swp16        9216 100G      up sw2              swp16
swp

admin@sw1:mgmt:~$ nv show qos roce
                    operational  applied  description
-----  -----  -----
enable            on                    Turn feature 'on' or
'off'. This feature is disabled by default.
mode              lossless  lossless  Roce Mode
congestion-control
  congestion-mode  ECN,RED    Congestion config mode
  enabled-tc      0,2,5     Congestion config enabled
Traffic Class
  max-threshold   195.31 KB  Congestion config max-
threshold
  min-threshold   39.06 KB   Congestion config min-
threshold
  probability     100
lldp-app-tlv
  priority        3          switch-priority of roce
  protocol-id     4791      L4 port number
  selector        UDP        L4 protocol
pfc
  pfc-priority    2, 5      switch-prio on which PFC
is enabled
  rx-enabled      enabled    PFC Rx Enabled status
  tx-enabled      enabled    PFC Tx Enabled status
trust

```

```
trust-mode          pcp,dscp          Trust Setting on the port
for packet classification
```

RoCE PCP/DSCP->SP mapping configurations

```
=====
```

	pcp	dscp	switch-prio
--	---	-----	-----
0	0	0,1,2,3,4,5,6,7	0
1	1	8,9,10,11,12,13,14,15	1
2	2	16,17,18,19,20,21,22,23	2
3	3	24,25,26,27,28,29,30,31	3
4	4	32,33,34,35,36,37,38,39	4
5	5	40,41,42,43,44,45,46,47	5
6	6	48,49,50,51,52,53,54,55	6
7	7	56,57,58,59,60,61,62,63	7

RoCE SP->TC mapping and ETS configurations

```
=====
```

	switch-prio	traffic-class	scheduler-weight
--	-----	-----	-----
0	0	0	DWRR-28%
1	1	0	DWRR-28%
2	2	2	DWRR-28%
3	3	0	DWRR-28%
4	4	0	DWRR-28%
5	5	5	DWRR-43%
6	6	0	DWRR-28%
7	7	0	DWRR-28%

RoCE pool config

```
=====
```

	name	mode	size	switch-priorities	
traffic-class					
--	-----	-----	----	-----	
0	lossy-default-ingress	Dynamic	50%	0,1,3,4,6,7	-
1	roce-reserved-ingress	Dynamic	50%	2,5	-
2	lossy-default-egress	Dynamic	50%	-	0
3	roce-reserved-egress	Dynamic	inf	-	2,5

Exception List

```
=====
```

description
--

----...

```
1  RoCE PFC Priority Mismatch.Expected pfc-priority: 3.
2  Congestion Config TC Mismatch.Expected enabled-tc: 0,3.
3  Congestion Config mode Mismatch.Expected congestion-mode:
ECN.
4  Congestion Config min-threshold Mismatch.Expected min-
threshold: 150000.
5  Congestion Config max-threshold Mismatch.Expected max-
threshold:
1500000.
6  Scheduler config mismatch for traffic-class mapped to
switch-prio0.
Expected scheduler-weight: DWRR-50%.
7  Scheduler config mismatch for traffic-class mapped to
switch-prio1.
Expected scheduler-weight: DWRR-50%.
8  Scheduler config mismatch for traffic-class mapped to
switch-prio2.
Expected scheduler-weight: DWRR-50%.
9  Scheduler config mismatch for traffic-class mapped to
switch-prio3.
Expected scheduler-weight: DWRR-50%.
10 Scheduler config mismatch for traffic-class mapped to
switch-prio4.
Expected scheduler-weight: DWRR-50%.
11 Scheduler config mismatch for traffic-class mapped to
switch-prio5.
Expected scheduler-weight: DWRR-50%.
12 Scheduler config mismatch for traffic-class mapped to
switch-prio6.
Expected scheduler-weight: strict-priority.
13 Scheduler config mismatch for traffic-class mapped to
switch-prio7.
Expected scheduler-weight: DWRR-50%.
14 Invalid reserved config for ePort.TC[2].Expected 0 Got 1024
15 Invalid reserved config for ePort.TC[5].Expected 0 Got 1024
16 Invalid traffic-class mapping for switch-priority 2.Expected
0 Got 2
17 Invalid traffic-class mapping for switch-priority 3.Expected
3 Got 0
18 Invalid traffic-class mapping for switch-priority 5.Expected
0 Got 5
19 Invalid traffic-class mapping for switch-priority 6.Expected
6 Got 0
Incomplete Command: set interface swp3-16 link fast-linkupp3-16 link
fast-linkup
Incomplete Command: set interface swp3-16 link fast-linkupp3-16 link
```

```
fast-linkup
```

```
Incomplete Command: set interface swp3-16 link fast-linkupp3-16 link  
fast-linkup
```



The exceptions listed do not affect performance and can be safely ignored.

6. Verify information for the transceiver in the interface:

```
admin@sw1:mgmt:~$ nv show platform transceiver
```

Interface SN	Identifier Vendor Rev	Vendor Name	Vendor PN	Vendor
swp1s0	0x00 None			
swp1s1	0x00 None			
swp1s2	0x00 None			
swp1s3	0x00 None			
swp2s0	0x11 (QSFP28)	CISCO-LEONI	L45593-D278-D20	
LCC2321GTTJ	00			
swp2s1	0x11 (QSFP28)	CISCO-LEONI	L45593-D278-D20	
LCC2321GTTJ	00			
swp2s2	0x11 (QSFP28)	CISCO-LEONI	L45593-D278-D20	
LCC2321GTTJ	00			
swp2s3	0x11 (QSFP28)	CISCO-LEONI	L45593-D278-D20	
LCC2321GTTJ	00			
swp3	0x00 None			
swp4	0x00 None			
swp5	0x00 None			
swp6	0x00 None			
.				
.				
.				
swp15	0x11 (QSFP28)	Amphenol	112-00595	
APF20279210117	B0			
swp16	0x11 (QSFP28)	Amphenol	112-00595	
APF20279210166	B0			

7. Verify that the nodes each have a connection to each switch:

```
admin@sw1:mgmt:~$ nv show interface lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
eth0	100M	Mgmt	mgmt-sw1	Eth110/1/29
swp2s1	25G	Trunk/L2	node1	e0a
swp15	100G	BondMember	sw2	swp15
swp16	100G	BondMember	sw2	swp16

8. Verify the health of cluster ports on the cluster.

a. Verify that cluster ports are up and healthy across all nodes in the cluster:

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

```
Node: node1
```

```
Ignore
```

Health	Health					Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
e3a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

```
Node: node2
```

```
Ignore
```

Health	Health					Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
e3a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

b. Verify the switch health from the cluster (this might not show switch sw2, since LIFs are not homed on e0d).

```

cluster1::*> network device-discovery show -protocol lldp
Node/          Local  Discovered
Protocol       Port   Device (LLDP: ChassisID)  Interface Platform
-----
node1/lldp
              e3a    sw1 (b8:ce:f6:19:1a:7e)   swp3          -
              e3b    sw2 (b8:ce:f6:19:1b:96)   swp3          -

node2/lldp
              e3a    sw1 (b8:ce:f6:19:1a:7e)   swp4          -
              e3b    sw2 (b8:ce:f6:19:1b:96)   swp4          -

cluster1::*> system switch ethernet show -is-monitoring-enabled
-operational true
Switch                Type                Address
Model
-----
sw1                    cluster-network    10.233.205.90
MSN2100-CB2RC
    Serial Number: MNXXXXXXGD
    Is Monitored: true
    Reason: None
    Software Version: Cumulus Linux version 5.4.0 running on
Mellanox
                    Technologies Ltd. MSN2100
    Version Source: LLDP

sw2                    cluster-network    10.233.205.91
MSN2100-CB2RC
    Serial Number: MNCXXXXXXGS
    Is Monitored: true
    Reason: None
    Software Version: Cumulus Linux version 5.4.0 running on
Mellanox
                    Technologies Ltd. MSN2100
    Version Source: LLDP

```

What's next?

[Install the CSHM file.](#)

Install the Ethernet Switch Health Monitor configuration file

Follow this procedure to install the applicable configuration file for Ethernet switch health monitoring of NVIDIA cluster switches.

Supported models are:

- MSN2100-CB2FC
- MSN2100-CB2RC
- X190006-PE
- X190006-PI



This installation procedure applies to ONTAP 9.10.1 and later.

Before you begin

- Verify that you need to download the configuration file by running `system switch ethernet show` and checking if **OTHER** is shown for your model.

If your model is still showing **OTHER** after applying the configuration file, contact NetApp support.

- Make sure that the ONTAP cluster is up and running.
- Enable SSH to use all of the features available in CSHM.
- Clear the `/mroot/etc/cshm_nod/nod_sign/` directory on all nodes:

- a. Enter the nodeshell:

```
system node run -node <name>
```

- b. Change to advanced privilege:

```
priv set advanced
```

- c. List the configuration files in the `/etc/cshm_nod/nod_sign` directory. If the directory exists and contains configuration files, it lists the file names.

```
ls /etc/cshm_nod/nod_sign
```

- d. Delete all configuration files corresponding to your connected switch models.

If you are unsure, remove all configuration files for the supported models listed above, then download and install the latest configuration files for those same models.

```
rm /etc/cshm_nod/nod_sign/<filename>
```

- e. Confirm that the deleted configuration files are no longer in the directory:

```
ls /etc/cshm_nod/nod_sign
```

Steps

1. Download the Ethernet switch health monitor configuration zip file based on the corresponding ONTAP release version. This file is available from the [NVIDIA Ethernet switches](#) page.

- a. On the NVIDIA SN2100 Software download page, select **Nvidia CSHM File**.
- b. On the Caution/Must read page, select the check box to agree.
- c. On the End User License Agreement page, select the check box to agree and click **Accept & Continue**.
- d. On the Nvidia CSHM File - Download page, select the applicable configuration file. The following files are available:

ONTAP 9.15.1 and later

- MSN2100-CB2FC-v1.4.zip
- MSN2100-CB2RC-v1.4.zip
- X190006-PE-v1.4.zip
- X190006-PI-v1.4.zip

ONTAP 9.11.1 through 9.14.1

- MSN2100-CB2FC_PRIOR_R9.15.1-v1.4.zip
- MSN2100-CB2RC_PRIOR_R9.15.1-v1.4.zip
- X190006-PE_PRIOR_9.15.1-v1.4.zip
- X190006-PI_PRIOR_9.15.1-v1.4.zip

2. Upload the applicable zip file to your internal web server.
3. Access the advanced mode setting from one of the ONTAP systems in the cluster.

```
set -privilege advanced
```

4. Run the switch health monitor configure command.

```
cluster1::> system cluster-switch configure-health-monitor -node *  
-package-url http://192.168.2.20/usr/download/[filename.zip]
```

5. Verify that the command output ends with the following text for your ONTAP version:

ONTAP 9.15.1 and later

Ethernet switch health monitoring installed the configuration file.

ONTAP 9.11.1 through 9.14.1

SHM installed the configuration file.

ONTAP 9.10.1

CSHM downloaded package processed successfully.

If an error occurs, contact NetApp support.

6. Wait up to twice the Ethernet switch health monitor polling interval, found by running `system switch`

ethernet polling-interval show, before completing the next step.

7. Run the command `system switch ethernet show` on the ONTAP system and make sure that the cluster switches are discovered with the monitored field set to **True** and the serial number field not showing **Unknown**.

```
cluster1::> system switch ethernet show
```

What's next?

[Configure switch health monitoring.](#)

Migrate switches

Migrate from a Cisco cluster switch to a NVIDIA SN2100 cluster switch

You can migrate Cisco cluster switches for an ONTAP cluster to NVIDIA SN2100 cluster switches. This is a nondisruptive procedure.

Review requirements

You must be aware of certain configuration information, port connections and cabling requirements when you are replacing some older Cisco cluster switches with NVIDIA SN2100 cluster switches. See [Overview of installation and configuration for NVIDIA SN2100 switches](#).

Supported switches

The following Cisco cluster switches are supported:

- Nexus 9336C-FX2
- Nexus 92300YC
- Nexus 5596UP
- Nexus 3232C
- Nexus 3132Q-V

For details of supported ports and their configurations, see the [Hardware Universe](#) .

What you'll need

Ensure that:

- The existing cluster is properly set up and functioning.
- All cluster ports are in the **up** state to ensure nondisruptive operations.
- The NVIDIA SN2100 cluster switches are configured and operating under the proper version of Cumulus Linux installed with the reference configuration file (RCF) applied.
- The existing cluster network configuration have the following:
 - A redundant and fully functional NetApp cluster using both older Cisco switches.
 - Management connectivity and console access to both the older Cisco switches and the new switches.
 - All cluster LIFs in the up state with the cluster LIFs are on their home ports.

- ISL ports enabled and cabled between the older Cisco switches and between the new switches.
- Some of the ports are configured on NVIDIA SN2100 switches to run at 40 GbE or 100 GbE.
- You have planned, migrated, and documented 40 GbE and 100 GbE connectivity from nodes to NVIDIA SN2100 cluster switches.



If you are changing the port speed of the e0a and e1a cluster ports on AFF A800 or AFF C800 systems, you might observe malformed packets being received after the speed conversion. See [Bug 1570339](#) and the Knowledge Base article [CRC errors on T6 ports after converting from 40GbE to 100GbE](#) for guidance.

Migrate the switches

About the examples

In this procedure, Cisco Nexus 3232C cluster switches are used for example commands and outputs.

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

- The existing Cisco Nexus 3232C cluster switches are *c1* and *c2*.
- The new NVIDIA SN2100 cluster switches are *sw1* and *sw2*.
- The nodes are *node1* and *node2*.
- The cluster LIFs are *node1_clus1* and *node1_clus2* on node 1, and *node2_clus1* and *node2_clus2* on node 2 respectively.
- The `cluster1::*>` prompt indicates the name of the cluster.
- The cluster ports used in this procedure are *e3a* and *e3b*.
- Breakout ports take the format: `swp[port]s[breakout port 0-3]`. For example, four breakout ports on `swp1` are *swp1s0*, *swp1s1*, *swp1s2*, and *swp1s3*.

About this task

This procedure covers the following scenario:

- Switch *c2* is replaced by switch *sw2* first.
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - Cabling between the nodes and *c2* are then disconnected from *c2* and reconnected to *sw2*.
- Switch *c1* is replaced by switch *sw1*.
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - Cabling between the nodes and *c1* are then disconnected from *c1* and reconnected to *sw1*.

Step 1: Prepare for migration

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
```

where *x* is the duration of the maintenance window in hours.

2. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (*>) appears.

3. Disable auto-revert on the cluster LIFs:

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

Step 2: Configure ports and cabling

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

Each port should display up for `Link` and healthy for `Health Status`.

- a. Display the network port attributes:

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

Show example

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

Node: node1

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper  Speed (Mbps)
Status      Status
-----
e3a         Cluster    Cluster    up    9000  auto/100000
healthy     false
e3b         Cluster    Cluster    up    9000  auto/100000
healthy     false

Node: node2

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper  Speed (Mbps)
Status      Status
-----
e3a         Cluster    Cluster    up    9000  auto/100000
healthy     false
e3b         Cluster    Cluster    up    9000  auto/100000
healthy     false
```

- b. Display information about the logical interfaces and their designated home nodes:

```
network interface show -vserver Cluster
```

Each LIF should display up/up for Status Admin/Oper and true for Is Home.

Show example

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

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

Cluster				
e3a	node1_clus1	up/up	169.254.209.69/16	node1
e3b	true			
e3a	node1_clus2	up/up	169.254.49.125/16	node1
e3b	true			
e3a	node2_clus1	up/up	169.254.47.194/16	node2
e3b	true			
e3a	node2_clus2	up/up	169.254.19.183/16	node2
e3b	true			

2. The cluster ports on each node are connected to existing cluster switches in the following way (from the nodes' perspective):

```
network device-discovery show -protocol lldp
```

Show example

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

Node/	Local	Discovered		
Protocol	Port	Device (LLDP: ChassisID)	Interface	
Platform				

node1	/lldp			
	e3a	c1 (6a:ad:4f:98:3b:3f)	Eth1/1	-
	e3b	c2 (6a:ad:4f:98:4c:a4)	Eth1/1	-
node2	/lldp			
	e3a	c1 (6a:ad:4f:98:3b:3f)	Eth1/2	-
	e3b	c2 (6a:ad:4f:98:4c:a4)	Eth1/2	-

3. The cluster ports and switches are connected in the following way (from the switches' perspective):

```
show cdp neighbors
```

Show example

```
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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e3a	Eth1/1	124	H	AFF-A400
node2 e3a	Eth1/2	124	H	AFF-A400
c2 Eth1/31	Eth1/31	179	S I s	N3K-C3232C
c2 Eth1/32	Eth1/32	175	S I s	N3K-C3232C

```
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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e3b	Eth1/1	124	H	AFF-A400
node2 e3b	Eth1/2	124	H	AFF-A400
c1 Eth1/31	Eth1/31	175	S I s	N3K-C3232C
c1 Eth1/32	Eth1/32	175	S I s	N3K-C3232C

4. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e3a
Cluster node1_clus2 169.254.49.125 node1 e3b
Cluster node2_clus1 169.254.47.194 node2 e3a
Cluster node2_clus2 169.254.19.183 node2 e3b
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)

```

5. On switch c2, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs.

```

(c2)# configure
Enter configuration commands, one per line. End with CNTL/Z.

(c2) (Config)# interface
(c2) (config-if-range)# shutdown <interface_list>
(c2) (config-if-range)# exit
(c2) (Config)# exit
(c2)#

```

6. Move the node cluster ports from the old switch c2 to the new switch sw2, using appropriate cabling supported by NVIDIA SN2100.
7. Display the network port attributes:

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

Show example

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

```
Node: node1
```

```
Ignore
```

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

e3a	Cluster	Cluster		up	9000	auto/100000		healthy	false
e3b	Cluster	Cluster		up	9000	auto/100000		healthy	false

```
Node: node2
```

```
Ignore
```

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

e3a	Cluster	Cluster		up	9000	auto/100000		healthy	false
e3b	Cluster	Cluster		up	9000	auto/100000		healthy	false

8. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

Show example

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

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	
node1	/lldp			
	e3a	c1 (6a:ad:4f:98:3b:3f)	Eth1/1	-
	e3b	sw2 (b8:ce:f6:19:1a:7e)	swp3	-
node2	/lldp			
	e3a	c1 (6a:ad:4f:98:3b:3f)	Eth1/2	-
	e3b	sw2 (b8:ce:f6:19:1b:96)	swp4	-

9. On switch sw2, verify that all node cluster ports are up:

```
net show interface
```

Show example

```
cumulus@sw2:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP
Summary					

.....					
UP	swp3	100G	9216	Trunk/L2	e3b
Master: bridge(UP)					
UP	swp4	100G	9216	Trunk/L2	e3b
Master: bridge(UP)					
UP	swp15	100G	9216	BondMember	sw1 (swp15)
Master: cluster_isl(UP)					
UP	swp16	100G	9216	BondMember	sw1 (swp16)
Master: cluster_isl(UP)					

10. On switch c1, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs.

```
(c1)# configure
```

Enter configuration commands, one per line. End with CNTL/Z.

```
(c1) (Config)# interface
```

```
(c1) (config-if-range)# shutdown <interface_list>
```

```
(c1) (config-if-range)# exit
```

```
(c1) (Config)# exit
```

```
(c1)#
```

11. Move the node cluster ports from the old switch c1 to the new switch sw1, using appropriate cabling supported by NVIDIA SN2100.
12. Verify the final configuration of the cluster:

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

Each port should display up for Link and healthy for Health Status.

Show example

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

```
Node: node1
```

```
Ignore
```

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

e3a	Cluster	Cluster		up	9000	auto/100000			
healthy	false								
e3b	Cluster	Cluster		up	9000	auto/100000			
healthy	false								

```
Node: node2
```

```
Ignore
```

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

e3a	Cluster	Cluster		up	9000	auto/100000			
healthy	false								
e3b	Cluster	Cluster		up	9000	auto/100000			
healthy	false								

13. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

Show example

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

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	

node1	/lldp			
	e3a	sw1 (b8:ce:f6:19:1a:7e)	swp3	-
	e3b	sw2 (b8:ce:f6:19:1b:96)	swp3	-
node2	/lldp			
	e3a	sw1 (b8:ce:f6:19:1a:7e)	swp4	-
	e3b	sw2 (b8:ce:f6:19:1b:96)	swp4	-

14. On switches sw1 and sw2, verify that all node cluster ports are up:

```
net show interface
```

Show example

```
cumulus@sw1:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP
Summary					

.....					
...					
...					
UP	swp3	100G	9216	Trunk/L2	e3a
Master: bridge(UP)					
UP	swp4	100G	9216	Trunk/L2	e3a
Master: bridge(UP)					
UP	swp15	100G	9216	BondMember	sw2 (swp15)
Master: cluster_isl(UP)					
UP	swp16	100G	9216	BondMember	sw2 (swp16)
Master: cluster_isl(UP)					

```
cumulus@sw2:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP
Summary					

.....					
...					
...					
UP	swp3	100G	9216	Trunk/L2	e3b
Master: bridge(UP)					
UP	swp4	100G	9216	Trunk/L2	e3b
Master: bridge(UP)					
UP	swp15	100G	9216	BondMember	sw1 (swp15)
Master: cluster_isl(UP)					
UP	swp16	100G	9216	BondMember	sw1 (swp16)
Master: cluster_isl(UP)					

15. Verify that both nodes each have one connection to each switch:

```
net show lldp
```


Show example

The following example shows the appropriate results for both switches:

```
cumulus@sw1:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
swp3	100G	Trunk/L2	node1	e3a
swp4	100G	Trunk/L2	node2	e3a
swp15	100G	BondMember	sw2	swp15
swp16	100G	BondMember	sw2	swp16

```
cumulus@sw2:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
swp3	100G	Trunk/L2	node1	e3b
swp4	100G	Trunk/L2	node2	e3b
swp15	100G	BondMember	sw1	swp15
swp16	100G	BondMember	sw1	swp16

Step 3: Verify the configuration

1. Enable auto-revert on the cluster LIFs:

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

2. Verify that all cluster network LIFs are back on their home ports:

```
network interface show
```

Show example

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

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

Cluster				
e3a	node1_clus1	up/up	169.254.209.69/16	node1
	true			
e3b	node1_clus2	up/up	169.254.49.125/16	node1
	true			
e3a	node2_clus1	up/up	169.254.47.194/16	node2
	true			
e3b	node2_clus2	up/up	169.254.19.183/16	node2
	true			

3. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

What's next?

[Configure switch health monitoring.](#)

Migrate to a two-node switched cluster with NVIDIA SN2100 cluster switches

If you have an existing two-node switchless cluster environment, you can migrate to a two-node switched cluster environment using NVIDIA SN2100 switches to enable you to scale beyond two nodes in the cluster.

The procedure you use depends on whether you have two dedicated cluster-network ports on each controller or a single cluster port on each controller. The process documented works for all nodes using optical or Twinax ports but is not supported on this switch if nodes are using onboard 10GBASE-T RJ45 ports for the cluster-network ports.

Review requirements

Two-node switchless configuration

Ensure that:

- The two-node switchless configuration are properly set up and functioning.
- The nodes are running ONTAP 9.10.1P3 and later.
- All cluster ports are in the **up** state.
- All cluster logical interfaces (LIFs) are in the **up** state and on their home ports.

NVIDIA SN2100 cluster switch configuration

Ensure that:

- Both switches have management network connectivity.
- There is console access to the cluster switches.
- NVIDIA SN2100 node-to-node switch and switch-to-switch connections use Twinax or fiber cables.



See [Review cabling and configuration considerations](#) for caveats and further details. The [Hardware Universe - Switches](#) also contains more information about cabling.

- Inter-Switch Link (ISL) cables are connected to ports swp15 and swp16 on both NVIDIA SN2100 switches.
- Initial customization of both the SN2100 switches are completed, so that:
 - SN2100 switches are running the latest version of Cumulus Linux
 - Reference Configuration Files (RCFs) are applied to the switches
 - Any site customization, such as SMTP, SNMP, and SSH are configured on the new switches.

The [Hardware Universe](#) contains the latest information about the actual cluster ports for your platforms.

Migrate the switches

About the examples

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

- The names of the SN2100 switches are *sw1* and *sw2*.
- The names of the cluster SVMs are *node1* and *node2*.
- The names of the LIFs are *node1_clus1* and *node1_clus2* on node 1, and *node2_clus1* and *node2_clus2* on node 2 respectively.
- The `cluster1::*>` prompt indicates the name of the cluster.
- The cluster ports used in this procedure are *e3a* and *e3b*.
- Breakout ports take the format: `swp[port]s[breakout port 0-3]`. For example, four breakout ports on swp1 are *swp1s0*, *swp1s1*, *swp1s2*, and *swp1s3*.

Step 1: Prepare for migration

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`

where *x* is the duration of the maintenance window in hours.

2. Change the privilege level to advanced, entering *y* when prompted to continue: `set -privilege advanced`

The advanced prompt (*>) appears.

Step 2: Configure ports and cabling

Cumulus Linux 4.4.x

1. Disable all node-facing ports (not ISL ports) on both the new cluster switches sw1 and sw2.

You must not disable the ISL ports.

The following commands disable the node-facing ports on switches sw1 and sw2:

```
cumulus@sw1:~$ net add interface swp1s0-3, swp2s0-3, swp3-14 link
down
cumulus@sw1:~$ net pending
cumulus@sw1:~$ net commit

cumulus@sw2:~$ net add interface swp1s0-3, swp2s0-3, swp3-14 link
down
cumulus@sw2:~$ net pending
cumulus@sw2:~$ net commit
```

2. Verify that the ISL and the physical ports on the ISL between the two SN2100 switches sw1 and sw2 are up on ports swp15 and swp16:

```
net show interface
```

The following commands show that the ISL ports are up on switches sw1 and sw2:

```
cumulus@sw1:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp15	100G	9216	BondMember	sw2 (swp15)	Master: cluster_isl (UP)
UP	swp16	100G	9216	BondMember	sw2 (swp16)	Master: cluster_isl (UP)

```
cumulus@sw2:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp15	100G	9216	BondMember	sw1 (swp15)	Master: cluster_isl (UP)
UP	swp16	100G	9216	BondMember	sw1 (swp16)	Master: cluster_isl (UP)

Cumulus Linux 5.x

1. Disable all node-facing ports (not ISL ports) on both new cluster switches sw1 and sw2.

You must not disable the ISL ports.

The following commands disable the node-facing ports on switches sw1 and sw2:

```
cumulus@sw1:~$ nv set interface swp1s0-3,swp2s0-3,swp3-14 link state  
down  
cumulus@sw1:~$ nv config apply  
cumulus@sw1:~$ nv config save  
  
cumulus@sw2:~$ nv set interface swp1s0-3,swp2s0-3,swp3-14 link state  
down  
cumulus@sw2:~$ nv config apply  
cumulus@sw2:~$ nv config save
```

2. Verify that the ISL and the physical ports on the ISL between the two SN2100 switches sw1 and sw2 are up on ports swp15 and swp16:

```
nv show interface
```

The following examples show that the ISL ports are up on switches sw1 and sw2:

```
cumulus@sw1:~$ nv show interface
```

```
Interface      MTU      Speed  State  Remote Host  Remote Port
Type          Summary
-----
...
...
+ swp14        9216           down
swp
+ swp15        9216    100G   up     ossg-rcf1    Intra-Cluster Switch
ISL Port swp15  swp
+ swp16        9216    100G   up     ossg-rcf2    Intra-Cluster Switch
ISL Port swp16  swp
```

```
cumulus@sw2:~$ nv show interface
```

```
Interface      MTU      Speed  State  Remote Host  Remote Port
Type          Summary
-----
...
...
+ swp14        9216           down
swp
+ swp15        9216    100G   up     ossg-rcf1    Intra-Cluster Switch
ISL Port swp15  swp
+ swp16        9216    100G   up     ossg-rcf2    Intra-Cluster Switch
ISL Port swp16  swp
```

3. Verify that all cluster ports are up:

```
network port show
```

Each port should display up for Link and healthy for Health Status.

Show example

```
cluster1::*> network port show
```

```
Node: node1
```

```
Ignore
```

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

e3a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/100000
healthy	false					

```
Node: node2
```

```
Ignore
```

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

e3a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/100000
healthy	false					

4. Verify that all cluster LIFs are up and operational:

```
network interface show
```

Each cluster LIF should display true for Is Home and have a Status Admin/Oper of up/up.

Show example

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

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

Cluster				
e3a	node1_clus1	up/up	169.254.209.69/16	node1
e3b	true			
e3a	node1_clus2	up/up	169.254.49.125/16	node1
e3b	true			
e3a	node2_clus1	up/up	169.254.47.194/16	node2
e3b	true			
e3a	node2_clus2	up/up	169.254.19.183/16	node2
e3b	true			

5. Disable auto-revert on the cluster LIFs:

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

Show example

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

Vserver	Logical	Auto-revert
Interface		

Cluster		
	node1_clus1	false
	node1_clus2	false
	node2_clus1	false
	node2_clus2	false

6. Disconnect the cable from cluster port e3a on node1, and then connect e3a to port 3 on cluster switch sw1, using the appropriate cabling supported by the SN2100 switches.

The [Hardware Universe - Switches](#) contains more information about cabling.

7. Disconnect the cable from cluster port e3a on node2, and then connect e3a to port 4 on cluster switch sw1,

using the appropriate cabling supported by the SN2100 switches.

Cumulus Linux 4.4.x

8. On switch sw1, enable all node-facing ports.

The following commands enable all node-facing ports on switch sw1.

```
cumulus@sw1:~$ net del interface swp1s0-3, swp2s0-3, swp3-14 link
down
cumulus@sw1:~$ net pending
cumulus@sw1:~$ net commit
```

9. On switch sw1, verify that all ports are up:

```
net show interface all
```

```
cumulus@sw1:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
...						
DN	swp1s0	10G	9216	Trunk/L2		Master: br_default(UP)
DN	swp1s1	10G	9216	Trunk/L2		Master: br_default(UP)
DN	swp1s2	10G	9216	Trunk/L2		Master: br_default(UP)
DN	swp1s3	10G	9216	Trunk/L2		Master: br_default(UP)
DN	swp2s0	25G	9216	Trunk/L2		Master: br_default(UP)
DN	swp2s1	25G	9216	Trunk/L2		Master: br_default(UP)
DN	swp2s2	25G	9216	Trunk/L2		Master: br_default(UP)
DN	swp2s3	25G	9216	Trunk/L2		Master: br_default(UP)
UP	swp3	100G	9216	Trunk/L2	node1 (e3a)	Master: br_default(UP)
UP	swp4	100G	9216	Trunk/L2	node2 (e3a)	Master: br_default(UP)
...						
...						
UP	swp15	100G	9216	BondMember	swp15	Master: cluster_isl(UP)
UP	swp16	100G	9216	BondMember	swp16	Master: cluster_isl(UP)
...						

Cumulus Linux 5.x

8. On switch sw1, enable all node-facing ports.

The following commands enable all node-facing ports on switch sw1.

```
cumulus@sw1:~$ nv set interface swp1s0-3,swp2s0-3,swp3-14 link state  
up  
cumulus@sw1:~$ nv config apply  
cumulus@sw1:~$ nv config save
```

9. On switch sw1, verify that all ports are up:

```
nv show interface
```

```
cumulus@sw1:~$ nv show interface
```

Interface	State	Speed	MTU	Type	Remote Host
Remote Port	Summary				
...					
...					
swp1s0	up	10G	9216	swp	odq-a300-1a
e0a					
swp1s1	up	10G	9216	swp	odq-a300-1b
e0a					
swp1s2	down	10G	9216	swp	
swp1s3	down	10G	9216	swp	
swp2s0	down	25G	9216	swp	
swp2s1	down	25G	9216	swp	
swp2s2	down	25G	9216	swp	
swp2s3	down	25G	9216	swp	
swp3	down		9216	swp	
swp4	down		9216	swp	
...					
...					
swp14	down		9216	swp	
swp15	up	100G	9216	swp	ossg-int-rcf10
swp15					
swp16	up	100G	9216	swp	ossg-int-rcf10
swp16					

10. Verify that all cluster ports are up:

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

Show example

The following example shows that all of the cluster ports are up on node1 and node2:

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

Node: node1

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
-----
e3a         Cluster   Cluster          up  9000  auto/100000
healthy    false
e3b         Cluster   Cluster          up  9000  auto/100000
healthy    false

Node: node2

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
-----
e3a         Cluster   Cluster          up  9000  auto/100000
healthy    false
e3b         Cluster   Cluster          up  9000  auto/100000
healthy    false
```

11. Display information about the status of the nodes in the cluster:

```
cluster show
```

Show example

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

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
node1	true	true	false
node2	true	true	false

12. Disconnect the cable from cluster port e3b on node1, and then connect e3b to port 3 on cluster switch sw2, using the appropriate cabling supported by the SN2100 switches.
13. Disconnect the cable from cluster port e3b on node2, and then connect e3b to port 4 on cluster switch sw2, using the appropriate cabling supported by the SN2100 switches.

Cumulus Linux 4.4.x

14. On switch sw2, enable all node-facing ports.

The following commands enable the node-facing ports on switch sw2:

```
cumulus@sw2:~$ net del interface swp1s0-3, swp2s0-3, swp3-14 link  
down  
cumulus@sw2:~$ net pending  
cumulus@sw2:~$ net commit
```

15. On switch sw2, verify that all ports are up:

```
net show interface all
```



```
cumulus@sw2:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
...						
DN	swp1s0	10G	9216	Trunk/L2		Master: br_default(UP)
DN	swp1s1	10G	9216	Trunk/L2		Master: br_default(UP)
DN	swp1s2	10G	9216	Trunk/L2		Master: br_default(UP)
DN	swp1s3	10G	9216	Trunk/L2		Master: br_default(UP)
DN	swp2s0	25G	9216	Trunk/L2		Master: br_default(UP)
DN	swp2s1	25G	9216	Trunk/L2		Master: br_default(UP)
DN	swp2s2	25G	9216	Trunk/L2		Master: br_default(UP)
DN	swp2s3	25G	9216	Trunk/L2		Master: br_default(UP)
UP	swp3	100G	9216	Trunk/L2	node1 (e3b)	Master: br_default(UP)
UP	swp4	100G	9216	Trunk/L2	node2 (e3b)	Master: br_default(UP)
...						
...						
UP	swp15	100G	9216	BondMember	swp15	Master: cluster_isl(UP)
UP	swp16	100G	9216	BondMember	swp16	Master: cluster_isl(UP)
...						

16. On both switches sw1 and sw2, verify that both nodes each have one connection to each switch:

```
net show lldp
```

The following example shows the appropriate results for both switches sw1 and sw2:

```
cumulus@sw1:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
swp3	100G	Trunk/L2	node1	e3a
swp4	100G	Trunk/L2	node2	e3a
swp15	100G	BondMember	sw2	swp15
swp16	100G	BondMember	sw2	swp16

```
cumulus@sw2:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
swp3	100G	Trunk/L2	node1	e3b
swp4	100G	Trunk/L2	node2	e3b
swp15	100G	BondMember	sw1	swp15
swp16	100G	BondMember	sw1	swp16

Cumulus Linux 5.x

14. On switch sw2, enable all node-facing ports.

The following commands enable the node-facing ports on switch sw2:

```
cumulus@sw2:~$ nv set interface swp1s0-3,swp2s0-3,swp3-14 link state  
up  
cumulus@sw2:~$ nv config apply  
cumulus@sw2:~$ nv config save
```

15. On switch sw2, verify that all ports are up:

```
nv show interface
```

```
cumulus@sw2:~$ nv show interface
```

Interface	State	Speed	MTU	Type	Remote Host
Remote Port	Summary				
-----	-----	-----	-----	-----	-----
-----	-----	-----	-----	-----	-----
...					
...					
swp1s0	up	10G	9216	swp	odq-a300-1a
e0a					
swp1s1	up	10G	9216	swp	odq-a300-1b
e0a					
swp1s2	down	10G	9216	swp	
swp1s3	down	10G	9216	swp	
swp2s0	down	25G	9216	swp	
swp2s1	down	25G	9216	swp	
swp2s2	down	25G	9216	swp	
swp2s3	down	25G	9216	swp	
swp3	down		9216	swp	
swp4	down		9216	swp	
...					
...					
swp14	down		9216	swp	
swp15	up	100G	9216	swp	ossq-int-rcf10
swp15					
swp16	up	100G	9216	swp	ossq-int-rcf10
swp16					

16. On both switches sw1 and sw2, verify that both nodes each have one connection to each switch:

```
nv show interface --view=lldp
```

The following examples show the appropriate results for both switches sw1 and sw2:

```
cumulus@sw1:~$ nv show interface --view=lldp
```

Interface	Speed	Type	Remote Host
Remote Port			
-----	-----	-----	-----
-----	-----	-----	-----
...			
...			
swp1s0	10G	swp	odq-a300-1a
e0a			
swp1s1	10G	swp	odq-a300-1b

```

e0a
swp1s2      10G    swp
swp1s3      10G    swp
swp2s0      25G    swp
swp2s1      25G    swp
swp2s2      25G    swp
swp2s3      25G    swp
swp3                swp
swp4                swp
...
...
swp14                swp
swp15      100G    swp      ossg-int-rcf10
swp15
swp16      100G    swp      ossg-int-rcf10
swp16

```

```
cumulus@sw2:~$ nv show interface --view=lldp
```

```

Interface      Speed  Type      Remote Host
Remote Port
-----
-----
...
...
swp1s0      10G    swp      odq-a300-1a
e0a
swp1s1      10G    swp      odq-a300-1b
e0a
swp1s2      10G    swp
swp1s3      10G    swp
swp2s0      25G    swp
swp2s1      25G    swp
swp2s2      25G    swp
swp2s3      25G    swp
swp3                swp
swp4                swp
...
...
swp14                swp
swp15      100G    swp      ossg-int-rcf10
swp15
swp16      100G    swp      ossg-int-rcf10
swp16

```

17. Display information about the discovered network devices in your cluster:

```
network device-discovery show -protocol lldp
```

Show example

```
cluster1::*> network device-discovery show -protocol lldp
Node/          Local  Discovered
Protocol       Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1          /lldp
               e3a   sw1 (b8:ce:f6:19:1a:7e)   swp3       -
               e3b   sw2 (b8:ce:f6:19:1b:96)   swp3       -
node2          /lldp
               e3a   sw1 (b8:ce:f6:19:1a:7e)   swp4       -
               e3b   sw2 (b8:ce:f6:19:1b:96)   swp4       -
```

18. Verify that all cluster ports are up:

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

Show example

The following example shows that all of the cluster ports are up on node1 and node2:

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

Node: node1

Ignore

Health
Speed(Mbps) Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e3a       Cluster      Cluster      up    9000  auto/10000
healthy  false
e3b       Cluster      Cluster      up    9000  auto/10000
healthy  false

Node: node2

Ignore

Health
Speed(Mbps) Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e3a       Cluster      Cluster      up    9000  auto/10000
healthy  false
e3b       Cluster      Cluster      up    9000  auto/10000
healthy  false
```

Step 3: Verify the configuration

1. Enable auto-revert on all cluster LIFs:

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

Show example

```
cluster1::*> net interface modify -vserver Cluster -lif * -auto
-revert true
```

Vserver	Logical Interface	Auto-revert
Cluster	node1_clus1	true
	node1_clus2	true
	node2_clus1	true
	node2_clus2	true

2. Verify that all interfaces display true for Is Home:

```
net interface show -vserver Cluster
```



This might take a minute to complete.

Show example

The following example shows that all LIFs are up on node1 and node2 and that Is Home results are true:

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

Vserver	Logical Interface	Status	Network	Current	Port
Cluster	node1_clus1	up/up	169.254.209.69/16	node1	e3a
true	node1_clus2	up/up	169.254.49.125/16	node1	e3b
true	node2_clus1	up/up	169.254.47.194/16	node2	e3a
true	node2_clus2	up/up	169.254.19.183/16	node2	e3b
true					

3. Verify that the settings are disabled:

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

Show example

The false output in the following example shows that the configuration settings are disabled:

```
cluster1::*> network options switchless-cluster show  
Enable Switchless Cluster: false
```

4. Verify the status of the node members in the cluster:

```
cluster show
```

Show example

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

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
node1	true	true	false
node2	true	true	false

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node1
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e3a
Cluster node1_clus2 169.254.49.125 node1 e3b
Cluster node2_clus1 169.254.47.194 node2 e3a
Cluster node2_clus2 169.254.19.183 node2 e3b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

6. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

What's next?

[Configure switch health monitoring.](#)

Replace switches

Replace a NVIDIA SN2100 cluster switch

Follow this procedure to replace a defective NVIDIA SN2100 switch in a cluster network. This is a nondisruptive procedure (NDU).

Review requirements

Existing cluster and network infrastructure

Ensure that:

- The existing cluster are verified as completely functional, with at least one fully connected cluster switch.
- All cluster ports are up.
- All cluster logical interfaces (LIFs) are up and on their home ports.
- The ONTAP `cluster ping-cluster -node node1` command indicates that basic connectivity and larger than PMTU communication are successful on all paths.

NVIDIA SN2100 replacement switch

Ensure that:

- Management network connectivity on the replacement switch are functional.
- Console access to the replacement switch are in place.
- The node connections are ports swp1 through swp14.
- All Inter-Switch Link (ISL) ports are disabled on ports swp15 and swp16.
- The desired reference configuration file (RCF) and Cumulus operating system image switch are loaded onto the switch.
- Initial customization of the switch is complete.

Also make sure that any previous site customizations, such as STP, SNMP, and SSH, are copied to the new switch.



You must execute the command for migrating a cluster LIF from the node where the cluster LIF is hosted.

Enable console logging

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

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

Replace the switch

About the examples

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

- The names of the existing NVIDIA SN2100 switches are *sw1* and *sw2*.
- The name of the new NVIDIA SN2100 switch is *nsw2*.
- The node names are *node1* and *node2*.
- The cluster ports on each node are named *e3a* and *e3b*.

- The cluster LIF names are *node1_clus1* and *node1_clus2* for node1, and *node2_clus1* and *node2_clus2* for node2.
- The prompt for changes to all cluster nodes is `cluster1::*>`
- Breakout ports take the format: `swp[port]s[breakout port 0-3]`. For example, four breakout ports on `swp1` are *swp1s0*, *swp1s1*, *swp1s2*, and *swp1s3*.

About the cluster network topology

This procedure is based on the following cluster network topology:

Show example topology

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

```
Node: node1
```

```
Ignore
```

Health	Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Speed (Mbps)	Health Status
	e3a	Cluster	Cluster		up	9000	auto/100000		healthy
false									
	e3b	Cluster	Cluster		up	9000	auto/100000		healthy
false									

```
Node: node2
```

```
Ignore
```

Health	Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Speed (Mbps)	Health Status
	e3a	Cluster	Cluster		up	9000	auto/100000		healthy
false									
	e3b	Cluster	Cluster		up	9000	auto/100000		healthy
false									

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

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
Cluster	node1_clus1	up/up	169.254.209.69/16	node1	e3a
true	node1_clus2	up/up	169.254.49.125/16	node1	e3b
true					

```

node2_clus1 up/up 169.254.47.194/16 node2 e3a
true
node2_clus2 up/up 169.254.19.183/16 node2 e3b
true

```

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

```

Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1      /lldp
           e3a    sw1 (b8:ce:f6:19:1a:7e)   swp3       -
           e3b    sw2 (b8:ce:f6:19:1b:96)   swp3       -
node2      /lldp
           e3a    sw1 (b8:ce:f6:19:1a:7e)   swp4       -
           e3b    sw2 (b8:ce:f6:19:1b:96)   swp4       -

```

+

```
cumulus@sw1:~$ net show lldp
```

```

LocalPort  Speed  Mode          RemoteHost      RemotePort
-----
swp3       100G   Trunk/L2      sw2             e3a
swp4       100G   Trunk/L2      sw2             e3a
swp15      100G   BondMember    sw2             swp15
swp16      100G   BondMember    sw2             swp16

```

```
cumulus@sw2:~$ net show lldp
```

```

LocalPort  Speed  Mode          RemoteHost      RemotePort
-----
swp3       100G   Trunk/L2      sw1             e3b
swp4       100G   Trunk/L2      sw1             e3b
swp15      100G   BondMember    sw1             swp15
swp16      100G   BondMember    sw1             swp16

```

Step 1: Prepare for replacement

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

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

where x is the duration of the maintenance window in hours.

2. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (***>**) appears.

3. Install the appropriate RCF and image on the switch, nsw2, and make any necessary site preparations.

If necessary, verify, download, and install the appropriate versions of the RCF and Cumulus software for the new switch.

- a. You can download the applicable Cumulus software for your cluster switches from the *NVIDIA Support* site. Follow the steps on the Download page to download the Cumulus Linux for the version of ONTAP software you are installing.
- b. The appropriate RCF is available from the *NVIDIA Cluster and Storage Switches* page. Follow the steps on the Download page to download the correct RCF for the version of ONTAP software you are installing.

Step 2: Configure ports and cabling

1. On the new switch nsw2, log in as admin and shut down all of the ports that will be connected to the node cluster interfaces (ports swp1 to swp14).

The LIFs on the cluster nodes should have already failed over to the other cluster port for each node.

Show example

```
cumulus@nsw2:~$ net add interface swp1s0-3, swp2s0-3, swp3-14 link
down
cumulus@nsw2:~$ net pending
cumulus@nsw2:~$ net commit
```

2. Disable auto-revert on the cluster LIFs:

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

Show example

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

```
Warning: Disabling the auto-revert feature of the cluster logical
interface may effect the availability of your cluster network. Are
you sure you want to continue? {y|n}: y
```

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

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

4. Shut down the ISL ports swp15 and swp16 on the SN2100 switch sw1.

Show example

```
cumulus@sw1:~$ net add interface swp15-16 link down
cumulus@sw1:~$ net pending
cumulus@sw1:~$ net commit
```

5. Remove all the cables from the SN2100 sw1 switch, and then connect them to the same ports on the SN2100 nsw2 switch.
6. Bring up the ISL ports swp15 and swp16 between the sw1 and nsw2 switches.

Show example

The following commands enable ISL ports swp15 and swp16 on switch sw1:

```
cumulus@sw1:~$ net del interface swp15-16 link down
cumulus@sw1:~$ net pending
cumulus@sw1:~$ net commit
```

The following example shows that the ISL ports are up on switch sw1:

```
cumulus@sw1:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp15	100G	9216	BondMember	nsw2 (swp15)	Master: cluster_isl (UP)
UP	swp16	100G	9216	BondMember	nsw2 (swp16)	Master: cluster_isl (UP)

+

The following example shows that the ISL ports are up on switch nsw2:

+

```
cumulus@nsw2:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp15	100G	9216	BondMember	sw1 (swp15)	Master: cluster_isl (UP)
UP	swp16	100G	9216	BondMember	sw1 (swp16)	Master: cluster_isl (UP)

7. Verify that port e3b is up on all nodes:

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

Show example

The output should be similar to the following:

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

Node: node1

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU   Admin/Oper  Speed (Mbps)
Status      Status
-----
e3a         Cluster   Cluster           up   9000  auto/100000
healthy    false
e3b         Cluster   Cluster           up   9000  auto/100000
healthy    false

Node: node2

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU   Admin/Oper  Speed (Mbps)
Status      Status
-----
e3a         Cluster   Cluster           up   9000  auto/100000
healthy    false
e3b         Cluster   Cluster           up   9000  auto/100000
healthy    false
```

8. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

Show example

```
cluster1::~*> network device-discovery show -protocol lldp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1      /lldp
           e3a    sw1  (b8:ce:f6:19:1a:7e)      swp3       -
           e3b    nsw2 (b8:ce:f6:19:1b:b6)      swp3       -
node2      /lldp
           e3a    sw1  (b8:ce:f6:19:1a:7e)      swp4       -
           e3b    nsw2 (b8:ce:f6:19:1b:b6)      swp4       -
```

9. Verify that all node cluster ports are up:

```
net show interface
```

Show example

```
cumulus@nsw2:~$ net show interface

State Name           Spd   MTU   Mode           LLDP
Summary
-----
...
...
UP      swp3                100G  9216  Trunk/L2
Master: bridge(UP)
UP      swp4                100G  9216  Trunk/L2
Master: bridge(UP)
UP      swp15               100G  9216  BondMember  sw1 (swp15)
Master: cluster_isl(UP)
UP      swp16               100G  9216  BondMember  sw1 (swp16)
Master: cluster_isl(UP)
```

10. Verify that both nodes each have one connection to each switch:

```
net show lldp
```

Show example

The following example shows the appropriate results for both switches:

```
cumulus@sw1:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
swp3	100G	Trunk/L2	node1	e3a
swp4	100G	Trunk/L2	node2	e3a
swp15	100G	BondMember	nsw2	swp15
swp16	100G	BondMember	nsw2	swp16

```
cumulus@nsw2:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
swp3	100G	Trunk/L2	node1	e3b
swp4	100G	Trunk/L2	node2	e3b
swp15	100G	BondMember	sw1	swp15
swp16	100G	BondMember	sw1	swp16

11. Enable auto-revert on the cluster LIFs:

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

12. On switch nsw2, bring up the ports connected to the network ports of the nodes.

Show example

```
cumulus@nsw2:~$ net del interface swp1-14 link down  
cumulus@nsw2:~$ net pending  
cumulus@nsw2:~$ net commit
```

13. Display information about the nodes in a cluster:

```
cluster show
```

Show example

This example shows that the node health for node1 and node2 in this cluster is true:

```
cluster1::*> cluster show
```

Node	Health	Eligibility
node1	true	true
node2	true	true

14. Verify that all physical cluster ports are up:

```
network port show ipspace Cluster
```

Show example

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

Node node1
Ignore
Speed (Mbps)
Health Health
Port IPspace Broadcast Domain Link MTU Admin/Oper
Status Status
-----
e3a Cluster Cluster up 9000 auto/10000
healthy false
e3b Cluster Cluster up 9000 auto/10000
healthy false

Node: node2

Ignore
Speed (Mbps)
Health Health
Port IPspace Broadcast Domain Link MTU Admin/Oper
Status Status
-----
e3a Cluster Cluster up 9000 auto/10000
healthy false
e3b Cluster Cluster up 9000 auto/10000
healthy false
```

Step 3: Verify the configuration

1. Verify that the cluster network is healthy.

Show example

```
cumulus@sw1:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
swp3	100G	Trunk/L2	node1	e3a
swp4	100G	Trunk/L2	node2	e3a
swp15	100G	BondMember	nsw2	swp15
swp16	100G	BondMember	nsw2	swp16

2. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

What's next?

[Configure switch health monitoring.](#)

Replace NVIDIA SN2100 cluster switches with switchless connections

You can migrate from a cluster with a switched cluster network to one where two nodes are directly connected for ONTAP 9.3 and later.

Review requirements

Guidelines

Review the following guidelines:

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


What you'll need

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

Migrate the switches

About this task

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



About the examples

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

Step 1: Prepare for migration

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

```
set -privilege advanced
```

The advanced prompt `*>` appears.

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

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

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

Show example

The following example output shows if the option is enabled.

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

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

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

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

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

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

Show example


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

Step 2: Configure ports and cabling

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

```
network port show -ipspace Cluster
```

In the following example for nodes with cluster ports "e0a" and "e0b", one group is identified as "node1:e0a" and "node2:e0a" and the other group as "node1:e0b" and "node2:e0b". Your nodes might be using different cluster ports because they vary by system.



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

Show example

```
cluster::> network port show -ipSPACE Cluster
Node: node1

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

Node: node2

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

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

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

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

Show example

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

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

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

4. Disable auto-revert for the cluster LIFs:

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

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

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

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

Show example

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

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

6. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

7. Verify that the cluster is healthy:

```
cluster ring show
```

All units must be either master or secondary.


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



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


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

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



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

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



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

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

The following example shows that the switchless cluster is enabled:

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

10. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```



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


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



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


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

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



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

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



Step 3: Verify the configuration

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

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

Show example

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

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

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

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

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

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

Show example

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

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

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

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

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

```
cluster show
```

Show example

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

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

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

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

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

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

7. Change the privilege level back to admin:

```
set -privilege admin
```

Storage switches

Cisco Nexus 9336C-FX2

Overview

Overview of installation and configuration for Cisco Nexus 9336C-FX2 storage switches

The Cisco Nexus 9336C-FX2 switch is part of the Cisco Nexus 9000 platform and can be installed in a NetApp system cabinet. The Nexus 9336C-FX2 is an Ethernet switch that allows you to switch data between controllers and disk shelves.

Initial configuration overview

To initially configure a Cisco Nexus 9336C-FX2 switch on systems running ONTAP, follow these steps:

1. [Complete cabling worksheet.](#)
2. [Install the switch.](#)
3. [Configure the switch.](#)
4. [Install the switch in a NetApp cabinet.](#)
5. [Prepare to install NX-OS and RCF.](#)
6. [Install the NX-OS software.](#)
7. [Install the RCF config file.](#)

Additional information

Before you begin installation or maintenance, be sure to review the following:

- [Configuration requirements](#)
- [Components and part numbers](#)
- [Required documentation](#)
- [Smart Call Home requirements](#)

Configuration requirements for Cisco Nexus 9336C-FX2 storage switches

For Cisco Nexus 9336C-FX2 switch installation and maintenance, be sure to review configuration and network requirements.

ONTAP support

From ONTAP 9.9.1, you can use Cisco Nexus 9336C-FX2 switches to combine storage and cluster functionality into a shared switch configuration.

If you want to build ONTAP clusters with more than two nodes, you need two supported network switches.

Configuration requirements

For configuration, you need the appropriate number and type of cables and cable connectors for your switches.

Depending on the type of switch you are initially configuring, you need to connect to the switch console port with the included console cable; you also need to provide specific network information.

Network requirements

You need the following network information for all switch configurations.

- IP subnet for management network traffic
- Host names and IP addresses for each of the storage system controllers and all applicable switches
- Most storage system controllers are managed through the e0M interface by connecting to the Ethernet service port (wrench icon). On AFF A800 and AFF A700s systems, the e0M interface uses a dedicated Ethernet port.
- Refer to the [Hardware Universe](#) for the latest information.

For more information about the initial configuration of your switch, see the following guide: [Cisco Nexus 9336C-FX2 Installation and Upgrade Guide](#).

Components and part numbers for Cisco Nexus 9336C-FX2 storage switches

For Cisco Nexus 9336C-FX2 switch installation and maintenance, be sure to review the list of components and part numbers.

The following table lists the part number and description for the 9336C-FX2 switch, fans, and power supplies:

Part number	Description
X190200-CS-PE	N9K-9336C-FX2, CS, PTSX, 36PT10/25/40/100GQSFP28
X190200-CS-PI	N9K-9336C-FX2, CS, PSIN, 36PT10/25/40/100GQSFP28
X190210-FE-PE	N9K-9336C, FTE, PTSX, 36PT10/25/40/100GQSFP28
X190210-FE-PI	N9K-9336C, FTE, PSIN, 36PT10/25/40/100GQSFP28
X190002	Accessory Kit X190001/X190003
X-NXA-PAC-1100W-PE2	N9K-9336C AC 1100W PSU - Port side exhaust airflow
X-NXA-PAC-1100W-PI2	N9K-9336C AC 1100W PSU - Port side Intake airflow
X-NXA-FAN-65CFM-PE	N9K-9336C 65CFM, Port side exhaust airflow
X-NXA-FAN-65CFM-PI	N9K-9336C 65CFM, Port side intake airflow

Documentation requirements for Cisco Nexus 9336C-FX2 storage switches

For Cisco Nexus 9336C-FX2 switch installation and maintenance, be sure to review specific switch and controller documentation to set up your Cisco 9336C-FX2 switches and ONTAP cluster.

Switch documentation

To set up the Cisco Nexus 9336C-FX2 switches, you need the following documentation from the [Cisco Nexus 9000 Series Switches Support](#) page:

Document title	Description
<i>Nexus 9000 Series Hardware Installation Guide</i>	Provides detailed information about site requirements, switch hardware details, and installation options.
<i>Cisco Nexus 9000 Series Switch Software Configuration Guides</i> (choose the guide for the NX-OS release installed on your switches)	Provides initial switch configuration information that you need before you can configure the switch for ONTAP operation.
<i>Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide</i> (choose the guide for the NX-OS release installed on your switches)	Provides information on how to downgrade the switch to ONTAP supported switch software, if necessary.
<i>Cisco Nexus 9000 Series NX-OS Command Reference Master Index</i>	Provides links to the various command references provided by Cisco.
<i>Cisco Nexus 9000 MIBs Reference</i>	Describes the Management Information Base (MIB) files for the Nexus 9000 switches.
<i>Nexus 9000 Series NX-OS System Message Reference</i>	Describes the system messages for Cisco Nexus 9000 series switches, those that are informational, and others that might help diagnose problems with links, internal hardware, or the system software.
<i>Cisco Nexus 9000 Series NX-OS Release Notes</i> (choose the notes for the NX-OS release installed on your switches)	Describes the features, bugs, and limitations for the Cisco Nexus 9000 Series.
Regulatory Compliance and Safety Information for Cisco Nexus 9000 Series	Provides international agency compliance, safety, and statutory information for the Nexus 9000 series switches.

ONTAP systems documentation

To set up an ONTAP system, you need the following documents for your version of the operating system from the [ONTAP 9 Documentation Center](#).

Name	Description
Controller-specific <i>Installation and Setup Instructions</i>	Describes how to install NetApp hardware.
ONTAP documentation	Provides detailed information about all aspects of the ONTAP releases.
Hardware Universe	Provides NetApp hardware configuration and compatibility information.

Rail kit and cabinet documentation

To install a Cisco 9336-FX2 switch in a NetApp cabinet, see the following hardware documentation.

Name	Description
42U System Cabinet, Deep Guide	Describes the FRUs associated with the 42U system cabinet, and provides maintenance and FRU replacement instructions.
Install a Cisco 9336-FX2 switch in a NetApp Cabinet	Describes how to install a Cisco Nexus 9336C-FX2 switch in a four-post NetApp cabinet.

Smart Call Home requirements

To use Smart Call Home feature, review the following guidelines.

Smart Call Home monitors the hardware and software components on your network. When a critical system configuration occurs, it generates an email-based notification and raises an alert to all the recipients that are configured in your destination profile. To use Smart Call Home, you must configure a cluster network switch to communicate using email with the Smart Call Home system. In addition, you can optionally set up your cluster network switch to take advantage of Cisco's embedded Smart Call Home support feature.

Before you can use Smart Call Home, be aware of the following considerations:

- An email server must be in place.
- The switch must have IP connectivity to the email server.
- The contact name (SNMP server contact), phone number, and street address information must be configured. This is required to determine the origin of messages received.
- A CCO ID must be associated with an appropriate Cisco SMARTnet Service contract for your company.
- Cisco SMARTnet Service must be in place for the device to be registered.

The [Cisco support site](#) contains information about the commands to configure Smart Call Home.

Install hardware

Complete the Cisco Nexus 9336C-FX2 cabling worksheet

If you want to document the supported platforms, download a PDF of this page and complete the cabling worksheet.

The sample cabling worksheet provides examples of recommended port assignments from the switches to the controllers. The blank worksheet provides a template that you can use in setting up your cluster.

Sample cabling worksheet

The sample port definition on each pair of switches is as follows:

Cluster switch A		Cluster switch B	
Switch port	Node and port usage	Switch port	Node and port usage
1	4x100GbE node 1	1	4x100GbE node 1
2	4x100GbE node 2	2	4x100GbE node 2
3	4x100GbE node 3	3	4x100GbE node 3
4	4x100GbE node 4	4	4x100GbE node 4
5	4x100GbE node 5	5	4x100GbE node 5
6	4x100GbE node 6	6	4x100GbE node 6
7	4x100GbE node 7	7	4x100GbE node 7
8	4x100GbE node 8	8	4x100GbE node 8
9	4x100GbE node 9	9	4x100GbE node 9
10	4x100GbE node 10	10	4x100GbE node 10
11	4x100GbE node 11	11	4x100GbE node 11
12	4x100GbE node 12	12	4x100GbE node 12
13	4x100GbE node 13	13	4x100GbE node 13
14	4x100GbE node 14	14	4x100GbE node 14
15	4x100GbE node 15	15	4x100GbE node 15
16	4x100GbE node 16	16	4x100GbE node 16
17	4x100GbE node 17	17	4x100GbE node 17
18	4x100GbE node 18	18	4x100GbE node 18

Cluster switch A		Cluster switch B	
19	4x100GbE node 19	19	4x100GbE node 19
20	4x100GbE node 20	20	4x100GbE node 20
21	4x100GbE node 21	21	4x100GbE node 21
22	4x100GbE node 22	22	4x100GbE node 22
23	4x100GbE node 23	23	4x100GbE node 23
24	4x100GbE node 24	24	4x100GbE node 24
25	4x100GbE node 25	25	4x100GbE node 25
26	4x100GbE node 26	26	4x100GbE node 26
27	4x100GbE node 27	27	4x100GbE node 27
28	4x100GbE node 28	28	4x100GbE node 28
29	4x100GbE node 29	29	4x100GbE node 29
30	4x100GbE node 30	30	4x100GbE node 30
31	4x100GbE node 31	31	4x100GbE node 31
32	4x100GbE node 32	32	4x100GbE node 32
33	4x100GbE node 33	33	4x100GbE node 33
30	4x100GbE node 30	30	4x100GbE node 33
34	4x100GbE node 34	34	4x100GbE node 34
35	4x100GbE node 35	35	4x100GbE node 35
36	4x100GbE node 36	36	4x100GbE node 36

Blank cabling worksheet

You can use the blank cabling worksheet to document the platforms that are supported as nodes in a cluster. The *Supported Cluster Connections* section of the [Hardware Universe](#) defines the cluster ports used by the platform.

Cluster switch A		Cluster switch B	
1		1	
2		2	
3		3	
4		4	
5		5	
6		6	
7		7	
8		8	
9		9	
10		10	
11		11	
12		12	
13		13	
14		14	
15		15	
16		16	
17		17	
18		18	
19		19	
20		20	
21		21	
22		22	

Cluster switch A		Cluster switch B	
23		23	
24		24	
25		25	
26		26	
27		27	
28		28	
29		29	
30		30	
31		31	
32		32	
33		33	
34		34	
35		35	
36		36	

Install the 9336C-FX2 storage switch

Follow this procedure to install the Cisco Nexus 9336C-FX2 storage switch.

What you'll need

- Access to an HTTP, FTP or TFTP server at the installation site to download the applicable NX-OS and reference configuration file (RCF) releases.
- Applicable NX-OS version, downloaded from the [Cisco Software Download](#) page.
- Applicable licenses, network and configuration information, and cables.
- Completed [cabling worksheets](#).
- Applicable NetApp cluster network and management network RCFs downloaded from the NetApp Support Site at mysupport.netapp.com. All Cisco cluster network and management network switches arrive with the standard Cisco factory-default configuration. These switches also have the current version of the NX-OS software but do not have the RCFs loaded.
- Required switch documentation. See [Required documentation](#) for more information.

Steps

1. Rack the cluster network and management network switches and controllers.

If you are installing your...	Then...
Cisco Nexus 9336C-FX2 in a NetApp system cabinet	See Install switch in NetApp cabinet for instructions to install the switch in a NetApp cabinet.
Equipment in a Telco rack	See the procedures provided in the switch hardware installation guides and the NetApp installation and setup instructions.

2. Cable the cluster network and management network switches to the controllers using the completed cabling worksheets.
3. Power on the cluster network and management network switches and controllers.

What's next?

Go to [Configure Cisco Nexus 9336C-FX2 storage switch](#).

Configure the 9336C-FX2 storage switch

Follow this procedure to configure the Cisco Nexus 9336C-FX2 switch.

What you'll need


- Access to an HTTP, FTP or TFTP server at the installation site to download the applicable NX-OS and reference configuration file (RCF) releases.
- Applicable NX-OS version, downloaded from the [Cisco software download](#) page.
- Applicable licenses, network and configuration information, and cables.
- Completed [cabling worksheets](#).
- Applicable NetApp cluster network and management network RCFs downloaded from the NetApp Support Site at mysupport.netapp.com. All Cisco cluster network and management network switches arrive with the standard Cisco factory-default configuration. These switches also have the current version of the NX-OS software but do not have the RCFs loaded.
- Required switch documentation. See [Required documentation](#) for more information.


Steps

1. Perform an initial configuration of the cluster network switches.

Provide applicable responses to the following initial setup questions when you first boot the switch. Your site's security policy defines the responses and services to enable.

Prompt	Response
Abort Auto Provisioning and continue with normal setup? (yes/no)	Respond with yes . The default is no.
Do you want to enforce secure password standard? (yes/no)	Respond with yes . The default is yes.

Prompt	Response
Enter the password for admin.	The default password is “admin”; you must create a new, strong password. A weak password can be rejected.
Would you like to enter the basic configuration dialog? (yes/no)	Respond with yes at the initial configuration of the switch.
Create another login account? (yes/no)	Your answer depends on your site’s policies on alternate administrators. The default is no .
Configure read-only SNMP community string? (yes/no)	Respond with no . The default is no.
Configure read-write SNMP community string? (yes/no)	Respond with no . The default is no.
Enter the switch name.	The switch name is limited to 63 alphanumeric characters.
Continue with Out-of-band (mgmt0) management configuration? (yes/no)	Respond with yes (the default) at that prompt. At the mgmt0 IPv4 address: prompt, enter your IP address: ip_address.
Configure the default-gateway? (yes/no)	Respond with yes . At the IPv4 address of the default-gateway: prompt, enter your default_gateway.
Configure advanced IP options? (yes/no)	Respond with no . The default is no.
Enable the telnet service? (yes/no)	Respond with no . The default is no.
Enabled SSH service? (yes/no)	Respond with yes . The default is yes. <div style="border: 1px solid gray; padding: 5px; display: inline-block;">  SSH is recommended when using Ethernet Switch Health Monitor (CSHM) for its log collection features. SSHv2 is also recommended for enhanced security. </div>
Enter the type of SSH key you want to generate (dsa/rsa/rsa1).	The default is rsa .
Enter the number of key bits (1024-2048).	Enter the number of key bits from 1024 to 2048.
Configure the NTP server? (yes/no)	Respond with no . The default is no.

Prompt	Response
Configure default interface layer (L3/L2)	Respond with L2 . The default is L2.
Configure default switch port interface state (shut/noshut)	Respond with noshut . The default is noshut.
Configure CoPP system profile (strict/moderate/lenient/dense)	Respond with strict . The default is strict.
Would you like to edit the configuration? (yes/no)	You should see the new configuration at this point. Review and make any necessary changes to the configuration you just entered. Respond with no at the prompt if you are satisfied with the configuration. Respond with yes if you want to edit your configuration settings.
Use this configuration and save it? (yes/no)	Respond with yes to save the configuration. This automatically updates the kickstart and system images. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  If you do not save the configuration at this stage, none of the changes will be in effect the next time you reboot the switch. </div>

2. Verify the configuration choices you made in the display that appears at the end of the setup, and make sure that you save the configuration.
3. Check the version on the cluster network switches, and if necessary, download the NetApp-supported version of the software to the switches from the [Cisco software download](#) page.

What's next?

Optionally, you can [install a Cisco Nexus 9336C-FX2 switch in a NetApp cabinet](#). Otherwise, go to [Prepare to install NX-OS and RCF](#).

Install a Cisco Nexus 9336C-FX2 switch in a NetApp cabinet

Depending on your configuration, you might need to install the Cisco Nexus 9336C-FX2 switch and pass-through panel in a NetApp cabinet. Standard brackets are included with the switch.

What you'll need

- 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 cabinet posts.
- You must use the Cisco standard rail kit to install the switch in a NetApp cabinet.



The jumper cords are not included with the pass-through kit and should be included with your switches. If they were not shipped with the switches, you can order them from NetApp (part number X1558A-R6).

Required documentation

Review the initial preparation requirements, kit contents, and safety precautions in the [Cisco Nexus 9000 Series Hardware Installation Guide](#).

Steps


1. Install the pass-through blanking panel in the NetApp cabinet.

The pass-through panel kit is available from NetApp (part number X8784-R6).

The NetApp pass-through panel kit contains the following hardware:


- One pass-through blanking panel
- Four 10-32 x .75 screws
- Four 10-32 clip nuts
 - a. Determine the vertical location of the switches and blanking panel in the cabinet.

In this procedure, the blanking panel will be installed in U40.
 - b. Install two clip nuts on each side in the appropriate square holes for front cabinet rails.
 - c. Center the panel vertically to prevent intrusion into adjacent rack space, and then tighten the screws.
 - d. Insert the female connectors of both 48-inch jumper cords from the rear of the panel and through the brush assembly.




(1) Female connector of the jumper cord.

2. Install the rack-mount brackets on the Nexus 9336C-FX2 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.
3. Install the clip nuts in the square hole locations for all four IEA posts.



The two 9336C-FX2 switches will always be mounted in the top 2U of the cabinet RU41 and 42.

4. Install the slider rails in the cabinet.
 - 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 cabinet posts.

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


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

5. Install the switch in the cabinet.



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) As the chassis is pushed toward the rear posts, align the two rear rack-mount guides with the slider rails.

(2) Gently slide the switch until the front rack-mount brackets are flush with the front posts.

b. Attach the switch to the cabinet.



(1) With one person holding the front of the chassis level, the other person should fully tighten the four rear screws to the cabinet 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.



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.

6. When the switches are installed, connect the jumper cords to the switch power inlets.
7. Connect the male plugs of both jumper cords to the closest available PDU outlets.



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

8. Connect the management port on each 9336C-FX2 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.

Configure software

Software install workflow for Cisco Nexus 9336C-FX2 storage switches

To install and configure software for a Cisco Nexus 9336C-FX2 storage switch, follow these steps:

1. [Prepare to install NX-OS and RCF.](#)
2. [Install the NX-OS software.](#)
3. [Install the RCF config file.](#)

Prepare to install NX-OS software and RCF

Before you install the NX-OS software and the Reference Configuration File (RCF), follow this procedure.

About the examples

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.

About this task

The procedure requires the use of both ONTAP commands and Cisco Nexus 9000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

Steps

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=x h` where x is the duration of the maintenance window in hours.



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

2. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (***>**) appears.

3. Display how many cluster interconnect interfaces are configured in each node for each cluster interconnect switch:

```
network device-discovery show -protocol lldp
```

Show example

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

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	

cluster1-02/lldp				
	e0a	cs1	Eth1/2	N9K-
C9336C				
	e0b	cs2	Eth1/2	N9K-
C9336C				
cluster1-01/lldp				
	e0a	cs1	Eth1/1	N9K-
C9336C				
	e0b	cs2	Eth1/1	N9K-
C9336C				

```
4 entries were displayed.
```

4. Check the administrative or operational status of each cluster interface.
 - a. Display the network port attributes:

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

Show example

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

Node: cluster1-02

Ignore

Health Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Speed (Mbps)
Status Status
-----
-----
e0a Cluster Cluster up 9000 auto/100000
healthy false
e0b Cluster Cluster up 9000 auto/100000
healthy false

Node: cluster1-01

Ignore

Health Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Speed (Mbps)
Status Status
-----
-----
e0a Cluster Cluster up 9000 auto/100000
healthy false
e0b Cluster Cluster up 9000 auto/100000
healthy false

4 entries were displayed.
```

b. Display information about the LIFs:

```
network interface show -vserver Cluster
```

Show example

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

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

4 entries were displayed.

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

Packet	Source	Destination
Node	Date	LIF
Loss		
-----	-----	-----
node1		
3/5/2024 19:21:18 -06:00	cluster1-01_clus2	cluster1-02-
clus1 none		
3/5/2024 19:21:20 -06:00	cluster1-01_clus2	cluster1-
02_clus2 none		
node2		
3/5/2024 19:21:18 -06:00	cluster1-02_clus2	cluster1-
01_clus1 none		
3/5/2024 19:21:20 -06:00	cluster1-02_clus2	cluster1-
01_clus2 none		

All ONTAP releases

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

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



```

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

```

6. Verify that the auto-revert command is enabled on all cluster LIFs:

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

Show example

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

Vserver	Logical Interface	Auto-revert
Cluster	cluster1-01_clus1	true
	cluster1-01_clus2	true
	cluster1-02_clus1	true
	cluster1-02_clus2	true

4 entries were displayed.

What's next?

[Install the NX-OS software.](#)

Install the NX-OS software

Follow this procedure to install the NX-OS software on the Nexus 9336C-FX2 cluster switch.

Before you begin, complete the procedure in [Prepare to install NX-OS and RCF](#).

Review requirements

What you'll need

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).

Suggested documentation

- [Cisco Ethernet switch page](#)

Consult the switch compatibility table for the supported ONTAP and NX-OS versions.

- [Software Upgrade and downgrade guides](#)

Refer to the appropriate software and upgrade guides available on the Cisco website for complete documentation on the Cisco switch upgrade and downgrade procedures.

- [Cisco Nexus 9000 and 3000 Upgrade and ISSU Matrix](#)

Provides information on Disruptive Upgrade/Downgrade for Cisco NX-OS software on Nexus 9000 Series Switches based on your current and target releases.

On the page, select **Disruptive Upgrade** and select your current release and target release from the

dropdown list.

About the examples

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.

Install the software

The procedure requires the use of both ONTAP commands and Cisco Nexus 9000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

Steps

1. Connect the cluster switch to the management network.
2. Use the ping command to verify connectivity to the server hosting the NX-OS software and the RCF.

Show example

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.
```

3. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
              e0a    cs1                      Ethernet1/7      N9K-
C9336C-FX2
              e0d    cs2                      Ethernet1/7      N9K-
C9336C-FX2
cluster1-02/cdp
              e0a    cs1                      Ethernet1/8      N9K-
C9336C-FX2
              e0d    cs2                      Ethernet1/8      N9K-
C9336C-FX2
cluster1-03/cdp
              e0a    cs1                      Ethernet1/1/1    N9K-
C9336C-FX2
              e0b    cs2                      Ethernet1/1/1    N9K-
C9336C-FX2
cluster1-04/cdp
              e0a    cs1                      Ethernet1/1/2    N9K-
C9336C-FX2
              e0b    cs2                      Ethernet1/1/2    N9K-
C9336C-FX2
cluster1::*>
```

4. 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
```

Show example

```
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

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
```

Show example

```
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
```

Show example

```
cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                                     Type                                     Address
Model
-----
cs1                                         cluster-network                         10.233.205.90
N9K-C9336C-FX2
  Serial Number: FOCXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                                     9.3(5)
  Version Source: CDP

cs2                                         cluster-network                         10.233.205.91
N9K-C9336C-FX2
  Serial Number: FOCXXXXXXGS
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                                     9.3(5)
  Version Source: CDP
cluster1::*>
```

5. Disable auto-revert on the cluster LIFs. The cluster LIFs fail over to the partner cluster switch and remain there as you perform the upgrade procedure on the targeted switch:

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

6. Copy the NX-OS software and EPLD images to the Nexus 9336C-FX2 switch.

Show example

```
cs2# copy sftp: bootflash: vrf management
Enter source filename: /code/nxos.9.3.5.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.9.3.5.bin /bootflash/nxos.9.3.5.bin
/code/nxos.9.3.5.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.9.3.5.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.9.3.5.img /bootflash/n9000-
epld.9.3.5.img
/code/n9000-epld.9.3.5.img 100% 161MB 9.5MB/s 00:16
sftp> exit
Copy complete, now saving to disk (please wait)...
Copy complete.
```

7. Verify the running version of the NX-OS software:

```
show version
```

Show example

```
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 08.38
  NXOS: version 9.3(4)
  BIOS compile time: 05/29/2020
  NXOS image file is: bootflash:///nxos.9.3.4.bin
  NXOS compile time: 4/28/2020 21:00:00 [04/29/2020 02:28:31]

Hardware
  cisco Nexus9000 C9336C-FX2 Chassis
  Intel(R) Xeon(R) CPU E5-2403 v2 @ 1.80GHz with 8154432 kB of
memory.
  Processor Board ID FOC20291J6K

  Device name: cs2
  bootflash: 53298520 kB
  Kernel uptime is 0 day(s), 0 hour(s), 3 minute(s), 42 second(s)
```

```
Last reset at 157524 usecs after Mon Nov  2 18:32:06 2020
Reason: Reset Requested by CLI command reload
System version: 9.3(4)
Service:

plugin
  Core Plugin, Ethernet Plugin

Active Package(s) :

cs2#
```

8. Install the NX-OS image.

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

Show example

```
cs2# install all nxos bootflash:nxos.9.3.5.bin
```

```
Installer will perform compatibility check first. Please wait.  
Installer is forced disruptive
```

```
Verifying image bootflash:/nxos.9.3.5.bin for boot variable "nxos".  
[] 100% -- SUCCESS
```

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

```
Preparing "nxos" version info using image bootflash:/nxos.9.3.5.bin.  
[] 100% -- SUCCESS
```

```
Preparing "bios" version info using image bootflash:/nxos.9.3.5.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(4)	9.3(5)
yes			
1	bios	v08.37(01/28/2020):v08.23(09/23/2015)	
v08.38(05/29/2020)		yes	

```
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.
```

9. Verify the new version of NX-OS software after the switch has rebooted:

```
show version
```

Show example

```
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 05.33
NXOS: version 9.3(5)
BIOS compile time: 09/08/2018
NXOS image file is: bootflash:///nxos.9.3.5.bin
NXOS compile time: 11/4/2018 21:00:00 [11/05/2018 06:11:06]
```

Hardware

```
cisco Nexus9000 C9336C-FX2 Chassis
Intel(R) Xeon(R) CPU E5-2403 v2 @ 1.80GHz with 8154432 kB of
memory.
Processor Board ID FOC20291J6K
```

```
Device name: cs2
bootflash: 53298520 kB
Kernel uptime is 0 day(s), 0 hour(s), 3 minute(s), 42 second(s)
```

```
Last reset at 277524 usecs after Mon Nov  2 22:45:12 2020
```

```
Reason: Reset due to upgrade
```

```
System version: 9.3(4)
```

```
Service:
```

```
plugin
```

```
Core Plugin, Ethernet Plugin
```

```
Active Package(s):
```

10. Upgrade the EPLD image and reboot the switch.

Show example




```
cs2# show version module 1 epld
```

EPLD Device	Version
MI FPGA	0x7
IO FPGA	0x17
MI FPGA2	0x2
GEM FPGA	0x2
GEM FPGA	0x2
GEM FPGA	0x2
GEM FPGA	0x2

```
cs2# install epld bootflash:n9000-epld.9.3.5.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	0x07	0x07	No
1	SUP	IO FPGA	0x17	0x19	Yes
1	SUP	MI FPGA2	0x02	0x02	No

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

EPLDs upgraded.

Module 1 EPLD upgrade is successful.

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

Show example

```
cs2# show version module 1 epld
```

EPLD	Device	Version
MI	FPGA	0x7
IO	FPGA	0x19
MI	FPGA2	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2

12. Verify the health of cluster ports on the cluster.

a. Verify that cluster ports are up and healthy across all nodes in the cluster:

```
network port show -role cluster
```

Show example

```
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

Health Health Speed (Mbps)
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.

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
          e0a    cs1                        Ethernet1/7
N9K-C9336C-FX2
          e0d    cs2                        Ethernet1/7
N9K-C9336C-FX2
cluster01-2/cdp
          e0a    cs1                        Ethernet1/8
N9K-C9336C-FX2
          e0d    cs2                        Ethernet1/8
N9K-C9336C-FX2
cluster01-3/cdp
          e0a    cs1                        Ethernet1/1/1
N9K-C9336C-FX2
          e0b    cs2                        Ethernet1/1/1
N9K-C9336C-FX2
cluster1-04/cdp
          e0a    cs1                        Ethernet1/1/2
N9K-C9336C-FX2
          e0b    cs2                        Ethernet1/1/2
N9K-C9336C-FX2

cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                Type                Address
Model
-----
-----
cs1                    cluster-network    10.233.205.90
N9K-C9336C-FX2
  Serial Number: FOCXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                    9.3(5)
  Version Source: CDP

cs2                    cluster-network    10.233.205.91
```

```

N9K-C9336C-FX2
  Serial Number: FOCXXXXXXGS
    Is Monitored: true
      Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                    9.3(5)
  Version Source: CDP

2 entries were displayed.

```

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.

```

13. Verify that the cluster is healthy:

```
cluster show
```

Show example

```

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::*>

```

14. Repeat steps 6 to 13 to install the NX-OS software on switch cs1.

15. Enable auto-revert on the cluster LIFs.

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

16. Verify that the cluster LIFs have reverted to their home port:

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

Show example

```
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::*>
```

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

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

What's next?

[Install RCF.](#)

Install the Reference Configuration File

You install the Reference Configuration File (RCF) after setting up the Nexus 9336C-FX2 storage switch for the first time.

See the Knowledge Base article [How to clear configuration on a Cisco interconnect switch while retaining](#)

[remote connectivity](#) for further information when installing your RCF.

Before you begin

Verify the following installations and connections:

- A console connection to the switch. The console connection is optional if you have remote access to the switch.
- Switch cs1 and switch cs2 are powered up and the initial switch setup is complete (the Management IP address and SSH is set up).
- The desired NX-OS version has been installed.
- ONTAP node cluster ports are not connected.

Step 1: Install the RCF on the switches

1. Login to switch cs1 using SSH or by using a serial console.
2. Copy the RCF to the bootflash of switch cs1 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 9000 Series NX-OS Command Reference](#).

Show example

This example shows TFTP being used to copy an RCF to the bootflash on switch cs1:

```
cs1# copy tftp: bootflash: vrf management
Enter source filename: Nexus_9336C_RCF_v1.6-Storage.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)...
```

3. Apply the RCF previously downloaded to the bootflash.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 9000 Series NX-OS Command Reference](#).

Show example

This example shows the RCF `Nexus_9336C_RCF_v1.6-Storage.txt` being installed on switch cs1:

```
cs1# copy NNexus_9336C_RCF_v1.6-Storage.txt running-config echo-
commands
```


4. Examine the banner output from the `show banner motd` command. You must read and follow these instructions to ensure the correct configuration and operation of the switch.

Show example

```
cs1# show banner motd

*****
*****
* NetApp Reference Configuration File (RCF)
*
* Switch      : Nexus N9K-C9336C-FX2
* Filename    : Nexus_9336C_RCF_v1.6-Storage.txt
* Date       : 10-23-2020
* Version    : v1.6
*
* Port Usage : Storage configuration
* Ports 1-36: 100GbE Controller and Shelf Storage Ports
*****
*****
```

5. Verify that the RCF 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.

6. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
7. After you verify the RCF versions and switch settings are correct, copy the running-config file to the startup-config file:

```
copy running-config startup-config
```

Show example

```
cs1# copy running-config startup-config  
[#####] 100% Copy complete
```

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 9000 Series NX-OS Command Reference](#).

8. Reboot switch cs1.

```
cs1# reload
```

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

9. Repeat steps 1 through 7 on switch cs2.
10. Connect the cluster ports of all nodes in the ONTAP cluster to switches cs1 and cs2.

Step 2: Verify the switch connections

1. Verify that the switch ports connected to the cluster ports are **up**.

```
show interface brief
```

Show example

```
cs1# show interface brief | grep up
mgmt0  --                up      <mgmt ip address>
1000   1500
Eth1/11      1          eth  trunk  up      none
100G(D)  --
Eth1/12      1          eth  trunk  up      none
100G(D)  --
Eth1/13      1          eth  trunk  up      none
100G(D)  --
Eth1/14      1          eth  trunk  up      none
100G(D)  --
Eth1/15      1          eth  trunk  up      none
100G(D)  --
Eth1/16      1          eth  trunk  up      none
100G(D)  --
Eth1/17      1          eth  trunk  up      none
100G(D)  --
Eth1/18      1          eth  trunk  up      none
100G(D)  --
Eth1/23      1          eth  trunk  up      none
100G(D)  --
Eth1/24      1          eth  trunk  up      none
100G(D)  --
Eth1/25      1          eth  trunk  up      none
100G(D)  --
Eth1/26      1          eth  trunk  up      none
100G(D)  --
Eth1/27      1          eth  trunk  up      none
100G(D)  --
Eth1/28      1          eth  trunk  up      none
100G(D)  --
Eth1/29      1          eth  trunk  up      none
100G(D)  --
Eth1/30      1          eth  trunk  up      none
100G(D)  --
```

2. Verify that the cluster nodes are in their correct cluster VLANs using the following commands:

```
show vlan brief
```

```
show interface trunk
```

Show example

```
cs1# show vlan brief
```

VLAN	Name	Status	Ports
1	default	active	Po999
30	VLAN0030	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4 Eth1/5, Eth1/6, Eth1/7, Eth1/8 Eth1/9, Eth1/10, Eth1/11 Eth1/12, Eth1/13, Eth1/14 Eth1/15, Eth1/16, Eth1/17 Eth1/18, Eth1/19, Eth1/20 Eth1/21, Eth1/22, Eth1/23 Eth1/24, Eth1/25, Eth1/26 Eth1/27, Eth1/28, Eth1/29 Eth1/30, Eth1/31, Eth1/32 Eth1/33, Eth1/34, Eth1/35 Eth1/36

```
cs1# show interface trunk
```

Port	Native Vlan	Status	Port Channel
Eth1/1	1	trunking	--
Eth1/2	1	trunking	--
Eth1/3	1	trunking	--
Eth1/4	1	trunking	--
Eth1/5	1	trunking	--
Eth1/6	1	trunking	--
Eth1/7	1	trunking	--
Eth1/8	1	trunking	--

Eth1/9	1	trunking	--
Eth1/10	1	trunking	--
Eth1/11	1	trunking	--
Eth1/12	1	trunking	--
Eth1/13	1	trunking	--
Eth1/14	1	trunking	--
Eth1/15	1	trunking	--
Eth1/16	1	trunking	--
Eth1/17	1	trunking	--
Eth1/18	1	trunking	--
Eth1/19	1	trunking	--
Eth1/20	1	trunking	--
Eth1/21	1	trunking	--
Eth1/22	1	trunking	--
Eth1/23	1	trunking	--
Eth1/24	1	trunking	--
Eth1/25	1	trunking	--
Eth1/26	1	trunking	--
Eth1/27	1	trunking	--
Eth1/28	1	trunking	--
Eth1/29	1	trunking	--
Eth1/30	1	trunking	--
Eth1/31	1	trunking	--
Eth1/32	1	trunking	--
Eth1/33	1	trunking	--
Eth1/34	1	trunking	--
Eth1/35	1	trunking	--
Eth1/36	1	trunking	--

Port	Vlans Allowed on Trunk
------	------------------------

Eth1/1	30
Eth1/2	30
Eth1/3	30
Eth1/4	30
Eth1/5	30
Eth1/6	30
Eth1/7	30
Eth1/8	30
Eth1/9	30
Eth1/10	30
Eth1/11	30
Eth1/12	30

```
Eth1/13      30
Eth1/14      30
Eth1/15      30
Eth1/16      30
Eth1/17      30
Eth1/18      30
Eth1/19      30
Eth1/20      30
Eth1/21      30
Eth1/22      30
Eth1/23      30
Eth1/24      30
Eth1/25      30
Eth1/26      30
Eth1/27      30
Eth1/28      30
Eth1/29      30
Eth1/30      30
Eth1/31      30
Eth1/32      30
Eth1/33      30
Eth1/34      30
Eth1/35      30
Eth1/36      30
```

```
-----
-----
Port          Vlans Err-disabled on Trunk
-----
```

```
-----
Eth1/1        none
Eth1/2        none
Eth1/3        none
Eth1/4        none
Eth1/5        none
Eth1/6        none
Eth1/7        none
Eth1/8        none
Eth1/9        none
Eth1/10       none
Eth1/11       none
Eth1/12       none
Eth1/13       none
Eth1/14       none
Eth1/15       none
Eth1/16       none
```

```
Eth1/17      none
Eth1/18      none
Eth1/19      none
Eth1/20      none
Eth1/21      none
Eth1/22      none
Eth1/23      none
Eth1/24      none
Eth1/25      none
Eth1/26      none
Eth1/27      none
Eth1/28      none
Eth1/29      none
Eth1/30      none
Eth1/31      none
Eth1/32      none
Eth1/33      none
Eth1/34      none
Eth1/35      none
Eth1/36      none
```

```
-----
-----
Port          STP Forwarding
-----
```

```
-----
Eth1/1        none
Eth1/2        none
Eth1/3        none
Eth1/4        none
Eth1/5        none
Eth1/6        none
Eth1/7        none
Eth1/8        none
Eth1/9        none
Eth1/10       none
Eth1/11       30
Eth1/12       30
Eth1/13       30
Eth1/14       30
Eth1/15       30
Eth1/16       30
Eth1/17       30
Eth1/18       30
Eth1/19       none
Eth1/20       none
```

```
Eth1/21      none
Eth1/22      none
Eth1/23      30
Eth1/24      30
Eth1/25      30
Eth1/26      30
Eth1/27      30
Eth1/28      30
Eth1/29      30
Eth1/30      30
Eth1/31      none
Eth1/32      none
Eth1/33      none
Eth1/34      none
Eth1/35      none
Eth1/36      none
```

```
-----
-----
Port          Vlans in spanning tree forwarding state and not pruned
-----
```

```
-----
Eth1/1        Feature VTP is not enabled
none
Eth1/2        Feature VTP is not enabled
none
Eth1/3        Feature VTP is not enabled
none
Eth1/4        Feature VTP is not enabled
none
Eth1/5        Feature VTP is not enabled
none
Eth1/6        Feature VTP is not enabled
none
Eth1/7        Feature VTP is not enabled
none
Eth1/8        Feature VTP is not enabled
none
Eth1/9        Feature VTP is not enabled
none
Eth1/10       Feature VTP is not enabled
none
Eth1/11       Feature VTP is not enabled
30
Eth1/12       Feature VTP is not enabled
30
```



```
Eth1/13      Feature VTP is not enabled
30
Eth1/14      Feature VTP is not enabled
30
Eth1/15      Feature VTP is not enabled
30
Eth1/16      Feature VTP is not enabled
30
Eth1/17      Feature VTP is not enabled
30
Eth1/18      Feature VTP is not enabled
30
Eth1/19      Feature VTP is not enabled
none
Eth1/20      Feature VTP is not enabled
none
Eth1/21      Feature VTP is not enabled
none
Eth1/22      Feature VTP is not enabled
none
Eth1/23      Feature VTP is not enabled
30
Eth1/24      Feature VTP is not enabled
30
Eth1/25      Feature VTP is not enabled
30
Eth1/26      Feature VTP is not enabled
30
Eth1/27      Feature VTP is not enabled
30
Eth1/28      Feature VTP is not enabled
30
Eth1/29      Feature VTP is not enabled
30
Eth1/30      Feature VTP is not enabled
30
Eth1/31      Feature VTP is not enabled
none
Eth1/32      Feature VTP is not enabled
none
Eth1/33      Feature VTP is not enabled
none
Eth1/34      Feature VTP is not enabled
none
Eth1/35      Feature VTP is not enabled
none
```

```
Eth1/36      Feature VTP is not enabled
none
```



For specific port and VLAN usage details, refer to the banner and important notes section in your RCF.

Step 3: Set up your ONTAP cluster

NetApp recommends that you use System Manager to set up new clusters.

System Manager provides a simple and easy workflow for cluster setup and configuration including assigning a node management IP address, initializing the cluster, creating a local tier, configuring protocols and provisioning initial storage.

Go to [Configure ONTAP on a new cluster with System Manager](#) for setup instructions.

What's next?

[Configure switch health monitoring.](#)

Replace a Cisco Nexus 9336C-FX2 storage switch

You can replace a defective Nexus 9336C-FX2 switch in a cluster network. This is a nondisruptive procedure.

What you'll need

Before installing the NX-OS software and RCFs on a Cisco Nexus 9336C-FX2 storage switch, ensure that:

- Your system can support Cisco Nexus 9336C-FX2 storage switches.
- You have consulted the switch compatibility table on the Cisco Ethernet Switch page for the supported ONTAP, NX-OS, and RCF versions.
- You have referred to the appropriate software and upgrade guides available on the Cisco web site.

Cisco Nexus 3000 Series Switches:

- You have downloaded the applicable RCFs.
- The existing network configuration has the following characteristics:
 - The Cisco Ethernet Switches page has the latest RCF and NX-OS versions on your switches.
 - Management connectivity must exist on both switches.
- The replacement Cisco Nexus 9336C-FX2 switch has the following characteristics:
 - Management network connectivity is functional.
 - Console access to the replacement switch is in place.
 - The appropriate RCF and NX-OS operating system image is loaded onto the switch.
 - Initial configuration of the switch is complete.

About this task

This procedure replaces the second Nexus 9336C-FX2 storage switch S2 with the new 9336C-FX switch NS2.

The two nodes are node1 and node2.

Steps to complete:

- Confirm the switch to be replaced is S2.
- Disconnect the cables from switch S2.
- Reconnect the cables to switch NS2.
- Verify all device configurations on switch NS2.



There can be dependencies between command syntax in the RCF and NX-OS versions.

Steps

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

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

x is the duration of the maintenance window in hours.

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
```

Show example

```
storage::*> storage port show -port-type ENET
```

Node	Port	Type	Mode	Speed (Gb/s)	State	Status	VLAN ID
node1							
	e3a	ENET	storage	100	enabled	online	30
	e3b	ENET	storage	0	enabled	offline	30
	e7a	ENET	storage	0	enabled	offline	30
	e7b	ENET	storage	0	enabled	offline	30
node2							
	e3a	ENET	storage	100	enabled	online	30
	e3b	ENET	storage	0	enabled	offline	30
	e7a	ENET	storage	0	enabled	offline	30
	e7b	ENET	storage	0	enabled	offline	30

```
storage::*>
```

3. Verify that storage switch S1 is available:

```
network device-discovery show
```

Show example

```
storage::*> network device-discovery show
Node/          Local Discovered
Protocol      Port  Device (LLDP: ChassisID)  Interface  Platform
-----      -
node1/cdp
      e3a   S1                Ethernet1/1  NX9336C
      e4a   node2             e4a          AFF-A700
      e4e   node2             e4e          AFF-A700
node1/lldp
      e3a   S1                Ethernet1/1  -
      e4a   node2             e4a          -
      e4e   node2             e4e          -
node2/cdp
      e3a   S1                Ethernet1/2  NX9336C
      e4a   node1             e4a          AFF-A700
      e4e   node1             e4e          AFF-A700
node2/lldp
      e3a   S1                Ethernet1/2  -
      e4a   node1             e4a          -
      e4e   node1             e4e          -
storage::*>
```

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
```

Show example

```
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           e3a
node2              Eth1/2     121        S           e3a
SHFGD2008000011   Eth1/5     121        S           e0a
SHFGD2008000011   Eth1/6     120        S           e0a
SHFGD2008000022   Eth1/7     120        S           e0a
SHFGD2008000022   Eth1/8     120        S           e0a
```

5. Verify the shelf ports in the storage system:

```
storage shelf port show -fields remote-device,remote-port
```

Show example

```
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  -            -  
3.20    2  Ethernet1/6  S1  
3.20    3  -            -  
3.30    0  Ethernet1/7  S1  
3.20    1  -            -  
3.30    2  Ethernet1/8  S1  
3.20    3  -            -  
storage::*>
```

6. Remove all cables attached to storage switch S2.
7. Reconnect all cables to the replacement switch NS2.
8. Recheck the health status of the storage node ports:

```
storage port show -port-type ENET
```

Show example

```
storage::*> storage port show -port-type ENET  
  
Node           Port Type  Mode    Speed      State  Status  VLAN  
-----  -----  
node1  
           e3a  ENET  storage  100    enabled  online   30  
           e3b  ENET  storage   0    enabled  offline  30  
           e7a  ENET  storage   0    enabled  offline  30  
           e7b  ENET  storage   0    enabled  offline  30  
node2  
           e3a  ENET  storage  100    enabled  online   30  
           e3b  ENET  storage   0    enabled  offline  30  
           e7a  ENET  storage   0    enabled  offline  30  
           e7b  ENET  storage   0    enabled  offline  30  
storage::*>
```

9. Verify that both switches are available:

```
network device-discovery show
```

Show example

```
storage::*> network device-discovery show
Node/      Local Discovered
Protocol  Port  Device (LLDP: ChassisID)  Interface  Platform
-----  -
node1/cdp
          e3a  S1                        Ethernet1/1 NX9336C
          e4a  node2                     e4a         AFF-A700
          e4e  node2                     e4e         AFF-A700
          e7b  NS2                       Ethernet1/1 NX9336C
node1/lldp
          e3a  S1                        Ethernet1/1 -
          e4a  node2                     e4a         -
          e4e  node2                     e4e         -
          e7b  NS2                       Ethernet1/1 -
node2/cdp
          e3a  S1                        Ethernet1/2 NX9336C
          e4a  node1                     e4a         AFF-A700
          e4e  node1                     e4e         AFF-A700
          e7b  NS2                       Ethernet1/2 NX9336C
node2/lldp
          e3a  S1                        Ethernet1/2 -
          e4a  node1                     e4a         -
          e4e  node1                     e4e         -
          e7b  NS2                       Ethernet1/2 -
storage::*>
```

10. Verify the shelf ports in the storage system:

```
storage shelf port show -fields remote-device,remote-port
```

Show example

```
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  
storage::*>
```

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
```

What's next?

[Configure switch health monitoring.](#)

NVIDIA SN2100

Overview

Overview of configuration process for NVIDIA SN2100 storage switches

The NVIDIA SN2100 is an Ethernet switch that allows you to switch data between controllers and disk shelves.

Initial configuration overview

To configure a NVIDIA SN2100 switch on systems running ONTAP, follow these steps:

1. [Install the hardware for the NVIDIA SN2100 switch.](#)

Instructions are available in the *NVIDIA Switch Installation Guide*.

2. [Configure the switch.](#)

Instructions are available in the NVIDIA documentation.

3. [Review cabling and configuration considerations.](#)

Review requirements for optical connections, the QSA adapter, and the switchport speed.

4. [Cable NS224 shelves as switch-attached storage.](#)

Follow these procedures if you have a system in which the NS224 drive shelves need to be cabled as switch-attached storage (not direct-attached storage).

5. [Install Cumulus Linux in Cumulus mode](#) or [install Cumulus Linux in ONIE mode.](#)

You can install Cumulus Linux (CL) OS when the switch is running either Cumulus Linux or ONIE.

6. [Install the Reference Configuration File script.](#)

There are two RCF scripts available for Clustering and Storage applications.

7. [Install the CSHM file.](#)

You can install the applicable configuration file for Ethernet switch health monitoring of NVIDIA cluster switches.

Additional information

Before you begin installation or maintenance, be sure to review the following:

- [Configuration requirements](#)
- [Components and part numbers](#)
- [Required documentation](#)

Configuration requirements for NVIDIA SN2100 switches

For NVIDIA SN2100 switch installation and maintenance, be sure to review all requirements.

Installation requirements

If you want to build ONTAP clusters with more than two nodes, you need two supported cluster network switches. You can use additional management switches, which are optional.

You install the NVIDIA SN2100 switch (X190006/X190106) in the NVIDIA dual/single switch cabinet with the standard brackets that are included with the switch.

For cabling guidelines, see [Cabling and configuration considerations](#).

ONTAP and Linux support

The NVIDIA SN2100 switch is a 10/25/40/100 Gb Ethernet switch running Cumulus Linux. The switch supports the following:

- ONTAP 9.10.1P3. The SN2100 switch serves Cluster and Storage applications in ONTAP 9.10.1P3 over different switch-pairs. From ONTAP 9.10.1P3, you can use NVIDIA SN2100 switches to combine storage and cluster functionality into a shared switch configuration.
- Cumulus Linux (CL) OS version 4.4.3. For current compatibility information, see the [NVIDIA Ethernet Switches](#) information page.
- You can install Cumulus Linux when the switch is running Cumulus Linux or ONIE.

Components and part numbers for NVIDIA SN2100 switches

For NVIDIA SN2100 switch installation and maintenance, be sure to review the list of components and part numbers for the cabinet and rail kit.

Cabinet details

You install the NVIDIA SN2100 switch (X190006/X190106) in the NVIDIA dual/single switch cabinet with the standard brackets that are included with the switch.

Rail kit details

The following table lists the part number and description for the MSN2100 switches and rail kits:

Part number	Description
X190006-PE	Cluster Switch, NVIDIA SN2100, 16PT 100G, PTSX
X190006-PI	Cluster Switch, NVIDIA SN2100, 16PT 100G, PSIN
X190106-FE-PE	Switch, NVIDIA SN2100, 16PT 100G, PTSX, Front End
X190106-FE-PI	Switch, NVIDIA SN2100, 16PT 100G, PSIN, Front End
X-MTEF-KIT-D	Rail Kit, NVIDIA Dual switch side by side
X-MTEF-KIT-E	Rail Kit, NVIDIA Single switch short depth



See NVIDIA documentation for details on [installing your SN2100 switch and rail kit](#).

Documentation requirements for NVIDIA SN2100 switches

For NVIDIA SN2100 switch installation and maintenance, be sure to review all the recommended documentation.

The following table lists the documentation available for the NVIDIA SN2100 switches.

Title	Description
Setup and configure your NVIDIA SN2100 switches	Describes how to setup and configure your NVIDIA SN2100 switches, including installing Cumulus Linux and applicable RCFs.
Migrate from a Cisco cluster switch to a NVIDIA SN2100 cluster switch	Describes how to migrate from environments that use Cisco cluster switches to environments that use NVIDIA SN2100 cluster switches.
Migrate from a Cisco storage switch to a NVIDIA storage switch	Describes how to migrate from environments that use Cisco storage switches to environments that use NVIDIA SN2100 storage switches.

Title	Description
Migrate to a two-node switched cluster with NVIDIA SN2100 cluster switches	Describes how to migrate to a two-node switched environment using NVIDIA SN2100 cluster switches.
Replace a NVIDIA SN2100 cluster switch	Describes the procedure to replace a defective NVIDIA SN2100 switch in a cluster and download Cumulus Linux and reference configuration file.
Replace a NVIDIA SN2100 storage switch	Describes the procedure to replace a defective NVIDIA SN2100 storage switch and download Cumulus Linux and reference configuration file.

Install hardware

Install the hardware for the NVIDIA SN2100 switch

To install the SN2100 hardware, refer to NVIDIA's documentation.

Steps

1. Review the [configuration requirements](#).
2. Follow the instructions in [NVIDIA Switch Installation Guide](#).

What's next?

[Configure the switch](#).

Configure the NVIDIA SN2100 switch

To configure the SN2100 switch, refer to NVIDIA's documentation.

Steps

1. Review the [configuration requirements](#).
2. Follow the instructions in [NVIDIA System Bring-Up..](#)

What's next?

[Review cabling and configuration considerations](#).

Review cabling and configuration considerations

Before configuring your NVIDIA SN2100 switch, review the following considerations.

NVIDIA port details

Switch ports	Ports usage
swp1s0-3	4x10GbE breakout cluster port nodes
swp2s0-3	4x25GbE breakout cluster port nodes

swp3-14	40/100GbE cluster port nodes
swp15-16	100GbE Inter-Switch Link (ISL) ports

See the [Hardware Universe](#) for more information on switch ports.

Link-up delays with optical connections

If you are experiencing link-up delays of more than five seconds, Cumulus Linux 5.4 and later includes support for fast link-up. You can configure the links by using the `nv set` command as follows:

```
nv set interface <interface-id> link fast-linkup on  
nv config apply  
reload the switchd
```

Show example

```
cumulus@cumulus-cs13:mgmt:~$ nv set interface swp5 link fast-linkup on  
cumulus@cumulus-cs13:mgmt:~$ nv config apply  
switchd need to reload on this config change  
  
Are you sure? [y/N] y  
applied [rev_id: 22]  
  
Only switchd reload required
```

Support for copper connections

The following configuration changes are required to fix this issue.

Cumulus Linux 4.4.3

1. Identify the name for each interface using 40GbE/100GbE copper cables:

```
cumulus@cumulus:mgmt:~$ net show interface pluggables
```

Interface	Identifier	Vendor Name	Vendor PN	Vendor SN
swp3	0x11 (QSFP28)	Molex	112-00576	93A2229911111
B0				
swp4	0x11 (QSFP28)	Molex	112-00576	93A2229922222
B0				

2. Add the following two lines to the `/etc/cumulus/switchd.conf` file for every port (`swp<n>`) that is using 40GbE/100GbE copper cables:

- `interface.swp<n>.enable_media_depended_linkup_flow=TRUE`
- `interface.swp<n>.enable_short_tuning=TRUE`

For example:

```
cumulus@cumulus:mgmt:~$ sudo nano /etc/cumulus/switchd.conf
.
.
interface.swp3.enable_media_depended_linkup_flow=TRUE
interface.swp3.enable_short_tuning=TRUE
interface.swp4.enable_media_depended_linkup_flow=TRUE
interface.swp4.enable_short_tuning=TRUE
```

3. Restart the `switchd` service:

```
cumulus@cumulus:mgmt:~$ sudo systemctl restart switchd.service
```

4. Confirm that the ports are up:

```
cumulus@cumulus:mgmt:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp3	100G	9216	Trunk/L2		Master: bridge (UP)
UP	swp4	100G	9216	Trunk/L2		Master: bridge (UP)

Cumulus Linux 5.x

1. Identify the name for each interface using 40GbE/100GbE copper cables:

```
cumulus@cumulus:mgmt:~$ nv show interface pluggables
```

Interface	Identifier	Vendor Name	Vendor PN	Vendor SN
swp3 B0	0x11 (QSFP28)	Molex	112-00576	93A2229911111
swp4 B0	0x11 (QSFP28)	Molex	112-00576	93A2229922222

2. Configure the links using the `nv set` command as follows:

- `nv set interface <interface-id> link fast-linkup on`
- `nv config apply`
- Reload the `switchd` service

For example:

```
cumulus@cumulus:mgmt:~$ nv set interface swp5 link fast-linkup on
cumulus@cumulus:mgmt:~$ nv config apply
switchd need to reload on this config change

Are you sure? [y/N] y
applied [rev_id: 22]

Only switchd reload required
```

3. Confirm that the ports are up:

```
cumulus@cumulus:mgmt:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp3	100G	9216	Trunk/L2		Master: bridge (UP)
UP	swp4	100G	9216	Trunk/L2		Master: bridge (UP)

See the Knowledge Base article [SN2100 switch fails to connect using 40/100GbE copper cables](#) for further details.

On Cumulus Linux 4.4.2, copper connections are not supported on SN2100 switches with X1151A NIC, X1146A NIC, or onboard 100GbE ports.

For example:

- AFF A800 on ports e0a and e0b
- AFF A320 on ports e0g and e0h

QSA adapter

When a QSA adapter is used to connect to the 10GbE/25GbE cluster ports on a platform, the link might not come up.

To resolve this issue, do the following:

- For 10GbE, manually set the swp1s0-3 link speed to 10000 and set auto-negotiation to off.
- For 25GbE, manually set the swp2s0-3 link speed to 25000 and set auto-negotiation to off.



When using 10GbE/25GbE QSA adapters, insert them in non-breakout 40GbE/100GbE ports (swp3-swp14). Do not insert the QSA adapter in a port that is configured for breakout.

Set interface speed on breakout ports

Depending on the transceiver in the switch port, you might need to set the speed on the switch interface to a fixed speed. If using 10GbE and 25GbE breakout ports, verify that auto-negotiation is off and set the interface speed on the switch.

Cumulus Linux 4.4.3

For example:

```
cumulus@cumulus:mgmt:~$ net add int swpls3 link autoneg off && net com
--- /etc/network/interfaces      2019-11-17 00:17:13.470687027 +0000
+++ /run/nclu/ifupdown2/interfaces.tmp  2019-11-24 00:09:19.435226258
+0000
@@ -37,21 +37,21 @@
     alias 10G Intra-Cluster Node
     link-autoneg off
     link-speed 10000 <---- port speed set
     mstpctl-bpduguard yes
     mstpctl-portadminedge yes
     mtu 9216

auto swpls3
iface swpls3
    alias 10G Intra-Cluster Node
-   link-autoneg off
+   link-autoneg on
    link-speed 10000 <---- port speed set
    mstpctl-bpduguard yes
    mstpctl-portadminedge yes
    mtu 9216

auto swp2s0
iface swp2s0
    alias 25G Intra-Cluster Node
    link-autoneg off
    link-speed 25000 <---- port speed set
```

Check the interface and port status to verify that the settings are applied:

```
cumulus@cumulus:mgmt:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary

.						
.						
UP	swp1s0	10G	9216	Trunk/L2	cs07 (e4c)	Master:
	br_default(UP)					
UP	swp1s1	10G	9216	Trunk/L2	cs07 (e4d)	Master:
	br_default(UP)					
UP	swp1s2	10G	9216	Trunk/L2	cs08 (e4c)	Master:
	br_default(UP)					
UP	swp1s3	10G	9216	Trunk/L2	cs08 (e4d)	Master:
	br_default(UP)					
.						
.						
UP	swp3	40G	9216	Trunk/L2	cs03 (e4e)	Master:
	br_default(UP)					
UP	swp4	40G	9216	Trunk/L2	cs04 (e4e)	Master:
	br_default(UP)					
DN	swp5	N/A	9216	Trunk/L2		Master:
	br_default(UP)					
DN	swp6	N/A	9216	Trunk/L2		Master:
	br_default(UP)					
DN	swp7	N/A	9216	Trunk/L2		Master:
	br_default(UP)					
.						
.						
UP	swp15	100G	9216	BondMember	cs01 (swp15)	Master:
	cluster_isl(UP)					
UP	swp16	100G	9216	BondMember	cs01 (swp16)	Master:
	cluster_isl(UP)					
.						
.						

Cumulus Linux 5.x

For example:


```
cumulus@cumulus:mgmt:~$ nv set interface swp1s3 link auto-negotiate off
cumulus@cumulus:mgmt:~$ nv set interface swp1s3 link speed 10G
cumulus@cumulus:mgmt:~$ nv show interface swp1s3

link

  auto-negotiate      off          off
off
  duplex              full         full
full
  speed               10G         10G
10G
  fec                 auto         auto
auto
  mtu                  9216        9216
9216
[breakout]

  state                up           up
up
```

Check the interface and port status to verify that the settings are applied:

```
cumulus@cumulus:mgmt:~$ nv show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp1s0	10G	9216	Trunk/L2	cs07 (e4c)	Master: br_default(UP)
UP	swp1s1	10G	9216	Trunk/L2	cs07 (e4d)	Master: br_default(UP)
UP	swp1s2	10G	9216	Trunk/L2	cs08 (e4c)	Master: br_default(UP)
UP	swp1s3	10G	9216	Trunk/L2	cs08 (e4d)	Master: br_default(UP)
UP	swp3	40G	9216	Trunk/L2	cs03 (e4e)	Master: br_default(UP)
UP	swp4	40G	9216	Trunk/L2	cs04 (e4e)	Master: br_default(UP)
DN	swp5	N/A	9216	Trunk/L2		Master: br_default(UP)
DN	swp6	N/A	9216	Trunk/L2		Master: br_default(UP)
DN	swp7	N/A	9216	Trunk/L2		Master: br_default(UP)
UP	swp15	100G	9216	BondMember	cs01 (swp15)	Master: cluster_isl(UP)
UP	swp16	100G	9216	BondMember	cs01 (swp16)	Master: cluster_isl(UP)

What's next?

[Cable NS224 shelves as switch-attached storage.](#)

Cable NS224 shelves as switch-attached storage

If you have a system in which the NS224 drive shelves need to be cabled as switch-attached storage (not direct-attached storage), use the information provided here.

- Cable NS224 drive shelves through storage switches:

[Information for cabling switch-attached NS224 drive shelves](#)

- Install your storage switches:

[AFF and FAS Switch Documentation](#)

- Confirm supported hardware, such as storage switches and cables, for your platform model:

[NetApp Hardware Universe](#)

Configure software

Software install workflow for NVIDIA SN2100 storage switches

To install and configure the software for a NVIDIA SN2100 switch, follow these steps:

1. [Install Cumulus Linux in Cumulus mode](#) or [install Cumulus Linux in ONIE mode](#).

You can install Cumulus Linux (CL) OS when the switch is running either Cumulus Linux or ONIE.

2. [Install the Reference Configuration File script](#).

There are two RCF scripts available for Clustering and Storage applications.

3. [Install the CSHM file](#).

You can install the applicable configuration file for Ethernet switch health monitoring of NVIDIA cluster switches.

Install Cumulus Linux in Cumulus mode

Follow this procedure to install Cumulus Linux (CL) OS when the switch is running in Cumulus mode.



Cumulus Linux (CL) OS can be installed either when the switch is running Cumulus Linux or ONIE (see [Install in ONIE mode](#)).

What you'll need

- Intermediate-level Linux knowledge.
- Familiarity with basic text editing, UNIX file permissions, and process monitoring. A variety of text editors are pre-installed, including `vi` and `nano`.
- Access to a Linux or UNIX shell. If you are running Windows, use a Linux environment as your command line tool for interacting with Cumulus Linux.
- The baud rate requirement must be set to 115200 on the serial console switch for NVIDIA SN2100 switch console access, as follows:
 - 115200 baud
 - 8 data bits
 - 1 stop bit
 - parity: none

- flow control: none

About this task

Be aware of the following:



Each time Cumulus Linux is installed, the entire file system structure is erased and rebuilt.



The default password for the cumulus user account is **cumulus**. The first time you log into Cumulus Linux, you must change this default password. Be sure to update any automation scripts before installing a new image. Cumulus Linux provides command line options to change the default password automatically during the installation process.

Steps

1. Log in to the switch.

First time log in to the switch requires username/password of **cumulus/cumulus** with `sudo` privileges.

Show example

```
cumulus login: cumulus
Password: cumulus
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>
```

2. Check the Cumulus Linux version:

```
net show system
```

Show example

```
cumulus@cumulus:mgmt:~$ net show system
Hostname..... cumulus
Build..... Cumulus Linux 4.4.3
Uptime..... 0:08:20.860000
Model..... Mlnx X86
CPU..... x86_64 Intel Atom C2558 2.40GHz
Memory..... 8GB
Disk..... 14.7GB
ASIC..... Mellanox Spectrum MT52132
Ports..... 16 x 100G-QSFP28
Part Number..... MSN2100-CB2FC
Serial Number.... MT2105T05177
Platform Name.... x86_64-mlnx_x86-r0
Product Name..... MSN2100
ONIE Version..... 2019.11-5.2.0020-115200
Base MAC Address. 04:3F:72:43:92:80
Manufacturer..... Mellanox
```

3. Configure the hostname, IP address, subnet mask, and default gateway. The new hostname only becomes effective after restarting the console/SSH session.



A Cumulus Linux switch provides at least one dedicated Ethernet management port called eth0. This interface is specifically for out-of-band management use. By default, the management interface uses DHCPv4 for addressing.



Do not use an underscore (_), apostrophe ('), or non-ASCII characters in the hostname.

Show example

```
cumulus@cumulus:mgmt:~$ net add hostname sw1
cumulus@cumulus:mgmt:~$ net add interface eth0 ip address
10.233.204.71
cumulus@cumulus:mgmt:~$ net add interface eth0 ip gateway
10.233.204.1
cumulus@cumulus:mgmt:~$ net pending
cumulus@cumulus:mgmt:~$ net commit
```

This command modifies both the /etc/hostname and /etc/hosts files.

4. Confirm that the hostname, IP address, subnet mask, and default gateway have been updated.

Show example

```
cumulus@sw1:mgmt:~$ hostname sw1
cumulus@sw1:mgmt:~$ ifconfig eth0
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 10.233.204.71 netmask 255.255.254.0 broadcast 10.233.205.255
inet6 fe80::bace:f6ff:fe19:1df6 prefixlen 64 scopeid 0x20<link>
ether b8:ce:f6:19:1d:f6 txqueuelen 1000 (Ethernet)
RX packets 75364 bytes 23013528 (21.9 MiB)
RX errors 0 dropped 7 overruns 0 frame 0
TX packets 4053 bytes 827280 (807.8 KiB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0 device
memory 0xdfc00000-dfc1ffff

cumulus@sw1::mgmt:~$ ip route show vrf mgmt
default via 10.233.204.1 dev eth0
unreachable default metric 4278198272
10.233.204.0/23 dev eth0 proto kernel scope link src 10.233.204.71
127.0.0.0/8 dev mgmt proto kernel scope link src 127.0.0.1
```

5. Configure the time zone using NTP interactive mode.

- a. On a terminal, run the following command:

```
cumulus@sw1:~$ sudo dpkg-reconfigure tzdata
```

- b. Follow the on-screen menu options to select the geographic area and region.
- c. To set the time zone for all services and daemons, reboot the switch.
- d. Verify that the date and time on the switch are correct and update if necessary.

6. Install Cumulus Linux 4.4.3:

```
cumulus@sw1:mgmt:~$ sudo onie-install -a -i http://<web-
server>/<path>/cumulus-linux-4.4.3-mlx-amd64.bin
```

The installer starts the download. Type **y** when prompted.

7. Reboot the NVIDIA SN2100 switch:

```
cumulus@sw1:mgmt:~$ sudo reboot
```

8. The installation starts automatically, and the following GRUB screens appear. Do **not** make any selections:

- Cumulus-Linux GNU/Linux

- ONIE: Install OS
- CUMULUS-INSTALL
- Cumulus-Linux GNU/Linux

9. Repeat steps 1 to 4 to log in.

10. Verify that the Cumulus Linux version is 4.4.3:

```
net show version
```

Show example

```
cumulus@sw1:mgmt:~$ net show version  
NCLU_VERSION=1.0-cl4.4.3u0  
DISTRIB_ID="Cumulus Linux"  
DISTRIB_RELEASE=4.4.3  
DISTRIB_DESCRIPTION="Cumulus Linux 4.4.3"
```

11. Create a new user and add this user to the `sudo` group. This user only becomes effective after the console/SSH session is restarted.

```
sudo adduser --ingroup netedit admin
```

Show example

```
cumulus@sw1:mgmt:~$ sudo adduser --ingroup netedit admin
[sudo] password for cumulus:
Adding user `admin' ...
Adding new user `admin' (1001) with group `netedit' ...
Creating home directory `/home/admin' ...
Copying files from `/etc/skel' ...
New password:
Retype new password:
passwd: password updated successfully
Changing the user information for admin
Enter the new value, or press ENTER for the default
Full Name []:
Room Number []:
Work Phone []:
Home Phone []:
Other []:
Is the information correct? [Y/n] y

cumulus@sw1:mgmt:~$ sudo adduser admin sudo
[sudo] password for cumulus:
Adding user `admin' to group `sudo' ...
Adding user admin to group sudo
Done.
cumulus@sw1:mgmt:~$ exit
logout
Connection to 10.233.204.71 closed.

[admin@cycrh6svl01 ~]$ ssh admin@10.233.204.71
admin@10.233.204.71's password:
Linux sw1 4.19.0-cl-1-amd64 #1 SMP Cumulus 4.19.206-1+cl4.4.3u1
(2021-09-09) x86_64
Welcome to NVIDIA Cumulus (R) Linux (R)

For support and online technical documentation, visit
http://www.cumulusnetworks.com/support

The registered trademark Linux (R) is used pursuant to a sublicense
from LMI, the exclusive licensee of Linus Torvalds, owner of the
mark on a world-wide basis.
admin@sw1:mgmt:~$
```

What's next?

[Install RCF script.](#)

Install Cumulus Linux in ONIE mode

Follow this procedure to install Cumulus Linux (CL) OS when the switch is running in ONIE mode.



Cumulus Linux (CL) OS can be installed either when the switch is running Cumulus Linux or ONIE (see [Install in Cumulus mode](#)).

About this task

You can install the Cumulus Linux using Open Network Install Environment (ONIE) that allows for automatic discovery of a network installer image. This facilitates the system model of securing switches with an operating system choice, such as Cumulus Linux. The easiest way to install Cumulus Linux with ONIE is with local HTTP discovery.



If your host is IPv6-enabled, make sure it is running a web server. If your host is IPv4-enabled, make sure it is running DHCP in addition to a web server.

This procedure demonstrates how to upgrade Cumulus Linux after the admin has booted in ONIE.

Steps

1. Download the Cumulus Linux installation file to the root directory of the web server. Rename this file `onie-installer`.
2. Connect your host to the management Ethernet port of the switch using an Ethernet cable.
3. Power on the switch. The switch downloads the ONIE image installer and boots. After the installation completes, the Cumulus Linux login prompt appears in the terminal window.



Each time Cumulus Linux is installed, the entire file system structure is erased and rebuilt.

4. Reboot the SN2100 switch:

```
cumulus@cumulus:mgmt:~$ sudo reboot
```

5. Press the **Esc** key at the GNU GRUB screen to interrupt the normal boot process, select **ONIE** and press **Enter**.
6. On the next screen displayed, select **ONIE: Install OS**.
7. The ONIE installer discovery process runs searching for the automatic installation. Press **Enter** to temporarily stop the process.
8. When the discovery process has stopped:

```
ONIE:/ # onie-stop  
discover: installer mode detected.  
Stopping: discover...start-stop-daemon: warning: killing process 427:  
No such process done.
```

9. If the DHCP service is running on your network, verify that the IP address, subnet mask, and the default gateway are correctly assigned:

```
ifconfig eth0
```

Show example

```
ONIE:/ # ifconfig eth0
eth0  Link encap:Ethernet  HWaddr B8:CE:F6:19:1D:F6
      inet addr:10.233.204.71  Bcast:10.233.205.255
Mask:255.255.254.0
      inet6 addr: fe80::bace:f6ff:fe19:1df6/64 Scope:Link
UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
RX packets:21344 errors:0 dropped:2135 overruns:0 frame:0
TX packets:3500 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:1000
RX bytes:6119398 (5.8 MiB)  TX bytes:472975 (461.8 KiB)
Memory:dfc00000-dfc1ffff

ONIE:/ # route
Kernel IP routing table
Destination      Gateway          Genmask         Flags Metric Ref
Use Iface

default          10.233.204.1    0.0.0.0         UG    0     0
0 eth0
10.233.204.0    *                255.255.254.0   U     0     0
0 eth0
```

10. If the IP addressing scheme is manually defined, do the following:

```
ONIE:/ # ifconfig eth0 10.233.204.71 netmask 255.255.254.0
ONIE:/ # route add default gw 10.233.204.1
```

11. Repeat step 9 to verify that the static information is correctly entered.
12. Install Cumulus Linux:

```
ONIE:/ # route

Kernel IP routing table

ONIE:/ # onie-nos-install http://<web-server>/<path>/cumulus-linux-
4.4.3-mlx-amd64.bin

Stopping: discover... done.
Info: Attempting
http://10.60.132.97/x/eng/testbedN,svl/nic/files/cumulus-linux-4.4.3-
mlx-amd64.bin ...
Connecting to 10.60.132.97 (10.60.132.97:80)
installer          100% |*|    552M  0:00:00 ETA
...
...
```

13. Once the installation has completed, log in to the switch:

Show example

```
cumulus login: cumulus
Password: cumulus
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password: cumulus
New password: <new_password>
Retype new password: <new_password>
```

14. Verify the Cumulus Linux version:

```
net show version
```

Show example

```
cumulus@cumulus:mgmt:~$ net show version
NCLU_VERSION=1.0-cl4.4.3u4
DISTRIB_ID="Cumulus Linux"
DISTRIB_RELEASE=4.4.3
DISTRIB_DESCRIPTION="Cumulus Linux 4.4.3"
```

What's next?

[Install RCF script.](#)

Install or upgrade the RCF script

Follow this procedure to install or upgrade the RCF script.

What you'll need

Before installing or upgrading the RCF script, make sure that the following are available on the switch:

- Cumulus Linux 4.4.3 is installed.
- IP address, subnet mask, and default gateway defined via DHCP or manually configured.

Current RCF script versions

There are two RCF scripts available for Clustering and Storage applications. The procedure for each is the same.

- Clustering: **MSN2100-RCF-v1.x-Cluster**
- Storage: **MSN2100-RCF-v1.x-Storage**



The following example procedure shows how to download and apply the RCF script for Cluster switches.



Example command output uses switch management IP address 10.233.204.71, netmask 255.255.254.0 and default gateway 10.233.204.1.

Steps

1. Display the available interfaces on the SN2100 switch:

```
net show interface all
```

Show example

```
cumulus@cumulus:mgmt:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
-----	-----	---	-----	-----	-----	-----
.....						
.....						
ADMDN	swp1	N/A	9216	NotConfigured		
ADMDN	swp2	N/A	9216	NotConfigured		
ADMDN	swp3	N/A	9216	NotConfigured		
ADMDN	swp4	N/A	9216	NotConfigured		
ADMDN	swp5	N/A	9216	NotConfigured		
ADMDN	swp6	N/A	9216	NotConfigured		
ADMDN	swp7	N/A	9216	NotConfigure		
ADMDN	swp8	N/A	9216	NotConfigured		
ADMDN	swp9	N/A	9216	NotConfigured		
ADMDN	swp10	N/A	9216	NotConfigured		
ADMDN	swp11	N/A	9216	NotConfigured		
ADMDN	swp12	N/A	9216	NotConfigured		
ADMDN	swp13	N/A	9216	NotConfigured		
ADMDN	swp14	N/A	9216	NotConfigured		
ADMDN	swp15	N/A	9216	NotConfigured		
ADMDN	swp16	N/A	9216	NotConfigured		

2. Copy the RCF python script to the switch:

```
admin@sw1:mgmt:~$ pwd
/home/cumulus
cumulus@cumulus:mgmt:~$ cd /tmp
cumulus@cumulus:mgmt:/tmp$ scp <user>@<host>:<path>/MSN2100-RCF-v1.8-
Cluster
ssologin@10.233.204.71's password:
MSN2100-RCF-v1.8-Cluster          100% 8607    111.2KB/s
00:00
```

3. Apply the RCF python script **MSN2100-RCF-v1.8-Cluster**:

```
cumulus@cumulus:mgmt:/tmp$ sudo python3 MSN2100-RCF-v1.8-Cluster
[sudo] password for cumulus:
...
Step 1: Creating the banner file
Step 2: Registering banner message
Step 3: Updating the MOTD file
Step 4: Ensuring passwordless use of cl-support command by admin
Step 5: Disabling apt-get
Step 6: Creating the interfaces
Step 7: Adding the interface config
Step 8: Disabling cdp
Step 9: Adding the lldp config
Step 10: Adding the RoCE base config
Step 11: Modifying RoCE Config
Step 12: Configure SNMP
Step 13: Reboot the switch
```

The RCF script completes the steps listed above.



For any RCF python script issues that cannot be corrected, contact [NetApp Support](#) for assistance.

4. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
5. Verify the configuration after the reboot:

```
net show interface all
```

Show example

```
cumulus@cumulus:mgmt:~$ net show interface all
```

State	Name	Spd	MTU	Mode	LLDP	Summary
...						
...						
DN	swp1s0	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp1s1	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp1s2	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp1s3	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp2s0	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp2s1	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp2s2	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp2s3	N/A	9216	Trunk/L2		Master: bridge (UP)
UP	swp3	100G	9216	Trunk/L2		Master: bridge (UP)
UP	swp4	100G	9216	Trunk/L2		Master: bridge (UP)
DN	swp5	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp6	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp7	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp8	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp9	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp10	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp11	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp12	N/A	9216	Trunk/L2		Master: bridge (UP)
DN	swp13	N/A	9216	Trunk/L2		Master: bridge (UP)

```

DN      swp14      N/A    9216    Trunk/L2      Master:
bridge(UP)
UP      swp15      N/A    9216    BondMember    Master:
bond_15_16(UP)
UP      swp16      N/A    9216    BondMember    Master:
bond_15_16(UP)
...
...

```

```
cumulus@cumulus:mgmt:~$ net show roce config
```

```
RoCE mode..... lossless
```

```
Congestion Control:
```

```
Enabled SPs.... 0 2 5
```

```
Mode..... ECN
```

```
Min Threshold.. 150 KB
```

```
Max Threshold.. 1500 KB
```

```
PFC:
```

```
Status..... enabled
```

```
Enabled SPs.... 2 5
```

```
Interfaces..... swp10-16,swp1s0-3,swp2s0-3,swp3-9
```

```

DSCP                                802.1p  switch-priority
-----
0 1 2 3 4 5 6 7                    0      0
8 9 10 11 12 13 14 15              1      1
16 17 18 19 20 21 22 23            2      2
24 25 26 27 28 29 30 31            3      3
32 33 34 35 36 37 38 39            4      4
40 41 42 43 44 45 46 47            5      5
48 49 50 51 52 53 54 55            6      6
56 57 58 59 60 61 62 63            7      7

```

```

switch-priority  TC  ETS
-----
0 1 3 4 6 7      0  DWRR 28%
2                  2  DWRR 28%
5                  5  DWRR 43%

```

6. Verify information for the transceiver in the interface:

```
net show interface pluggables
```


Show example

```
cumulus@cumulus:mgmt:~$ net show interface pluggables
Interface Identifier      Vendor Name  Vendor PN      Vendor SN
Vendor Rev
-----
swp3      0x11 (QSFP28)  Amphenol     112-00574
APF20379253516 B0
swp4      0x11 (QSFP28)  AVAGO        332-00440      AF1815GU05Z
A0
swp15     0x11 (QSFP28)  Amphenol     112-00573
APF21109348001 B0
swp16     0x11 (QSFP28)  Amphenol     112-00573
APF21109347895 B0
```

7. Verify that the nodes each have a connection to each switch:

```
net show lldp
```

Show example

```
cumulus@cumulus:mgmt:~$ net show lldp
LocalPort  Speed  Mode           RemoteHost      RemotePort
-----
swp3       100G   Trunk/L2      sw1              e3a
swp4       100G   Trunk/L2      sw2              e3b
swp15      100G   BondMember    sw13             swp15
swp16      100G   BondMember    sw14             swp16
```

8. 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
```

Show example

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

Node: node1

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
e3a         Cluster    Cluster          up   9000  auto/10000
healthy     false
e3b         Cluster    Cluster          up   9000  auto/10000
healthy     false

Node: node2

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
e3a         Cluster    Cluster          up   9000  auto/10000
healthy     false
e3b         Cluster    Cluster          up   9000  auto/10000
healthy     false
```

- b. Verify the switch health from the cluster (this might not show switch sw2, since LIFs are not homed on e0d).

Show example

```
cluster1::*> network device-discovery show -protocol lldp
Node/          Local  Discovered
Protocol       Port   Device (LLDP: ChassisID)  Interface Platform
-----
node1/lldp
              e3a    sw1 (b8:ce:f6:19:1a:7e)   swp3      -
              e3b    sw2 (b8:ce:f6:19:1b:96)   swp3      -
node2/lldp
              e3a    sw1 (b8:ce:f6:19:1a:7e)   swp4      -
              e3b    sw2 (b8:ce:f6:19:1b:96)   swp4      -

cluster1::*> system switch ethernet show -is-monitoring-enabled
-operational true
Switch          Type          Address
Model
-----
sw1              cluster-network  10.233.205.90
MSN2100-CB2RC
  Serial Number: MNXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cumulus Linux version 4.4.3 running on
Mellanox
                    Technologies Ltd. MSN2100
  Version Source: LLDP

sw2              cluster-network  10.233.205.91
MSN2100-CB2RC
  Serial Number: MNCXXXXXXGS
  Is Monitored: true
  Reason: None
  Software Version: Cumulus Linux version 4.4.3 running on
Mellanox
                    Technologies Ltd. MSN2100
  Version Source: LLDP
```

What's next?

[Install the CSHM file.](#)

Install the Ethernet Switch Health Monitor configuration file

Follow this procedure to install the applicable configuration file for Ethernet switch health monitoring of NVIDIA cluster switches.

Supported models are:

- MSN2100-CB2FC
- MSN2100-CB2RC
- X190006-PE
- X190006-PI



This installation procedure applies to ONTAP 9.10.1 and later.

Before you begin

- Verify that you need to download the configuration file by running `system switch ethernet show` and checking if **OTHER** is shown for your model.

If your model is still showing **OTHER** after applying the configuration file, contact NetApp support.

- Make sure that the ONTAP cluster is up and running.
- Enable SSH to use all of the features available in CSHM.
- Clear the `/mroot/etc/cshm_nod/nod_sign/` directory on all nodes:

- a. Enter the nodeshell:

```
system node run -node <name>
```

- b. Change to advanced privilege:

```
priv set advanced
```

- c. List the configuration files in the `/etc/cshm_nod/nod_sign` directory. If the directory exists and contains configuration files, it lists the file names.

```
ls /etc/cshm_nod/nod_sign
```

- d. Delete all configuration files corresponding to your connected switch models.

If you are unsure, remove all configuration files for the supported models listed above, then download and install the latest configuration files for those same models.

```
rm /etc/cshm_nod/nod_sign/<filename>
```

- e. Confirm that the deleted configuration files are no longer in the directory:

```
ls /etc/cshm_nod/nod_sign
```

Steps

1. Download the Ethernet switch health monitor configuration zip file based on the corresponding ONTAP release version. This file is available from the [NVIDIA Ethernet switches](#) page.

- a. On the NVIDIA SN2100 Software download page, select **Nvidia CSHM File**.
- b. On the Caution/Must read page, select the check box to agree.
- c. On the End User License Agreement page, select the check box to agree and click **Accept & Continue**.
- d. On the Nvidia CSHM File - Download page, select the applicable configuration file. The following files are available:

ONTAP 9.15.1 and later

- MSN2100-CB2FC-v1.4.zip
- MSN2100-CB2RC-v1.4.zip
- X190006-PE-v1.4.zip
- X190006-PI-v1.4.zip

ONTAP 9.11.1 through 9.14.1

- MSN2100-CB2FC_PRIOR_R9.15.1-v1.4.zip
- MSN2100-CB2RC_PRIOR_R9.15.1-v1.4.zip
- X190006-PE_PRIOR_9.15.1-v1.4.zip
- X190006-PI_PRIOR_9.15.1-v1.4.zip

2. Upload the applicable zip file to your internal web server.
3. Access the advanced mode setting from one of the ONTAP systems in the cluster.

```
set -privilege advanced
```

4. Run the switch health monitor configure command.

```
cluster1::> system cluster-switch configure-health-monitor -node *  
-package-url http://192.168.2.20/usr/download/[filename.zip]
```

5. Verify that the command output ends with the following text for your ONTAP version:

ONTAP 9.15.1 and later

Ethernet switch health monitoring installed the configuration file.

ONTAP 9.11.1 through 9.14.1

SHM installed the configuration file.

ONTAP 9.10.1

CSHM downloaded package processed successfully.

If an error occurs, contact NetApp support.

6. Wait up to twice the Ethernet switch health monitor polling interval, found by running `system switch`

ethernet polling-interval show, before completing the next step.

7. Run the command `system switch ethernet show` on the ONTAP system and make sure that the cluster switches are discovered with the monitored field set to **True** and the serial number field not showing **Unknown**.

```
cluster1::> system switch ethernet show
```

What's next?

[Configure switch health monitoring.](#)

Migrate switches

Migrate from a Cisco storage switch to a NVIDIA SN2100 storage switch

You can migrate older Cisco switches for an ONTAP cluster to NVIDIA SN2100 storage switches. This is a non-disruptive procedure.

Review requirements

The following storage switches are supported:

- Cisco Nexus 9336C-FX2
- Cisco Nexus 3232C
- See the [Hardware Universe](#) for full details of supported ports and their configurations.

What you'll need

Ensure that:

- The existing cluster is properly set up and functioning.
- All storage ports are in the up state to ensure nondisruptive operations.
- The NVIDIA SN2100 storage switches are configured and operating under the proper version of Cumulus Linux installed with the reference configuration file (RCF) applied.
- The existing storage network configuration has the following:
 - A redundant and fully functional NetApp cluster using both older Cisco switches.
 - Management connectivity and console access to both the older Cisco switches and the new switches.
 - All cluster LIFs in the up state with the cluster LIFs are on their home ports.
 - ISL ports enabled and cabled between the older Cisco switches and between the new switches.
- See the [Hardware Universe](#) for full details of supported ports and their configurations.
- Some of the ports are configured on NVIDIA SN2100 switches to run at 100 GbE.
- You have planned, migrated, and documented 100 GbE connectivity from nodes to NVIDIA SN2100 storage switches.

Migrate the switches

About the examples

In this procedure, Cisco Nexus 9336C-FX2 storage switches are used for example commands and outputs.

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

- The existing Cisco Nexus 9336C-FX2 storage switches are *S1* and *S2*.
- The new NVIDIA SN2100 storage switches are *sw1* and *sw2*.
- The nodes are *node1* and *node2*.
- The cluster LIFs are *node1_clus1* and *node1_clus2* on node 1, and *node2_clus1* and *node2_clus2* on node 2 respectively.
- The `cluster1::*>` prompt indicates the name of the cluster.
- The network ports used in this procedure are *e5a* and *e5b*.
- Breakout ports take the format: *swp1s0-3*. For example four breakout ports on *swp1* are *swp1s0*, *swp1s1*, *swp1s2*, and *swp1s3*.
- Switch *S2* is replaced by switch *sw2* first and then switch *S1* is replaced by switch *sw1*.
 - Cabling between the nodes and *S2* are then disconnected from *S2* and reconnected to *sw2*.
 - Cabling between the nodes and *S1* are then disconnected from *S1* and reconnected to *sw1*.

Step 1: Prepare for migration

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

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

where *x* is the duration of the maintenance window in hours.

2. Change the privilege level to advanced, entering *y* when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (**>*) appears.

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

Each port should display enabled for *Status*.

Step 2: Configure cables and ports

1. Display the network port attributes:

```
storage port show
```

Show example

```
cluster1::*> storage port show
```

Node	Port	Type	Mode	Speed (Gb/s)	State	Status	VLAN ID
node1							
	e0c	ENET	storage	100	enabled	online	30
	e0d	ENET	storage	0	enabled	offline	30
	e5a	ENET	storage	0	enabled	offline	30
	e5b	ENET	storage	100	enabled	online	30
node2							
	e0c	ENET	storage	100	enabled	online	30
	e0d	ENET	storage	0	enabled	offline	30
	e5a	ENET	storage	0	enabled	offline	30
	e5b	ENET	storage	100	enabled	online	30

```
cluster1::*>
```

2. Verify that the storage ports on each node are connected to existing storage switches in the following way (from the nodes' perspective) using the command:

```
network device-discovery show -protocol lldp
```

Show example

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

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	
node1 /lldp				
	e0c	S1 (7c:ad:4f:98:6d:f0)	Eth1/1	-
	e5b	S2 (7c:ad:4f:98:8e:3c)	Eth1/1	-
node2 /lldp				
	e0c	S1 (7c:ad:4f:98:6d:f0)	Eth1/2	-
	e5b	S2 (7c:ad:4f:98:8e:3c)	Eth1/2	-

3. On switch S1 and S2, make sure that the storage ports and switches are connected in the following way (from the switches' perspective) using the command:

```
show lldp neighbors
```


Show example

```
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 Port ID	Local Intf	Holdtime	Capability
node1 e0c	Eth1/1	121	S
node2 e0c	Eth1/2	121	S
SHFGD1947000186 e0a	Eth1/10	120	S
SHFGD1947000186 e0a	Eth1/11	120	S
SHFGB2017000269 e0a	Eth1/12	120	S
SHFGB2017000269 e0a	Eth1/13	120	S

```
S2# 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 Port ID	Local Intf	Holdtime	Capability
node1 e5b	Eth1/1	121	S
node2 e5b	Eth1/2	121	S
SHFGD1947000186 e0b	Eth1/10	120	S
SHFGD1947000186 e0b	Eth1/11	120	S
SHFGB2017000269 e0b	Eth1/12	120	S
SHFGB2017000269 e0b	Eth1/13	120	S

- On switch sw2, shut down the ports connected to the storage ports and nodes of the disk shelves.

Show example

```
cumulus@sw2:~$ net add interface swp1-16 link down
cumulus@sw2:~$ net pending
cumulus@sw2:~$ net commit
```

- Move the node storage ports of the controller and disk shelves from the old switch S2 to the new switch sw2, using appropriate cabling supported by NVIDIA SN2100.
- On switch sw2, bring up the ports connected to the storage ports of the nodes and the disk shelves.

Show example

```
cumulus@sw2:~$ net del interface swp1-16 link down
cumulus@sw2:~$ net pending
cumulus@sw2:~$ net commit
```

- Verify that the storage ports on each node are now connected to the switches in the following way, from the nodes' perspective:

```
network device-discovery show -protocol lldp
```

Show example

```
cluster1::~*> network device-discovery show -protocol lldp
```

Node/ Protocol	Local Port	Discovered Device (LLDP: ChassisID)	Interface	Platform
node1	/lldp			
	e0c	S1 (7c:ad:4f:98:6d:f0)	Eth1/1	-
	e5b	sw2 (b8:ce:f6:19:1a:7e)	swp1	-
node2	/lldp			
	e0c	S1 (7c:ad:4f:98:6d:f0)	Eth1/2	-
	e5b	sw2 (b8:ce:f6:19:1a:7e)	swp2	-

- Verify the network port attributes:

```
storage port show
```

Show example

```
cluster1::*> storage port show
```

Node	Port	Type	Mode	Speed (Gb/s)	State	Status	VLAN ID
node1							
	e0c	ENET	storage	100	enabled	online	30
	e0d	ENET	storage	0	enabled	offline	30
	e5a	ENET	storage	0	enabled	offline	30
	e5b	ENET	storage	100	enabled	online	30
node2							
	e0c	ENET	storage	100	enabled	online	30
	e0d	ENET	storage	0	enabled	offline	30
	e5a	ENET	storage	0	enabled	offline	30
	e5b	ENET	storage	100	enabled	online	30

```
cluster1::*>
```

9. On switch sw2, verify that all node storage ports are up:

```
net show interface
```

Show example

```
cumulus@sw2:~$ net show interface

State  Name      Spd   MTU   Mode      LLDP
Summary
-----
...
...
UP      swp1      100G  9216  Trunk/L2  node1 (e5b)
Master: bridge(UP)
UP      swp2      100G  9216  Trunk/L2  node2 (e5b)
Master: bridge(UP)
UP      swp3      100G  9216  Trunk/L2  SHFFG1826000112 (e0b)
Master: bridge(UP)
UP      swp4      100G  9216  Trunk/L2  SHFFG1826000112 (e0b)
Master: bridge(UP)
UP      swp5      100G  9216  Trunk/L2  SHFFG1826000102 (e0b)
Master: bridge(UP)
UP      swp6      100G  9216  Trunk/L2  SHFFG1826000102 (e0b)
Master: bridge(UP)
...
...
```

10. On switch sw1, shut down the ports connected to the storage ports of the nodes and the disk shelves.

Show example

```
cumulus@sw1:~$ net add interface swp1-16 link down
cumulus@sw1:~$ net pending
cumulus@sw1:~$ net commit
```

11. Move the node storage ports of the controller and the disk shelves from the old switch S1 to the new switch sw1, using appropriate cabling supported by NVIDIA SN2100.
12. On switch sw1, bring up the ports connected to the storage ports of the nodes and the disk shelves.

Show example

```
cumulus@sw1:~$ net del interface swp1-16 link down
cumulus@sw1:~$ net pending
cumulus@sw1:~$ net commit
```

13. Verify that the storage ports on each node are now connected to the switches in the following way, from the nodes' perspective:

```
network device-discovery show -protocol lldp
```

Show example

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

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	

node1	/lldp			
	e0c	sw1 (b8:ce:f6:19:1b:96)	swp1	-
	e5b	sw2 (b8:ce:f6:19:1a:7e)	swp1	-
node2	/lldp			
	e0c	sw1 (b8:ce:f6:19:1b:96)	swp2	-
	e5b	sw2 (b8:ce:f6:19:1a:7e)	swp2	-

Step 3: Verify the configuration

1. Verify the final configuration:

```
storage port show
```

Each port should display enabled for State and enabled for Status.

Show example

```
cluster1::*> storage port show
```

Node	Port	Type	Mode	Speed (Gb/s)	State	Status	VLAN ID
node1							
	e0c	ENET	storage	100	enabled	online	30
	e0d	ENET	storage	0	enabled	offline	30
	e5a	ENET	storage	0	enabled	offline	30
	e5b	ENET	storage	100	enabled	online	30
node2							
	e0c	ENET	storage	100	enabled	online	30
	e0d	ENET	storage	0	enabled	offline	30
	e5a	ENET	storage	0	enabled	offline	30
	e5b	ENET	storage	100	enabled	online	30

```
cluster1::*>
```

2. On switch sw2, verify that all node storage ports are up:

```
net show interface
```

Show example

```
cumulus@sw2:~$ net show interface
```

```
State Name      Spd   MTU   Mode      LLDP
Summary
-----
...
...
UP      swp1      100G  9216   Trunk/L2  node1 (e5b)
Master: bridge(UP)
UP      swp2      100G  9216   Trunk/L2  node2 (e5b)
Master: bridge(UP)
UP      swp3      100G  9216   Trunk/L2  SHFFG1826000112 (e0b)
Master: bridge(UP)
UP      swp4      100G  9216   Trunk/L2  SHFFG1826000112 (e0b)
Master: bridge(UP)
UP      swp5      100G  9216   Trunk/L2  SHFFG1826000102 (e0b)
Master: bridge(UP)
UP      swp6      100G  9216   Trunk/L2  SHFFG1826000102 (e0b)
Master: bridge(UP)
...
...
```

3. Verify that both nodes each have one connection to each switch:

```
net show lldp
```

Show example

The following example shows the appropriate results for both switches:

```
cumulus@sw1:~$ net show lldp
LocalPort  Speed  Mode      RemoteHost      RemotePort
-----
...
swp1       100G   Trunk/L2  node1           e0c
swp2       100G   Trunk/L2  node2           e0c
swp3       100G   Trunk/L2  SHFFG1826000112 e0a
swp4       100G   Trunk/L2  SHFFG1826000112 e0a
swp5       100G   Trunk/L2  SHFFG1826000102 e0a
swp6       100G   Trunk/L2  SHFFG1826000102 e0a

cumulus@sw2:~$ net show lldp
LocalPort  Speed  Mode      RemoteHost      RemotePort
-----
...
swp1       100G   Trunk/L2  node1           e5b
swp2       100G   Trunk/L2  node2           e5b
swp3       100G   Trunk/L2  SHFFG1826000112 e0b
swp4       100G   Trunk/L2  SHFFG1826000112 e0b
swp5       100G   Trunk/L2  SHFFG1826000102 e0b
swp6       100G   Trunk/L2  SHFFG1826000102 e0b
```

1. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

What's next?

[Configure switch health monitoring.](#)

Replace a NVIDIA SN2100 storage switch

You can replace a defective NVIDIA SN2100 storage switch. This is a nondisruptive procedure.

What you'll need

Before installing the Cumulus software and RCFs on a NVIDIA SN2100 storage switch, ensure that:

- Your system can support NVIDIA SN2100 storage switches.

- You have downloaded the applicable RCFs.

The [Hardware Universe](#) provides full details of supported ports and their configurations.

The existing network configuration must have the following characteristics:

- Make sure that all troubleshooting steps have been completed to confirm that your switch needs replacing.
- Management connectivity must exist on both switches.



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

The replacement NVIDIA SN2100 switch must have the following characteristics:

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

Procedure summary

This procedure replaces the second NVIDIA SN2100 storage switch sw2 with the new NVIDIA SN2100 switch nsw2. The two nodes are node1 and node2.

Steps to complete:

- Confirm the switch to be replaced is sw2.
- Disconnect the cables from switch sw2.
- Reconnect the cables to switch nsw2.
- Verify all device configurations on switch nsw2.

Steps

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

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

x is the duration of the maintenance window in hours.

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

```
set -privilege advanced
```

3. 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
```

Show example

```
cluster1::*> storage port show -port-type ENET
```

Node	Port	Type	Mode	Speed (Gb/s)	State	Status	VLAN ID

node1	e3a	ENET	storage	100	enabled	online	30
	e3b	ENET	storage	0	enabled	offline	30
	e7a	ENET	storage	0	enabled	offline	30
	e7b	ENET	storage	100	enabled	online	30
node2	e3a	ENET	storage	100	enabled	online	30
	e3b	ENET	storage	0	enabled	offline	30
	e7a	ENET	storage	0	enabled	offline	30
	e7b	ENET	storage	100	enabled	online	30

```
cluster1::*>
```

4. Verify that storage switch sw1 is available:

```
network device-discovery show -protocol lldp
```

Show example

```
cluster1::*> network device-discovery show -protocol lldp
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
node1/lldp
    e0M      sw1 (00:ea:bd:68:6a:e8)  Eth1/46      -
    e0b      sw2 (6c:b2:ae:5f:a5:b2)  Ethernet1/16 -
    e0c      SHFFG1827000286 (d0:39:ea:1c:16:92)
                                     e0a          -
    e0e      sw3 (6c:b2:ae:5f:a5:ba)  Ethernet1/18 -
    e0f      SHFFG1827000286 (00:a0:98:fd:e4:a9)
                                     e0b          -
    e0g      sw4 (28:ac:9e:d5:4a:9c)  Ethernet1/11 -
    e0h      sw5 (6c:b2:ae:5f:a5:ca)  Ethernet1/22 -
    e1a      sw6 (00:f6:63:10:be:7c)  Ethernet1/33 -
    e1b      sw7 (00:f6:63:10:be:7d)  Ethernet1/34 -
    e2a      sw8 (b8:ce:f6:91:3d:88)  Ethernet1/35 -
Press <space> to page down, <return> for next line, or 'q' to
quit...
10 entries were displayed.
```

5. Run the `net show interface` command on the working switch to confirm that you can see both nodes and all shelves:

```
net show interface
```

Show example

```
cumulus@sw1:~$ net show interface

State  Name      Spd   MTU   Mode      LLDP
Summary
-----
...
...
UP      swp1      100G  9216  Trunk/L2  node1 (e3a)
Master: bridge(UP)
UP      swp2      100G  9216  Trunk/L2  node2 (e3a)
Master: bridge(UP)
UP      swp3      100G  9216  Trunk/L2  SHFFG1826000112 (e0b)
Master: bridge(UP)
UP      swp4      100G  9216  Trunk/L2  SHFFG1826000112 (e0b)
Master: bridge(UP)
UP      swp5      100G  9216  Trunk/L2  SHFFG1826000102 (e0b)
Master: bridge(UP)
UP      swp6      100G  9216  Trunk/L2  SHFFG1826000102 (e0b)
Master: bridge(UP)
...
...
```

6. Verify the shelf ports in the storage system:

```
storage shelf port show -fields remote-device, remote-port
```

Show example

```
cluster1::*> storage shelf port show -fields remote-device, remote-  
port  
shelf    id  remote-port  remote-device  
-----  --  -  
3.20     0  swp3         sw1  
3.20     1  -           -  
3.20     2  swp4         sw1  
3.20     3  -           -  
3.30     0  swp5         sw1  
3.20     1  -           -  
3.30     2  swp6         sw1  
3.20     3  -           -  
cluster1::*>
```

7. Remove all cables attached to storage switch sw2.
8. Reconnect all cables to the replacement switch nsw2.
9. Recheck the health status of the storage node ports:

```
storage port show -port-type ENET
```

Show example

```
cluster1::*> storage port show -port-type ENET  
  
Node      Port Type  Mode   Speed      State   Status  VLAN  
-----  -  
node1  
          e3a  ENET  storage 100    enabled  online  30  
          e3b  ENET  storage  0     enabled  offline 30  
          e7a  ENET  storage  0     enabled  offline 30  
          e7b  ENET  storage 100    enabled  online  30  
node2  
          e3a  ENET  storage 100    enabled  online  30  
          e3b  ENET  storage  0     enabled  offline 30  
          e7a  ENET  storage  0     enabled  offline 30  
          e7b  ENET  storage 100    enabled  online  30  
cluster1::*>
```

10. Verify that both switches are available:

```
net device-discovery show -protocol lldp
```

Show example

```
cluster1::*> network device-discovery show -protocol lldp
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
node1/lldp
e0M           sw1 (00:ea:bd:68:6a:e8)    Eth1/46          -
e0b           sw2 (6c:b2:ae:5f:a5:b2)    Ethernet1/16     -
e0c           SHFFG1827000286 (d0:39:ea:1c:16:92)
                                           e0a              -
e0e           sw3 (6c:b2:ae:5f:a5:ba)    Ethernet1/18     -
e0f           SHFFG1827000286 (00:a0:98:fd:e4:a9)
                                           e0b              -
e0g           sw4 (28:ac:9e:d5:4a:9c)    Ethernet1/11     -
e0h           sw5 (6c:b2:ae:5f:a5:ca)    Ethernet1/22     -
e1a           sw6 (00:f6:63:10:be:7c)    Ethernet1/33     -
e1b           sw7 (00:f6:63:10:be:7d)    Ethernet1/34     -
e2a           sw8 (b8:ce:f6:91:3d:88)    Ethernet1/35     -
Press <space> to page down, <return> for next line, or 'q' to
quit...
10 entries were displayed.
```

11. Verify the shelf ports in the storage system:

```
storage shelf port show -fields remote-device, remote-port
```

Show example

```
cluster1::*> storage shelf port show -fields remote-device, remote-  
port  
shelf    id    remote-port    remote-device  
-----  --    -  
3.20     0    swp3           sw1  
3.20     1    swp3           nsw2  
3.20     2    swp4           sw1  
3.20     3    swp4           nsw2  
3.30     0    swp5           sw1  
3.20     1    swp5           nsw2  
3.30     2    swp6           sw1  
3.20     3    swp6           nsw2  
cluster1::*>
```

12. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

What's next?

[Configure switch health monitoring.](#)

Shared switches

Cisco Nexus 9336C-FX2

Overview

Overview of installation and configuration for Cisco Nexus 9336C-FX2 shared switches

The Cisco Nexus 9336C-FX2 shared switch is part of the Cisco Nexus 9000 platform and can be installed in a NetApp system cabinet. Shared switches allow you to combine cluster and storage functionality into a shared switch configuration, by supporting the use of shared cluster and storage Reference Configuration Files.

Initial configuration overview

To initially configure a Cisco Nexus 9336C-FX2 switch on systems running ONTAP, follow these steps:

1. [Complete cabling worksheet.](#)
2. [Install the switch.](#)
3. [Configure the switch.](#)
4. [Install the switch in a NetApp cabinet.](#)
5. [Prepare to install the NX-OS and RCF.](#)
6. [Install the NX-OS software.](#)
7. [Install the RCF config file.](#)
8. [Enable SSH.](#)

Additional information

Before you begin installation or maintenance, be sure to review the following:

- [Configuration requirements](#)
- [Components and part numbers](#)
- [Required documentation](#)

Configuration requirements for Cisco Nexus 9336C-FX2 shared switches

For Cisco Nexus 9336C-FX2 switch installation and maintenance, be sure to review configuration and network requirements.

ONTAP support

From ONTAP 9.9.1, you can use Cisco Nexus 9336C-FX2 switches to combine storage and cluster functionality into a shared switch configuration.

If you want to build ONTAP clusters with more than two nodes, you need two supported network switches.

Configuration requirements

For configuration, you need the appropriate number and type of cables and cable connectors for your switches.

Depending on the type of switch you are initially configuring, you need to connect to the switch console port with the included console cable; you also need to provide specific network information.

Network requirements

You need the following network information for all switch configurations.

- IP subnet for management network traffic
- Host names and IP addresses for each of the storage system controllers and all applicable switches
- Most storage system controllers are managed through the e0M interface by connecting to the Ethernet service port (wrench icon). On AFF A800 and AFF A700s systems, the e0M interface uses a dedicated Ethernet port.
- Refer to the [Hardware Universe](#) for the latest information.

For more information about the initial configuration of your switch, see the following guide: [Cisco Nexus 9336C-FX2 Installation and Upgrade Guide](#).

Components and part numbers for Cisco Nexus 9336C-FX2 shared switches

For Cisco Nexus 9336C-FX2 switch installation and maintenance, be sure to review the list of components and part numbers.

The following table lists the part number and description for the 9336C-FX2 switch, fans, and power supplies:

Part number	Description
X190200-CS-PE	N9K-9336C-FX2, CS, PTSX, 36PT10/25/40/100GQSFP28
X190200-CS-PI	N9K-9336C-FX2, CS, PSIN, 36PT10/25/40/100GQSFP28
X190002	Accessory Kit X190001/X190003
X-NXA-PAC-1100W-PE2	N9K-9336C AC 1100W PSU - Port side exhaust airflow
X-NXA-PAC-1100W-PI2	N9K-9336C AC 1100W PSU - Port side Intake airflow
X-NXA-FAN-65CFM-PE	N9K-9336C 65CFM, Port side exhaust airflow
X-NXA-FAN-65CFM-PI	N9K-9336C 65CFM, Port side intake airflow

Documentation requirements for Cisco Nexus 9336C-FX2 shared switches

For Cisco Nexus 9336C-FX2 switch installation and maintenance, be sure to review specific switch and controller documentation to set up your Cisco 9336-FX2 switches and ONTAP cluster.

To set up the Cisco Nexus 9336C-FX2 shared switches, see the [Cisco Nexus 9000 Series Switches Support](#) page.

Document title	Description
Nexus 9000 Series Hardware Installation Guide	Provides detailed information about site requirements, switch hardware details, and installation options.
Cisco Nexus 9000 Series Switch Software Configuration Guides (choose the guide for the NX-OS release installed on your switches)	Provides initial switch configuration information that you need before you can configure the switch for ONTAP operation.
Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide (choose the guide for the NX-OS release installed on your switches)	Provides information on how to downgrade the switch to ONTAP supported switch software, if necessary.
Cisco Nexus 9000 Series NX-OS Command Reference Master Index	Provides links to the various command references provided by Cisco.
Cisco Nexus 9000 MIBs Reference	Describes the Management Information Base (MIB) files for the Nexus 9000 switches.
Nexus 9000 Series NX-OS System Message Reference	Describes the system messages for Cisco Nexus 9000 series switches, those that are informational, and others that might help diagnose problems with links, internal hardware, or the system software.
Cisco Nexus 9000 Series NX-OS Release Notes (choose the notes for the NX-OS release installed on your switches)	Describes the features, bugs, and limitations for the Cisco Nexus 9000 Series.
Regulatory Compliance and Safety Information for Cisco Nexus 9000 Series	Provides international agency compliance, safety, and statutory information for the Nexus 9000 series switches.

Install hardware


Complete the Cisco Nexus 9336C-FX2 cabling worksheet

Use the following cabling images to complete the cabling between the controllers and the switches.

Cable NS224 storage as switch-attached

If you want to cable NS224 storage as switch-attached, follow the switch-attached diagram:

Switch Attached

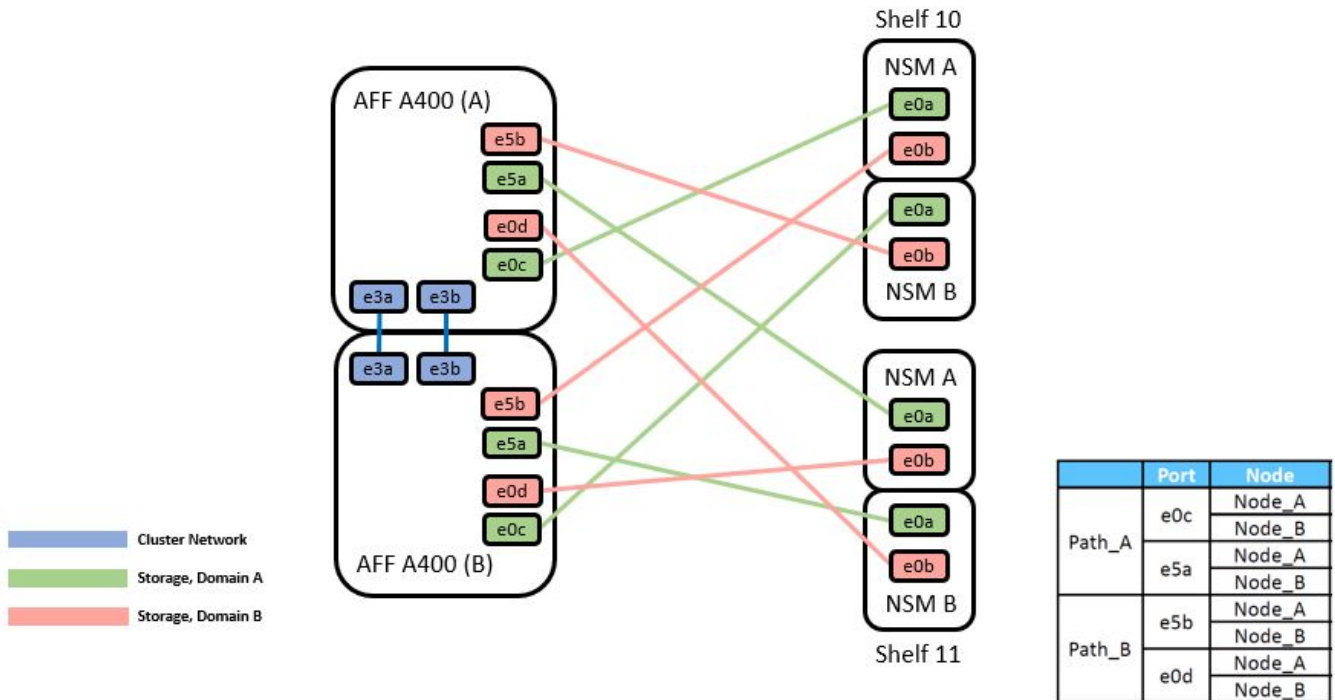


See the [Hardware Universe](#) for more information on switch ports.

Cable NS224 storage as direct-attached

If you want to cable NS224 storage as direct-attached instead of using the shared switch storage ports, follow the direct-attached diagram:

Direct Attached



See the [Hardware Universe](#) for more information on switch ports.

Cisco Nexus 9336C-FX2 cabling worksheet

If you want to document the supported platforms, you must complete the blank cabling worksheet by using completed sample cabling worksheet as a guide.

The sample port definition on each pair of switches is as follows:

Switch A			Switch B		
Switch Port	Port Role	Port Usage	Switch Port	Port Role	Port Usage
1	Cluster	40/100GbE	1	Cluster	40/100GbE
2	Cluster	40/100GbE	2	Cluster	40/100GbE
3	Cluster	40/100GbE	3	Cluster	40/100GbE
4	Cluster	40/100GbE	4	Cluster	40/100GbE
5	Cluster	40/100GbE	5	Cluster	40/100GbE
6	Cluster	40/100GbE	6	Cluster	40/100GbE
7	Cluster	40/100GbE	7	Cluster	40/100GbE
8	Cluster	40/100GbE	8	Cluster	40/100GbE
9	Cluster	40GbE w/4x10GbE b/o	9	Cluster	40GbE w/4x10GbE b/o
10	Cluster	100GbE w/4x25GbE b/o	10	Cluster	100GbE w/4x25GbE b/o
11	Storage	100GbE	11	Storage	100GbE
12	Storage	100GbE	12	Storage	100GbE
13	Storage	100GbE	13	Storage	100GbE
14	Storage	100GbE	14	Storage	100GbE
15	Storage	100GbE	15	Storage	100GbE
16	Storage	100GbE	16	Storage	100GbE
17	Storage	100GbE	17	Storage	100GbE
18	Storage	100GbE	18	Storage	100GbE
19	Storage	100GbE	19	Storage	100GbE
20	Storage	100GbE	20	Storage	100GbE
21	Storage	100GbE	21	Storage	100GbE
22	Storage	100GbE	22	Storage	100GbE
23	Storage	100GbE	23	Storage	100GbE
24	Storage	100GbE	24	Storage	100GbE
25	Storage	100GbE	25	Storage	100GbE
26	Storage	100GbE	26	Storage	100GbE
27	Storage	100GbE	27	Storage	100GbE
28	Storage	100GbE	28	Storage	100GbE
29	Storage	100GbE	29	Storage	100GbE
30	Storage	100GbE	30	Storage	100GbE
31	Storage	100GbE	31	Storage	100GbE
32	Storage	100GbE	32	Storage	100GbE
33	Storage	100GbE	33	Storage	100GbE
34	Storage	100GbE	34	Storage	100GbE
35	ISL	100GbE	35	ISL	100GbE
36	ISL	100GbE	36	ISL	100GbE

Where:

- 100G ISL to switch A port 35
- 100G ISL to switch A port 36
- 100G ISL to switch B port 35
- 100G ISL to switch B port 36

Blank cabling worksheet

You can use the blank cabling worksheet to document the platforms that are supported as nodes in a cluster. The Supported Cluster Connections table of the Hardware Universe defines the cluster ports used by the platform.

Switch Port	Switch A Port Role	Port Usage	Switch Port	Switch B Port Role	Port Usage
1			1		
2			2		
3			3		
4			4		
5			5		
6			6		
7			7		
8			8		
9			9		
10			10		
11			11		
12			12		
13			13		
14			14		
15			15		
16			16		
17			17		
18			18		
19			19		
20			20		
21			21		
22			22		
23			23		
24			24		
25			25		
26			26		
27			27		
28			28		
29			29		
30			30		
31			31		
32			32		
33			33		
34			34		
35			35		
36			36		

Where:

- 100G ISL to switch A port 35
- 100G ISL to switch A port 36
- 100G ISL to switch B port 35
- 100G ISL to switch B port 36

Install Cisco Nexus 9336C-FX2 shared switches

Follow these instructions to configure Cisco Nexus 9336C-FX2 shared switches.

What you'll need

- Required shared switch documentation, controller documentation and ONTAP documentation. See [Documentation requirements for Cisco Nexus 9336C-FX2 shared switches](#) and [NetApp ONTAP documentation](#).
- Applicable licenses, network and configuration information, and cables.
- Completed cabling worksheets. See [Complete the Cisco Nexus 9336C-FX2 cabling worksheet](#). For more information on cabling, refer to the [Hardware Universe](#).

Steps

1. Rack the switches, controllers and NS224 NVMe storage shelves.

See the [Racking instructions](#) to learn how to rack the switch in a NetApp cabinet.

2. Power on the switches, controllers and NS224 NVMe storage shelves.

What's next?

Go to [Configure Cisco Nexus 9336C-FX2 shared switch](#).

Configure Cisco Nexus 9336C-FX2 shared switches

Follow these instructions to configure Cisco Nexus 9336C-FX2 shared switches.

What you'll need

- Required shared switch documentation, controller documentation and ONTAP documentation. See [Documentation requirements for Cisco Nexus 9336C-FX2 shared switches](#) and [NetApp ONTAP documentation](#).
- Applicable licenses, network and configuration information, and cables.
- Completed cabling worksheets. See [Complete the Cisco Nexus 9336C-FX2 cabling worksheet](#). For more information on cabling, refer to the [Hardware Universe](#).

Steps

1. Perform an initial configuration of the switches.

For configuration, you need the appropriate number and type of cables and cable connectors for your switches.

Depending on the type of switch you are initially configuring, you need to connect to the switch console port with the included console cable; you also need to provide specific network information.

2. Boot the switch.

Provide the applicable responses to the following initial setup questions when you first boot the switch.

Your site's security policy defines the responses and services to enable.

- a. Abort Auto Provisioning and continue with normal setup? (yes/no)

Respond with **yes**. The default is no.

b. Do you want to enforce secure password standard? (yes/no)

Respond with **yes**. The default is yes.

c. Enter the password for admin.

The default password is admin; you must create a new, strong password.

A weak password can be rejected.

d. Would you like to enter the basic configuration dialog? (yes/no)

Respond with **yes** at the initial configuration of the switch.

e. Create another login account? (yes/no)

Your answer depends on your site's policies on alternate administrators. The default is no.

f. Configure read-only SNMP community string? (yes/no)

Respond with **no**. The default is no.

g. Configure read-write SNMP community string? (yes/no)

Respond with **no**. The default is no.

h. Enter the switch name.

The switch name is limited to 63 alphanumeric characters.

i. Continue with out-of-band (mgmt0) management configuration? (yes/no)

Respond with **yes** (the default) at that prompt. At the mgmt0 IPv4 address: prompt, enter your IP address: ip_address

j. Configure the default-gateway? (yes/no)

Respond with **yes**. At the IPv4 address of the default-gateway: prompt, enter your default_gateway.

k. Configure advanced IP options? (yes/no)

Respond with **no**. The default is no.

l. Enable the telnet service? (yes/no)

Respond with **no**. The default is no.

m. Enable SSH service? (yes/no)

Respond with **yes**. The default is yes.



SSH is recommended when using Ethernet Switch Health Monitor (CSHM) for its log collection features. SSHv2 is also recommended for enhanced security.

n. Enter the type of SSH key you want to generate (dsa/rsa/rsa1). The default is rsa.

o. Enter the number of key bits (1024- 2048).

p. Configure the NTP server? (yes/no)

Respond with **no**. The default is no.

q. Configure default interface layer (L3/L2):

Respond with **L2**. The default is L2.

r. Configure default switch port interface state (shut/noshut):

Respond with **noshut**. The default is noshut.

s. Configure CoPP system profile (strict/moderate/lenient/dense):

Respond with **strict**. The default is strict.

t. Would you like to edit the configuration? (yes/no)

You should see the new configuration at this point. Review and make any necessary changes to the configuration you just entered. Respond with no at the prompt if you are satisfied with the configuration. Respond with **yes** if you want to edit your configuration settings.

u. Use this configuration and save it? (yes/no)

Respond with **yes** to save the configuration. This automatically updates the kickstart and system images.

3. Verify the configuration choices you made in the display that appears at the end of the setup, and make sure that you save the configuration.



If you do not save the configuration at this stage, none of the changes will be in effect the next time you reboot the switch.

4. Check the version on the cluster network switches, and if necessary, download the NetApp-supported version of the software to the switches from the [Cisco software download](#) page.

What's next?

Depending on your configuration, you can [install switch in NetApp cabinet](#). Otherwise, go to [Prepare to install NX-OS and RCF](#).

Install a Cisco Nexus 9336C-FX2 switch in a NetApp cabinet

Depending on your configuration, you might need to install the Cisco Nexus 9336C-FX2 switch and pass-through panel in a NetApp cabinet. Standard brackets are included with the switch.

What you'll need

- 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 cabinet posts.
- You must use the Cisco standard rail kit to install the switch in a NetApp cabinet.



The jumper cords are not included with the pass-through kit and should be included with your switches. If they were not shipped with the switches, you can order them from NetApp (part number X1558A-R6).

Required documentation

Review the initial preparation requirements, kit contents, and safety precautions in the [Cisco Nexus 9000 Series Hardware Installation Guide](#).

Steps


1. Install the pass-through blanking panel in the NetApp cabinet.

The pass-through panel kit is available from NetApp (part number X8784-R6).

The NetApp pass-through panel kit contains the following hardware:


- One pass-through blanking panel
- Four 10-32 x .75 screws
- Four 10-32 clip nuts
 - a. Determine the vertical location of the switches and blanking panel in the cabinet.

In this procedure, the blanking panel will be installed in U40.
 - b. Install two clip nuts on each side in the appropriate square holes for front cabinet rails.
 - c. Center the panel vertically to prevent intrusion into adjacent rack space, and then tighten the screws.
 - d. Insert the female connectors of both 48-inch jumper cords from the rear of the panel and through the brush assembly.




(1) Female connector of the jumper cord.

2. Install the rack-mount brackets on the Nexus 9336C-FX2 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.
3. Install the clip nuts in the square hole locations for all four IEA posts.



The two 9336C-FX2 switches will always be mounted in the top 2U of the cabinet RU41 and 42.

4. Install the slider rails in the cabinet.
 - 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 cabinet posts.

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


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

5. Install the switch in the cabinet.



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) As the chassis is pushed toward the rear posts, align the two rear rack-mount guides with the slider rails.

(2) Gently slide the switch until the front rack-mount brackets are flush with the front posts.

b. Attach the switch to the cabinet.



(1) With one person holding the front of the chassis level, the other person should fully tighten the four rear screws to the cabinet 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.



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.

6. When the switches are installed, connect the jumper cords to the switch power inlets.
7. Connect the male plugs of both jumper cords to the closest available PDU outlets.



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

8. Connect the management port on each 9336C-FX2 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.

Configure software

Software install workflow for Cisco Nexus 9336C-FX2 shared switches

To install and configure software for a Cisco Nexus 9336C-FX2 shared switch, follow these steps:

1. [Prepare to install NX-OS and RCF.](#)
2. [Install the NX-OS software.](#)
3. [Install the RCF.](#)

Prepare to install NX-OS software and RCF

Before you install the NX-OS software and the Reference Configuration File (RCF), follow this procedure.

About the examples

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.

About this task

The procedure requires the use of both ONTAP commands and Cisco Nexus 9000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

Steps

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=x h`

where x is the duration of the maintenance window in hours.



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

2. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (`*>`) appears.

3. Display how many cluster interconnect interfaces are configured in each node for each cluster interconnect switch:

```
network device-discovery show -protocol cdp
```

Show example

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

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	
cluster1-02/cdp	e0a	cs1	Eth1/2	N9K-
C9336C	e0b	cs2	Eth1/2	N9K-
C9336C				
cluster1-01/cdp	e0a	cs1	Eth1/1	N9K-
C9336C	e0b	cs2	Eth1/1	N9K-
C9336C				

```
4 entries were displayed.
```

4. Check the administrative or operational status of each cluster interface.
 - a. Display the network port attributes:

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

Show example

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

Node: cluster1-02

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

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
```

Show example

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

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

4 entries were displayed.

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

Packet	Source	Destination
Node	Date	LIF
Loss		
-----	-----	-----
node1		
3/5/2022 19:21:18 -06:00	cluster1-01_clus2	cluster1-02-
clus1 none		
3/5/2022 19:21:20 -06:00	cluster1-01_clus2	cluster1-
02_clus2 none		
node2		
3/5/2022 19:21:18 -06:00	cluster1-02_clus2	cluster1-
01_clus1 none		
3/5/2022 19:21:20 -06:00	cluster1-02_clus2	cluster1-
01_clus2 none		

All ONTAP releases

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

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

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

6. Verify that the auto-revert command is enabled on all cluster LIFs:

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

Show example

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

Vserver	Logical Interface	Auto-revert
Cluster	cluster1-01_clus1	true
	cluster1-01_clus2	true
	cluster1-02_clus1	true
	cluster1-02_clus2	true

4 entries were displayed.

What's next?

[Install the NX-OS software.](#)

Install the NX-OS software

Follow this procedure to install the NX-OS software on the Nexus 9336C-FX2 shared switch.

Before you begin, complete the procedure in [Prepare to install NX-OS and RCF](#).

Review requirements

What you'll need

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).

Suggested documentation

- [Cisco Ethernet switch page](#)

Consult the switch compatibility table for the supported ONTAP and NX-OS versions.

- [Software Upgrade and downgrade guides](#)

Refer to the appropriate software and upgrade guides available on the Cisco website for complete documentation on the Cisco switch upgrade and downgrade procedures.

- [Cisco Nexus 9000 and 3000 Upgrade and ISSU Matrix](#)

Provides information on Disruptive Upgrade/Downgrade for Cisco NX-OS software on Nexus 9000 Series Switches based on your current and target releases.

On the page, select **Disruptive Upgrade** and select your current release and target release from the

dropdown list.

About the examples

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.

Install the software

The procedure requires the use of both ONTAP commands and Cisco Nexus 9000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

Steps

1. Connect the cluster switch to the management network.
2. Use the ping command to verify connectivity to the server hosting the NX-OS software and the RCF.

Show example

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.
```

3. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
              e0a    cs1                      Ethernet1/7      N9K-
C9336C-FX2
              e0d    cs2                      Ethernet1/7      N9K-
C9336C-FX2
cluster1-02/cdp
              e0a    cs1                      Ethernet1/8      N9K-
C9336C-FX2
              e0d    cs2                      Ethernet1/8      N9K-
C9336C-FX2
cluster1-03/cdp
              e0a    cs1                      Ethernet1/1/1    N9K-
C9336C-FX2
              e0b    cs2                      Ethernet1/1/1    N9K-
C9336C-FX2
cluster1-04/cdp
              e0a    cs1                      Ethernet1/1/2    N9K-
C9336C-FX2
              e0b    cs2                      Ethernet1/1/2    N9K-
C9336C-FX2
cluster1::*>
```

4. 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
```

Show example

```
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
```

```
Health Health Speed (Mbps)
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
```

Show example

```
cluster1::*> network interface show -role cluster
          Logical          Status      Network
Current   Current Is
Vserver   Interface          Admin/Oper Address/Mask      Node
Port      Home
-----
-----
Cluster
cluster1-01  cluster1-01_clus1  up/up      169.254.3.4/23
              e0a      true
cluster1-01  cluster1-01_clus2  up/up      169.254.3.5/23
              e0d      true
cluster1-02  cluster1-02_clus1  up/up      169.254.3.8/23
              e0a      true
cluster1-02  cluster1-02_clus2  up/up      169.254.3.9/23
              e0d      true
cluster1-03  cluster1-03_clus1  up/up      169.254.1.3/23
              e0a      true
cluster1-03  cluster1-03_clus2  up/up      169.254.1.1/23
              e0b      true
cluster1-04  cluster1-04_clus1  up/up      169.254.1.6/23
              e0a      true
cluster1-04  cluster1-04_clus2  up/up      169.254.1.7/23
              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
```


Show example

```
cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                                     Type                                     Address
Model
-----
cs1                                         cluster-network                         10.233.205.90
N9K-C9336C-FX2
  Serial Number: FOCXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                                     9.3(5)
  Version Source: CDP

cs2                                         cluster-network                         10.233.205.91
N9K-C9336C-FX2
  Serial Number: FOCXXXXXXGS
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                                     9.3(5)
  Version Source: CDP
cluster1::*>
```

5. Disable auto-revert on the cluster LIFs. The cluster LIFs fail over to the partner cluster switch and remain there as you perform the upgrade procedure on the targeted switch:

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

6. Copy the NX-OS software and EPLD images to the Nexus 9336C-FX2 switch.

Show example

```
cs2# copy sftp: bootflash: vrf management
Enter source filename: /code/nxos.9.3.5.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.9.3.5.bin /bootflash/nxos.9.3.5.bin
/code/nxos.9.3.5.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.9.3.5.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.9.3.5.img /bootflash/n9000-
epld.9.3.5.img
/code/n9000-epld.9.3.5.img 100% 161MB 9.5MB/s 00:16
sftp> exit
Copy complete, now saving to disk (please wait)...
Copy complete.
```

7. Verify the running version of the NX-OS software:

```
show version
```

Show example

```
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 08.38
  NXOS: version 9.3(4)
  BIOS compile time: 05/29/2020
  NXOS image file is: bootflash:///nxos.9.3.4.bin
  NXOS compile time: 4/28/2020 21:00:00 [04/29/2020 02:28:31]

Hardware
  cisco Nexus9000 C9336C-FX2 Chassis
  Intel(R) Xeon(R) CPU E5-2403 v2 @ 1.80GHz with 8154432 kB of
memory.
  Processor Board ID FOC20291J6K

  Device name: cs2
  bootflash: 53298520 kB
  Kernel uptime is 0 day(s), 0 hour(s), 3 minute(s), 42 second(s)
```

```
Last reset at 157524 usecs after Mon Nov  2 18:32:06 2020
Reason: Reset Requested by CLI command reload
System version: 9.3(4)
Service:

plugin
  Core Plugin, Ethernet Plugin

Active Package(s) :

cs2#
```

8. Install the NX-OS image.

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

Show example

```
cs2# install all nxos bootflash:nxos.9.3.5.bin
```

```
Installer will perform compatibility check first. Please wait.  
Installer is forced disruptive
```

```
Verifying image bootflash:/nxos.9.3.5.bin for boot variable "nxos".  
[] 100% -- SUCCESS
```

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

```
Preparing "nxos" version info using image bootflash:/nxos.9.3.5.bin.  
[] 100% -- SUCCESS
```

```
Preparing "bios" version info using image bootflash:/nxos.9.3.5.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(4)	9.3(5)
yes			
1	bios	v08.37(01/28/2020):v08.23(09/23/2015)	
v08.38(05/29/2020)		yes	

```
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.
```

9. Verify the new version of NX-OS software after the switch has rebooted:

```
show version
```

Show example

```
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 05.33
NXOS: version 9.3(5)
BIOS compile time: 09/08/2018
NXOS image file is: bootflash:///nxos.9.3.5.bin
NXOS compile time: 11/4/2018 21:00:00 [11/05/2018 06:11:06]
```

Hardware

```
cisco Nexus9000 C9336C-FX2 Chassis
Intel(R) Xeon(R) CPU E5-2403 v2 @ 1.80GHz with 8154432 kB of
memory.
Processor Board ID FOC20291J6K

Device name: cs2
bootflash: 53298520 kB
Kernel uptime is 0 day(s), 0 hour(s), 3 minute(s), 42 second(s)
```

```
Last reset at 277524 usecs after Mon Nov  2 22:45:12 2020
```

```
Reason: Reset due to upgrade
```

```
System version: 9.3(4)
```

```
Service:
```

```
plugin
```

```
Core Plugin, Ethernet Plugin
```

```
Active Package(s):
```

10. Upgrade the EPLD image and reboot the switch.

Show example



```
cs2# show version module 1 epld
```

EPLD	Device	Version
MI	FPGA	0x7
IO	FPGA	0x17
MI	FPGA2	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2

```
cs2# install epld bootflash:n9000-epld.9.3.5.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	0x07	0x07	No
1	SUP	IO FPGA	0x17	0x19	Yes
1	SUP	MI FPGA2	0x02	0x02	No

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

EPLDs upgraded.

Module 1 EPLD upgrade is successful.

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

Show example

```
cs2# show version module 1 epld
```

EPLD	Device	Version
MI	FPGA	0x7
IO	FPGA	0x19
MI	FPGA2	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2

12. Verify the health of cluster ports on the cluster.

a. Verify that cluster ports are up and healthy across all nodes in the cluster:

```
network port show -role cluster
```

Show example

```
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
```

```
Health Health Speed (Mbps)
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.

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
          e0a    cs1                        Ethernet1/7
N9K-C9336C-FX2
          e0d    cs2                        Ethernet1/7
N9K-C9336C-FX2
cluster01-2/cdp
          e0a    cs1                        Ethernet1/8
N9K-C9336C-FX2
          e0d    cs2                        Ethernet1/8
N9K-C9336C-FX2
cluster01-3/cdp
          e0a    cs1                        Ethernet1/1/1
N9K-C9336C-FX2
          e0b    cs2                        Ethernet1/1/1
N9K-C9336C-FX2
cluster1-04/cdp
          e0a    cs1                        Ethernet1/1/2
N9K-C9336C-FX2
          e0b    cs2                        Ethernet1/1/2
N9K-C9336C-FX2

cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                Type                Address
Model
-----
-----
cs1                    cluster-network    10.233.205.90
N9K-C9336C-FX2
  Serial Number: FOCXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                    9.3(5)
  Version Source: CDP

cs2                    cluster-network    10.233.205.91
```

```

N9K-C9336C-FX2
  Serial Number: FOCXXXXXXGS
    Is Monitored: true
      Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                9.3(5)
  Version Source: CDP

2 entries were displayed.

```

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.

```

13. Verify that the cluster is healthy:

```
cluster show
```

Show example

```

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::*>

```

14. Repeat steps 6 to 13 to install the NX-OS software on switch cs1.

15. Enable auto-revert on the cluster LIFs.

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

16. Verify that the cluster LIFs have reverted to their home port:

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

Show example

```
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::*>
```

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

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

What's next?

[Install the RCF.](#)

Install the Reference Configuration File (RCF)

You can install the reference configuration file (RCF) after setting up the Nexus 9336C-FX2 switch for the first time.

Before you begin, complete the procedure in [Prepare to install NX-OS and RCF.](#)

Before you begin

Verify the following installations and connections:

- A console connection to the switch. The console connection is optional if you have remote access to the switch.
- Switch cs1 and switch cs2 are powered up and the initial switch setup is complete (the Management IP address and SSH is set up).
- The desired NX-OS version has been installed.
- Inter-switch link (ISL) connections between switches are connected.
- ONTAP node cluster ports are not connected.

Step 1: Install the RCF on the switches

1. Login to switch cs1 using SSH or by using a serial console.
2. Copy the RCF to the bootflash of switch cs1 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 9000 Series NX-OS Command Reference](#).

Show example

This example shows TFTP being used to copy an RCF to the bootflash on switch cs1:

```
cs1# copy tftp: bootflash: vrf management
Enter source filename: Nexus_9336C_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)...
```

3. Apply the RCF previously downloaded to the bootflash.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 9000 Series NX-OS Command Reference](#).

Show example

This example shows the RCF file `Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt` being installed on switch cs1:

```
cs1# copy Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt running-
config echo-commands
```

4. Examine the banner output from the `show banner motd` command. You must read and follow these instructions to ensure the proper configuration and operation of the switch.

Show example

```
cs1# show banner motd

*****
*****
* NetApp Reference Configuration File (RCF)
*
* Switch   : Nexus N9K-C9336C-FX2
* Filename : Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt
* Date     : 10-23-2020
* Version  : v1.6
*
* Port Usage:
* Ports 1- 3: Breakout mode (4x10G) Intra-Cluster Ports, int
e1/1/1-4, e1/2/1-4
, e1/3/1-4
* Ports 4- 6: Breakout mode (4x25G) Intra-Cluster/HA Ports, int
e1/4/1-4, e1/5/
1-4, e1/6/1-4
* Ports 7-34: 40/100GbE Intra-Cluster/HA Ports, int e1/7-34
* Ports 35-36: Intra-Cluster ISL Ports, int e1/35-36
*
* Dynamic breakout commands:
* 10G: interface breakout module 1 port <range> map 10g-4x
* 25G: interface breakout module 1 port <range> map 25g-4x
*
* Undo breakout commands and return interfaces to 40/100G
configuration in confi
g mode:
* no interface breakout module 1 port <range> map 10g-4x
* no interface breakout module 1 port <range> map 25g-4x
* interface Ethernet <interfaces taken out of breakout mode>
* inherit port-profile 40-100G
* priority-flow-control mode auto
* service-policy input HA
* exit
*
*****
*****
```

5. 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.

6. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
7. After you verify the RCF versions and switch settings are correct, copy the running-config file to the startup-config file.

```
copy running-config startup-config
```

Show example

```
cs1# copy running-config startup-config
[#####] 100% Copy complete
```

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 9000 Series NX-OS Command Reference](#).

8. Reboot switch cs1.

```
cs1# reload
```

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

9. Repeat steps 1 through 7 on switch cs2.
10. Connect the cluster ports of all nodes in the ONTAP cluster to switches cs1 and cs2.

Step 2: Verify the switch connections

1. Verify that the switch ports connected to the cluster ports are **up**.

```
show interface brief
```

Show example

```
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)  --
.
.
```

2. Verify that the cluster nodes are in their correct cluster VLANs using the following commands:

```
show vlan brief
```

```
show interface trunk
```

Show example

```
cs1# show vlan brief
```

VLAN Name	Status	Ports
1 default	active	Pol, Eth1/1, Eth1/2, Eth1/3 Eth1/4, Eth1/5, Eth1/6, Eth1/7 Eth1/8, Eth1/35, Eth1/36 Eth1/9/1, Eth1/9/2, Eth1/9/3 Eth1/9/4, Eth1/10/1, Eth1/10/2 Eth1/10/3, Eth1/10/4
17 VLAN0017	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4 Eth1/5, Eth1/6, Eth1/7, Eth1/8 Eth1/9/1, Eth1/9/2, Eth1/9/3 Eth1/9/4, Eth1/10/1, Eth1/10/2 Eth1/10/3, Eth1/10/4
18 VLAN0018	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4 Eth1/5, Eth1/6, Eth1/7, Eth1/8 Eth1/9/1, Eth1/9/2, Eth1/9/3 Eth1/9/4, Eth1/10/1, Eth1/10/2 Eth1/10/3, Eth1/10/4
31 VLAN0031	active	Eth1/11, Eth1/12, Eth1/13 Eth1/14, Eth1/15, Eth1/16 Eth1/17, Eth1/18, Eth1/19 Eth1/20, Eth1/21, Eth1/22
32 VLAN0032	active	Eth1/23, Eth1/24, Eth1/25

```

Eth1/28                               Eth1/26, Eth1/27,
Eth1/31                               Eth1/29, Eth1/30,
Eth1/34                               Eth1/32, Eth1/33,
33   VLAN0033                         active   Eth1/11, Eth1/12,
Eth1/13                               Eth1/14, Eth1/15,
Eth1/16                               Eth1/17, Eth1/18,
Eth1/19                               Eth1/20, Eth1/21,
Eth1/22                               Eth1/23, Eth1/24,
34   VLAN0034                         active   Eth1/26, Eth1/27,
Eth1/25                               Eth1/29, Eth1/30,
Eth1/28                               Eth1/32, Eth1/33,
Eth1/31
Eth1/34

```

```
cs1# show interface trunk
```

```

-----
Port          Native  Status      Port
              Vlan               Channel
-----
Eth1/1        1       trunking    --
Eth1/2        1       trunking    --
Eth1/3        1       trunking    --
Eth1/4        1       trunking    --
Eth1/5        1       trunking    --
Eth1/6        1       trunking    --
Eth1/7        1       trunking    --
Eth1/8        1       trunking    --
Eth1/9/1      1       trunking    --
Eth1/9/2      1       trunking    --
Eth1/9/3      1       trunking    --
Eth1/9/4      1       trunking    --
Eth1/10/1     1       trunking    --
Eth1/10/2     1       trunking    --
Eth1/10/3     1       trunking    --
Eth1/10/4     1       trunking    --
Eth1/11       33      trunking    --

```

Eth1/12	33	trunking	--
Eth1/13	33	trunking	--
Eth1/14	33	trunking	--
Eth1/15	33	trunking	--
Eth1/16	33	trunking	--
Eth1/17	33	trunking	--
Eth1/18	33	trunking	--
Eth1/19	33	trunking	--
Eth1/20	33	trunking	--
Eth1/21	33	trunking	--
Eth1/22	33	trunking	--
Eth1/23	34	trunking	--
Eth1/24	34	trunking	--
Eth1/25	34	trunking	--
Eth1/26	34	trunking	--
Eth1/27	34	trunking	--
Eth1/28	34	trunking	--
Eth1/29	34	trunking	--
Eth1/30	34	trunking	--
Eth1/31	34	trunking	--
Eth1/32	34	trunking	--
Eth1/33	34	trunking	--
Eth1/34	34	trunking	--
Eth1/35	1	trnk-bndl	Pol
Eth1/36	1	trnk-bndl	Pol
Pol	1	trunking	--

Port	Vlans Allowed on Trunk
------	------------------------

Eth1/1	1,17-18
Eth1/2	1,17-18
Eth1/3	1,17-18
Eth1/4	1,17-18
Eth1/5	1,17-18
Eth1/6	1,17-18
Eth1/7	1,17-18
Eth1/8	1,17-18
Eth1/9/1	1,17-18
Eth1/9/2	1,17-18
Eth1/9/3	1,17-18
Eth1/9/4	1,17-18
Eth1/10/1	1,17-18
Eth1/10/2	1,17-18
Eth1/10/3	1,17-18
Eth1/10/4	1,17-18

```
Eth1/11      31, 33
Eth1/12      31, 33
Eth1/13      31, 33
Eth1/14      31, 33
Eth1/15      31, 33
Eth1/16      31, 33
Eth1/17      31, 33
Eth1/18      31, 33
Eth1/19      31, 33
Eth1/20      31, 33
Eth1/21      31, 33
Eth1/22      31, 33
Eth1/23      32, 34
Eth1/24      32, 34
Eth1/25      32, 34
Eth1/26      32, 34
Eth1/27      32, 34
Eth1/28      32, 34
Eth1/29      32, 34
Eth1/30      32, 34
Eth1/31      32, 34
Eth1/32      32, 34
Eth1/33      32, 34
Eth1/34      32, 34
Eth1/35      1
Eth1/36      1
Po1          1
..
..
..
..
..
```



For specific port and VLAN usage details, refer to the banner and important notes section in your RCF.

3. Verify that the ISL between cs1 and cs2 is functional:

```
show port-channel summary
```


Show example

```
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      Eth1/35 (P)      Eth1/36 (P)
cs1#
```

Step 3: Set up your ONTAP cluster

NetApp recommends that you use System Manager to set up new clusters.

System Manager provides a simple and easy workflow for cluster set up and configuration including assigning a node management IP address, initializing the cluster, creating a local tier, configuring protocols and provisioning initial storage.

Go to [Configure ONTAP on a new cluster with System Manager](#) for setup instructions.

What's next?

[Configure switch health monitoring.](#)

Upgrade your reference configuration file (RCF)

You upgrade your RCF version when you have an existing version of the RCF file installed on your operational switches.

Before you begin

Make sure you have the following:

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).
- The current RCF.
- If you are updating your RCF version, you need a boot configuration in the RCF that reflects the desired boot images.

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.



No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure migrates all of the cluster LIFs to the operational partner switch while performing the steps on the target switch.



Before installing a new switch software version and RCFs, you must erase the switch settings and perform basic configuration. You must be connected to the switch using the serial console or have preserved basic configuration information before erasing the switch settings.

Step 1: Prepare for the upgrade

1. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
              e0a    cs1                      Ethernet1/7      N9K-
C9336C
              e0d    cs2                      Ethernet1/7      N9K-
C9336C
cluster1-02/cdp
              e0a    cs1                      Ethernet1/8      N9K-
C9336C
              e0d    cs2                      Ethernet1/8      N9K-
C9336C
cluster1-03/cdp
              e0a    cs1                      Ethernet1/1/1    N9K-
C9336C
              e0b    cs2                      Ethernet1/1/1    N9K-
C9336C
cluster1-04/cdp
              e0a    cs1                      Ethernet1/1/2    N9K-
C9336C
              e0b    cs2                      Ethernet1/1/2    N9K-
C9336C
cluster1::*>
```

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
```

Show example

```
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
```

```
Health Health Speed (Mbps)
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
```

Show example

```
cluster1::*> network interface show -role cluster
          Logical          Status      Network
Current   Current Is
Vserver   Interface          Admin/Oper Address/Mask      Node
Port      Home
-----
Cluster
cluster1-01  cluster1-01_clus1  up/up      169.254.3.4/23
             e0a      true
cluster1-01  cluster1-01_clus2  up/up      169.254.3.5/23
             e0d      true
cluster1-02  cluster1-02_clus1  up/up      169.254.3.8/23
             e0a      true
cluster1-02  cluster1-02_clus2  up/up      169.254.3.9/23
             e0d      true
cluster1-03  cluster1-03_clus1  up/up      169.254.1.3/23
             e0a      true
cluster1-03  cluster1-03_clus2  up/up      169.254.1.1/23
             e0b      true
cluster1-04  cluster1-04_clus1  up/up      169.254.1.6/23
             e0a      true
cluster1-04  cluster1-04_clus2  up/up      169.254.1.7/23
             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
```

Show example

```
cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                                     Type                               Address
Model
-----
-----
cs1                                       cluster-network                    10.233.205.90
N9K-C9336C
  Serial Number: FOCXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                                     9.3(5)
  Version Source: CDP

cs2                                       cluster-network                    10.233.205.91
N9K-C9336C
  Serial Number: FOCXXXXXXGS
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                                     9.3(5)
  Version Source: CDP
cluster1::*>
```

3. Disable auto-revert on the cluster LIFs.

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

Step 2: Configure ports

1. On cluster switch cs1, shut down the ports connected to the cluster ports of the nodes.

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

```
cs1(config-if-range)# shutdown
```



Make sure to shutdown **all** connected cluster ports to avoid any network connection issues. See the Knowledge Base article [Node out of quorum when migrating cluster LIF during switch OS upgrade](#) for further details.

2. Verify that the cluster LIFs have failed over to the ports hosted on cluster switch cs1. This might take a few seconds.

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

Show example

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

Current Is	Logical	Status	Network	Current
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::*>
```

3. Verify that the cluster is healthy:

```
cluster show
```


Show example

```
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::*>
```

4. If you have not already done so, save a copy of the current switch configuration by copying the output of the following command to a text file:

```
show running-config
```

- a. Record any custom additions between the current running-config and the RCF file in use (such as an SNMP configuration for your organization).
- b. For NX-OS 10.2 and later, use the `show diff running-config` command to compare with the saved RCF file in the bootflash. Otherwise, use a third party diff or compare tool.

5. Save basic configuration details to the `write_erase.cfg` file on the bootflash.

```
switch# show run | i "username admin password" > bootflash:write_erase.cfg
```

```
switch# show run | section "vrf context management" >>
bootflash:write_erase.cfg
```

```
switch# show run | section "interface mgmt0" >> bootflash:write_erase.cfg
```

```
switch# show run | section "switchname" >> bootflash:write_erase.cfg
```

6. Issue the write erase command to erase the current saved configuration:

```
switch# write erase
```

Warning: This command will erase the startup-configuration.

```
Do you wish to proceed anyway? (y/n) [n] y
```

7. Copy the previously saved basic configuration into the startup configuration.

```
switch# copy write_erase.cfg startup-config
```

8. Perform a reboot of the switch:

```
switch# reload
```

This command will reboot the system. (y/n)? [n] y

9. After the management IP address is reachable again, log in to the switch through SSH.

You might need to update host file entries related to the SSH keys.

10. Copy the RCF to the bootflash of switch cs1 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 9000 Series NX-OS Command Reference](#) guides.

Show example

This example shows TFTP being used to copy an RCF to the bootflash on switch cs1:

```
cs1# copy tftp: bootflash: vrf management
Enter source filename: Nexus_9336C_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)...
```

11. Apply the RCF previously downloaded to the bootflash.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 9000 Series NX-OS Command Reference](#) guides.

Show example

This example shows the RCF file `Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt` being installed on switch cs1:

```
cs1# copy Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt running-
config echo-commands
```

12. Examine the banner output from the `show banner motd` command. You must read and follow these instructions to ensure the proper configuration and operation of the switch.

Show example

```
cs1# show banner motd

*****
*****
* NetApp Reference Configuration File (RCF)
*
* Switch    : Nexus N9K-C9336C-FX2
* Filename  : Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt
* Date      : 10-23-2020
* Version   : v1.6
*
* Port Usage:
* Ports 1- 3: Breakout mode (4x10G) Intra-Cluster Ports, int
e1/1/1-4, e1/2/1-4
, e1/3/1-4
* Ports 4- 6: Breakout mode (4x25G) Intra-Cluster/HA Ports, int
e1/4/1-4, e1/5/
1-4, e1/6/1-4
* Ports 7-34: 40/100GbE Intra-Cluster/HA Ports, int e1/7-34
* Ports 35-36: Intra-Cluster ISL Ports, int e1/35-36
*
* Dynamic breakout commands:
* 10G: interface breakout module 1 port <range> map 10g-4x
* 25G: interface breakout module 1 port <range> map 25g-4x
*
* Undo breakout commands and return interfaces to 40/100G
configuration in confi
g mode:
* no interface breakout module 1 port <range> map 10g-4x
* no interface breakout module 1 port <range> map 25g-4x
* interface Ethernet <interfaces taken out of breakout mode>
* inherit port-profile 40-100G
* priority-flow-control mode auto
* service-policy input HA
* exit
*
*****
*****
```

13. 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.

14. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
15. After you verify the RCF versions, custom additions, 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 9000 Series NX-OS Command Reference](#) guides.

```
cs1# copy running-config startup-config
```

```
[ ] 100% Copy complete
```

16. Reboot switch cs1. You can ignore the “cluster switch health monitor” alerts and “cluster ports down” events reported on the nodes while the switch reboots.

```
cs1# reload
```

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

17. Verify the health of cluster ports on the cluster.
 - a. Verify that cluster ports are up and healthy across all nodes in the cluster:

```
network port show -role cluster
```

Show example

```
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
```

```
Health Health Speed (Mbps)
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.

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
          e0a   cs1                        Ethernet1/7
N9K-C9336C
          e0d   cs2                        Ethernet1/7
N9K-C9336C
cluster01-2/cdp
          e0a   cs1                        Ethernet1/8
N9K-C9336C
          e0d   cs2                        Ethernet1/8
N9K-C9336C
cluster01-3/cdp
          e0a   cs1                        Ethernet1/1/1
N9K-C9336C
          e0b   cs2                        Ethernet1/1/1
N9K-C9336C
cluster1-04/cdp
          e0a   cs1                        Ethernet1/1/2
N9K-C9336C
          e0b   cs2                        Ethernet1/1/2
N9K-C9336C

cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch          Type          Address
Model
-----
-----
cs1              cluster-network  10.233.205.90
NX9-C9336C
  Serial Number: FOCXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                  9.3(5)
  Version Source: CDP

cs2              cluster-network  10.233.205.91
```

```

NX9-C9336C
  Serial Number: FOCXXXXXXGS
    Is Monitored: true
      Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                9.3(5)
  Version Source: CDP

2 entries were displayed.

```

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.

```

18. Verify that the cluster is healthy:

```
cluster show
```

Show example

```

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::*>

```

19. Repeat steps 1 to 18 on switch cs2.

20. Enable auto-revert on the cluster LIFs.

```

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

```


Step 3: Verify the cluster network configuration and cluster health

1. Verify that the switch ports connected to the cluster ports are **up**.

```
show interface brief
```

Show example

```
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)  --
.
.
```

2. Verify that the expected nodes are still connected:

```
show cdp neighbors
```

Show example

```
cs1# 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
node1             Eth1/1        133     H           FAS2980
e0a
node2             Eth1/2        133     H           FAS2980
e0a
cs1               Eth1/35       175     R S I s     N9K-C9336C
Eth1/35
cs1               Eth1/36       175     R S I s     N9K-C9336C
Eth1/36

Total entries displayed: 4
```

3. Verify that the cluster nodes are in their correct cluster VLANs using the following commands:

```
show vlan brief
```

```
show interface trunk
```

Show example

```
cs1# show vlan brief
```

VLAN Name	Status	Ports
1 default	active	Pol, Eth1/1, Eth1/2, Eth1/3 Eth1/4, Eth1/5, Eth1/6, Eth1/7 Eth1/8, Eth1/35, Eth1/36 Eth1/9/1, Eth1/9/2, Eth1/9/3 Eth1/9/4, Eth1/10/1, Eth1/10/2 Eth1/10/3, Eth1/10/4
17 VLAN0017	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4 Eth1/5, Eth1/6, Eth1/7, Eth1/8 Eth1/9/1, Eth1/9/2, Eth1/9/3 Eth1/9/4, Eth1/10/1, Eth1/10/2 Eth1/10/3, Eth1/10/4
18 VLAN0018	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4 Eth1/5, Eth1/6, Eth1/7, Eth1/8 Eth1/9/1, Eth1/9/2, Eth1/9/3 Eth1/9/4, Eth1/10/1, Eth1/10/2 Eth1/10/3, Eth1/10/4
31 VLAN0031	active	Eth1/11, Eth1/12, Eth1/13 Eth1/14, Eth1/15, Eth1/16 Eth1/17, Eth1/18, Eth1/19 Eth1/20, Eth1/21, Eth1/22
32 VLAN0032	active	Eth1/23, Eth1/24, Eth1/25

```

Eth1/28                               Eth1/26, Eth1/27,
Eth1/31                               Eth1/29, Eth1/30,
Eth1/34                               Eth1/32, Eth1/33,
33   VLAN0033                         active   Eth1/11, Eth1/12,
Eth1/13                               Eth1/14, Eth1/15,
Eth1/16                               Eth1/17, Eth1/18,
Eth1/19                               Eth1/20, Eth1/21,
Eth1/22                               Eth1/23, Eth1/24,
34   VLAN0034                         active   Eth1/26, Eth1/27,
Eth1/25                               Eth1/29, Eth1/30,
Eth1/28                               Eth1/32, Eth1/33,
Eth1/31
Eth1/34

```

```
cs1# show interface trunk
```

```

-----
Port                Native  Status      Port
                   Vlan              Channel
-----
Eth1/1              1      trunking    --
Eth1/2              1      trunking    --
Eth1/3              1      trunking    --
Eth1/4              1      trunking    --
Eth1/5              1      trunking    --
Eth1/6              1      trunking    --
Eth1/7              1      trunking    --
Eth1/8              1      trunking    --
Eth1/9/1            1      trunking    --
Eth1/9/2            1      trunking    --
Eth1/9/3            1      trunking    --
Eth1/9/4            1      trunking    --
Eth1/10/1           1      trunking    --
Eth1/10/2           1      trunking    --
Eth1/10/3           1      trunking    --
Eth1/10/4           1      trunking    --
Eth1/11             33     trunking    --

```

Eth1/12	33	trunking	--
Eth1/13	33	trunking	--
Eth1/14	33	trunking	--
Eth1/15	33	trunking	--
Eth1/16	33	trunking	--
Eth1/17	33	trunking	--
Eth1/18	33	trunking	--
Eth1/19	33	trunking	--
Eth1/20	33	trunking	--
Eth1/21	33	trunking	--
Eth1/22	33	trunking	--
Eth1/23	34	trunking	--
Eth1/24	34	trunking	--
Eth1/25	34	trunking	--
Eth1/26	34	trunking	--
Eth1/27	34	trunking	--
Eth1/28	34	trunking	--
Eth1/29	34	trunking	--
Eth1/30	34	trunking	--
Eth1/31	34	trunking	--
Eth1/32	34	trunking	--
Eth1/33	34	trunking	--
Eth1/34	34	trunking	--
Eth1/35	1	trnk-bndl	Pol
Eth1/36	1	trnk-bndl	Pol
Pol	1	trunking	--

Port	Vlans Allowed on Trunk
------	------------------------

Eth1/1	1,17-18
Eth1/2	1,17-18
Eth1/3	1,17-18
Eth1/4	1,17-18
Eth1/5	1,17-18
Eth1/6	1,17-18
Eth1/7	1,17-18
Eth1/8	1,17-18
Eth1/9/1	1,17-18
Eth1/9/2	1,17-18
Eth1/9/3	1,17-18
Eth1/9/4	1,17-18
Eth1/10/1	1,17-18
Eth1/10/2	1,17-18
Eth1/10/3	1,17-18
Eth1/10/4	1,17-18

```
Eth1/11      31, 33
Eth1/12      31, 33
Eth1/13      31, 33
Eth1/14      31, 33
Eth1/15      31, 33
Eth1/16      31, 33
Eth1/17      31, 33
Eth1/18      31, 33
Eth1/19      31, 33
Eth1/20      31, 33
Eth1/21      31, 33
Eth1/22      31, 33
Eth1/23      32, 34
Eth1/24      32, 34
Eth1/25      32, 34
Eth1/26      32, 34
Eth1/27      32, 34
Eth1/28      32, 34
Eth1/29      32, 34
Eth1/30      32, 34
Eth1/31      32, 34
Eth1/32      32, 34
Eth1/33      32, 34
Eth1/34      32, 34
Eth1/35      1
Eth1/36      1
Po1          1
..
..
..
..
..
```



For specific port and VLAN usage details, refer to the banner and important notes section in your RCF.

4. Verify that the ISL between cs1 and cs2 is functional:

```
show port-channel summary
```

Show example

```
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        Eth1/35 (P)      Eth1/36 (P)
cs1#
```

5. Verify that the cluster LIFs have reverted to their home port:

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

Show example

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

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

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

6. Verify that the cluster is healthy:

```
cluster show
```


Show example

```
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::*>
```

7. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

Packet	Source	Destination
Node	Date	LIF
Loss		
-----	-----	-----
node1		
3/5/2022 19:21:18 -06:00	cluster1-01_clus2	cluster1-02-
clus1 none		
3/5/2022 19:21:20 -06:00	cluster1-01_clus2	cluster1-
02_clus2 none		
node2		
3/5/2022 19:21:18 -06:00	cluster1-02_clus2	cluster1-
01_clus1 none		
3/5/2022 19:21:20 -06:00	cluster1-02_clus2	cluster1-
01_clus2 none		

All ONTAP releases

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

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

```

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

```

What's next?

[Configure switch health monitoring.](#)

Migrate switches

Migrate from a switchless cluster with direct-attached storage

You can migrate from a switchless cluster with direct-attached storage by adding two new shared switches.

The procedure you use depends on whether you have two dedicated cluster-network ports on each controller or a single cluster port on each controller. The process documented works for all nodes using optical or Twinax ports, but is not supported on this switch if nodes are using onboard 10Gb BASE-T RJ45 ports for the cluster-network ports.

Most systems require two dedicated cluster-network ports on each controller. See [Cisco Ethernet Switches](#) for more information.

If you have an existing two-node switchless cluster environment, you can migrate to a two-node switched cluster environment using Cisco Nexus 9336C-FX2 switches to enable you to scale beyond two nodes in the cluster.

Review requirements

Ensure that:

- For the two-node switchless configuration:
 - The two-node switchless configuration is properly set up and functioning.
 - The nodes are running ONTAP 9.8 and later.
 - All cluster ports are in the **up** state.
 - All cluster logical interfaces (LIFs) are in the **up** state and on their **home** ports.
- For the Cisco Nexus 9336C-FX2 switch configuration:
 - Both switches have management network connectivity.
 - There is console access to the cluster switches.
 - Nexus 9336C-FX2 node-to-node switch and switch-to-switch connections use Twinax or fiber cables.
 - The NetApp [Hardware Universe](#) contains more information about cabling.
 - Inter-Switch Link (ISL) cables are connected to ports 1/35 and 1/36 on both 9336C-FX2 switches.
- Initial customization of the 9336C-FX2 switches are completed. So that the:
 - 9336C-FX2 switches are running the latest version of software
 - Reference Configuration Files (RCFs) have been applied to the switches
 - Any site customization, such as SMTP, SNMP, and SSH is configured on the new switches.

Migrate the switches

About the examples

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

- The names of the 9336C-FX2 switches are *cs1* and *cs2*.
- The names of the cluster SVMs are *node1* and *node2*.
- The names of the LIFs are *node1_clus1* and *node1_clus2* on node 1, and *node2_clus1* and *node2_clus2* on node 2 respectively.

- The cluster1::*> prompt indicates the name of the cluster.
- The cluster ports used in this procedure are e3a and e3b, as per the AFF A400 controller. The [Hardware Universe](#) contains the latest information about the actual cluster ports for your platforms.

Step 1: Migrate from a switchless cluster with direct-attached

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.`

where x is the duration of the maintenance window in hours.



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

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

```
set -privilege advanced
```

The advanced prompt (*>) appears.

3. Disable all node-facing ports (not ISL ports) on both the new cluster switches cs1 and cs2. You must not disable the ISL ports.

Show example

The following example shows that node-facing ports 1 through 34 are disabled on switch cs1:

```
cs1# config
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# interface e1/1-34
cs1(config-if-range)# shutdown
```

4. Verify that the ISL and the physical ports on the ISL between the two 9336C-FX2 switches cs1 and cs2 are up on ports 1/35 and 1/36:

```
show port-channel summary
```

Show example

The following example shows that the ISL ports are up on switch cs1:

```
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      Eth1/35(P)  Eth1/36(P)
```

The following example shows that the ISL ports are up on switch cs2:

```
cs2# 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      Eth1/35(P)  Eth1/36(P)
```

5. Display the list of neighboring devices:

```
show cdp neighbors
```

This command provides information about the devices that are connected to the system.

Show example

The following example lists the neighboring devices on switch cs1:

```
cs1# 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
cs2            Eth1/35       175    R S I s      N9K-C9336C
Eth1/35
cs2            Eth1/36       175    R S I s      N9K-C9336C
Eth1/36
Total entries displayed: 2
```

The following example lists the neighboring devices on switch cs2:

```
cs2# 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
cs1            Eth1/35       177    R S I s      N9K-C9336C
Eth1/35
cs1            ) Eth1/36       177    R S I s      N9K-C9336C
Eth1/36
Total entries displayed: 2
```

6. Verify that all cluster ports are up:

```
network port show - ipspace Cluster
```

Each port should display up for Link and healthy for Health Status.

Show example

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

Node: node1

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper
Status
-----
e3a       Cluster      Cluster      up   9000  auto/100000
healthy
e3b       Cluster      Cluster      up   9000  auto/100000
healthy

Node: node2

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper
Status
-----
e3a       Cluster      Cluster      up   9000  auto/100000
healthy
e3b       Cluster      Cluster      up   9000  auto/100000
healthy
4 entries were displayed.
```

7. Verify that all cluster LIFs are up and operational:

```
network interface show - vserver Cluster
```

Each cluster LIF should display true for `Is Home` and have a `Status Admin/Oper` of `up/up`.

Show example

```
cluster1::*> network interface show -vserver Cluster
          Logical      Status      Network      Current
Current Is
Vserver   Interface  Admin/Oper  Address/Mask  Node
Port      Home
-----
Cluster
          node1_clus1  up/up      169.254.209.69/16  node1
e3a       true
          node1_clus2  up/up      169.254.49.125/16  node1
e3b       true
          node2_clus1  up/up      169.254.47.194/16  node2
e3a       true
          node2_clus2  up/up      169.254.19.183/16  node2
e3b       true
4 entries were displayed.
```

8. Verify that auto-revert is enabled on all cluster LIFs:

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

Show example

```
cluster1::*> network interface show -vserver Cluster -fields auto-
revert
          Logical
Vserver   Interface  Auto-revert
-----
Cluster
          node1_clus1  true
          node1_clus2  true
          node2_clus1  true
          node2_clus2  true
4 entries were displayed.
```

9. Disconnect the cable from cluster port e3a on node1, and then connect e3a to port 1 on cluster switch cs1, using the appropriate cabling supported by the 9336C-FX2 switches.

The NetApp [Hardware Universe](#) contains more information about cabling.

10. Disconnect the cable from cluster port e3a on node2, and then connect e3a to port 2 on cluster switch cs1, using the appropriate cabling supported by the 9336C-FX2 switches.
11. Enable all node-facing ports on cluster switch cs1.

Show example

The following example shows that ports 1/1 through 1/34 are enabled on switch cs1:

```
cs1# config
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# interface e1/1-34
cs1(config-if-range)# no shutdown
```

12. Verify that all cluster LIFs are **up**, operational, and display as true for Is Home:

```
network interface show - vserver Cluster
```

Show example

The following example shows that all the LIFs are **up** on node1 and node2 and that Is Home results are **true**:

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

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			
-----	-----	-----	-----	-----
-----	-----			
Cluster				
true	node1_clus1	up/up	169.254.209.69/16	node1 e3a
true	node1_clus2	up/up	169.254.49.125/16	node1 e3b
true	node2_clus1	up/up	169.254.47.194/16	node2 e3a
true	node2_clus2	up/up	169.254.19.183/16	node2 e3b
true				

4 entries were displayed.

13. Display information about the status of the nodes in the cluster:

```
cluster show
```

Show example

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

```
cluster1::*> cluster show
Node           Health  Eligibility  Epsilon
-----
node1          true    true         false
node2          true    true         false
2 entries were displayed.
```

14. Disconnect the cable from cluster port e3b on node1, and then connect e3b to port 1 on cluster switch cs2, using the appropriate cabling supported by the 9336C-FX2 switches.
15. Disconnect the cable from cluster port e3b on node2, and then connect e3b to port 2 on cluster switch cs2, using the appropriate cabling supported by the 9336C-FX2 switches.
16. Enable all node-facing ports on cluster switch cs2.

Show example

The following example shows that ports 1/1 through 1/34 are enabled on switch cs2:

```
cs2# config
Enter configuration commands, one per line. End with CNTL/Z.
cs2(config)# interface e1/1-34
cs2(config-if-range)# no shutdown
```

17. Verify that all cluster ports are up:

```
network port show - ipspace Cluster
```

Show example

The following example shows that all the cluster ports are up on node1 and node2:

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

Node: node1

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
e3a         Cluster   Cluster      up    9000  auto/100000
healthy    false
e3b         Cluster   Cluster      up    9000  auto/100000
healthy    false

Node: node2

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
e3a         Cluster   Cluster      up    9000  auto/100000
healthy    false
e3b         Cluster   Cluster      up    9000  auto/100000
healthy    false
4 entries were displayed.
```

18. Verify that all interfaces display true for Is Home:

```
network interface show - vserver Cluster
```



This might take several minutes to complete.

Show example

The following example shows that all LIFs are **up** on node1 and node2 and that Is Home results are true:

```
cluster1::*> network interface show -vserver Cluster
          Logical      Status      Network      Current
Current Is
Vserver  Interface      Admin/Oper  Address/Mask  Node      Port
Home
-----
-----
Cluster
true     node1_clus1    up/up      169.254.209.69/16  node1     e3a
true     node1_clus2    up/up      169.254.49.125/16  node1     e3b
true     node2_clus1    up/up      169.254.47.194/16  node2     e3a
true     node2_clus2    up/up      169.254.19.183/16  node2     e3b
true
4 entries were displayed.
```

19. Verify that both nodes each have one connection to each switch:

```
show cdp neighbors
```

Show example

The following example shows the appropriate results for both switches:

```
cs1# 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
node1          Eth1/1         133     H           AFFA400
e3a
node2          Eth1/2         133     H           AFFA400
e3a
cs2            Eth1/35        175     R S I s     N9K-C9336C
Eth1/35
cs2            Eth1/36        175     R S I s     N9K-C9336C
Eth1/36
Total entries displayed: 4
cs2# 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
node1          Eth1/1         133     H           AFFA400
e3b
node2          Eth1/2         133     H           AFFA400
e3b
cs1            Eth1/35        175     R S I s     N9K-C9336C
Eth1/35
cs1            Eth1/36        175     R S I s     N9K-C9336C
Eth1/36
Total entries displayed: 4
```

20. Display information about the discovered network devices in your cluster:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port  Device (LLDP: ChassisID)  Interface
Platform
-----
node2         /cdp
              e3a   cs1                       0/2          N9K-
C9336C
              e3b   cs2                       0/2          N9K-
C9336C
node1         /cdp
              e3a   cs1                       0/1          N9K-
C9336C
              e3b   cs2                       0/1          N9K-
C9336C
4 entries were displayed.
```

21. Verify that the storage configuration of HA pair 1 (and HA pair 2) is correct and error free:

```
system switch ethernet show
```

Show example

```
storage::*> system switch ethernet show
Switch                Type                Address
Model
-----
sh1
                        storage-network      172.17.227.5
C9336C
    Serial Number: FOC221206C2
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        9.3(5)
    Version Source: CDP
sh2
                        storage-network      172.17.227.6
C9336C
    Serial Number: FOC220443LZ
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        9.3(5)
    Version Source: CDP
2 entries were displayed.
storage::*>
```

22. Verify that the settings are disabled:

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



It might take several minutes for the command to complete. Wait for the '3-minute lifetime to expire' announcement.

The `false` output in the following example shows that the configuration settings are disabled:

Show example

```
cluster1::*> network options switchless-cluster show  
Enable Switchless Cluster: false
```

23. Verify the status of the node members in the cluster:

```
cluster show
```

Show example

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

```
cluster1::*> cluster show  
Node           Health  Eligibility  Epsilon  
-----  
node1          true    true         false  
node2          true    true         false
```

24. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e3a
Cluster node1_clus2 169.254.49.125 node1 e3b
Cluster node2_clus1 169.254.47.194 node2 e3a
Cluster node2_clus2 169.254.19.183 node2 e3b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

25. Change the privilege level back to admin:

```
set -privilege admin
```

Step 2: Set up the shared switch

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

- The names of the two shared switches are *sh1* and *sh2*.
- The nodes are *node1* and *node2*.



The procedure requires the use of both ONTAP commands and Cisco Nexus 9000 Series Switches commands, ONTAP commands are used unless otherwise indicated.

1. Verify that the storage configuration of HA pair 1 (and HA pair 2) is correct and error free:

```
system switch ethernet show
```

Show example

```
storage::*> system switch ethernet show
Switch                Type                Address
Model
-----
sh1
                    storage-network    172.17.227.5
C9336C

    Serial Number: FOC221206C2
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                    9.3(5)
    Version Source: CDP
sh2
                    storage-network    172.17.227.6
C9336C

    Serial Number: FOC220443LZ
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                    9.3(5)
    Version Source: CDP
2 entries were displayed.
storage::*>
```

2. Verify that the storage node ports are healthy and operational:

```
storage port show -port-type ENET
```

Show example

```
storage::*> storage port show -port-type ENET
                               Speed
VLAN
Node   Port   Type   Mode   (Gb/s)   State   Status
ID
-----
node1
30     e0c    ENET   storage 100      enabled online
30     e0d    ENET   storage 100      enabled online
30     e5a    ENET   storage 100      enabled online
30     e5b    ENET   storage 100      enabled online
node2
30     e0c    ENET   storage 100      enabled online
30     e0d    ENET   storage 100      enabled online
30     e5a    ENET   storage 100      enabled online
30     e5b    ENET   storage 100      enabled online
```

3. Move the HA pair 1, NSM224 path A ports to sh1 port range 11-22.
4. Install a cable from HA pair 1, node1, path A to sh1 port range 11-22. For example, the path A storage port on an AFF A400 is e0c.
5. Install a cable from HA pair 1, node2, path A to sh1 port range 11-22.
6. Verify that the node ports are healthy and operational:

```
storage port show -port-type ENET
```

Show example

```
storage::*> storage port show -port-type ENET
                               Speed
VLAN
Node   Port   Type   Mode   (Gb/s)   State   Status
ID
-----
node1
30     e0c    ENET   storage 100      enabled online
30     e0d    ENET   storage 0        enabled offline
30     e5a    ENET   storage 0        enabled offline
30     e5b    ENET   storage 100      enabled online
node2
30     e0c    ENET   storage 100      enabled online
30     e0d    ENET   storage 0        enabled offline
30     e5a    ENET   storage 0        enabled offline
30     e5b    ENET   storage 100      enabled online
```

7. Check that there are no storage switch or cabling issues with the cluster:

```
system health alert show -instance
```

Show example

```
storage::*> system health alert show -instance
There are no entries matching your query.
```

8. Move the HA pair 1, NSM224 path B ports to sh2 port range 11-22.
9. Install a cable from HA pair 1, node1, path B to sh2 port range 11-22. For example, the path B storage port on an AFF A400 is e5b.
10. Install a cable from HA pair 1, node2, path B to sh2 port range 11-22.

11. Verify that the node ports are healthy and operational:

```
storage port show -port-type ENET
```

Show example

```
storage::*> storage port show -port-type ENET
                               Speed
VLAN
Node   Port   Type   Mode   (Gb/s)   State   Status
ID
-----
node1
30     e0c    ENET   storage 100      enabled online
30     e0d    ENET   storage 0        enabled offline
30     e5a    ENET   storage 0        enabled offline
30     e5b    ENET   storage 100     enabled online
node2
30     e0c    ENET   storage 100     enabled online
30     e0d    ENET   storage 0        enabled offline
30     e5a    ENET   storage 0        enabled offline
30     e5b    ENET   storage 100     enabled online
```

12. Verify that the storage configuration of HA pair 1 is correct and error free:

```
system switch ethernet show
```

Show example

```
storage::*> system switch ethernet show
Switch                               Type                               Address
Model
-----
sh1
                                     storage-network                    172.17.227.5
C9336C
    Serial Number: FOC221206C2
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                                     9.3(5)
    Version Source: CDP
sh2
                                     storage-network                    172.17.227.6
C9336C
    Serial Number: FOC220443LZ
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                                     9.3(5)
    Version Source: CDP
2 entries were displayed.
storage::*>
```

13. Reconfigure the unused (controller) secondary storage ports on HA pair 1 from storage to networking. If more than one NS224 was direct attached, there will be ports that should be reconfigured.

Show example

```
storage port modify -node [node name] -port [port name] -mode
network
```

To place storage ports into a broadcast domain:

- `network port broadcast-domain create` (to create a new domain, if needed)
- `network port broadcast-domain add-ports` (to add ports to an existing domain)

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

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

What's next?

[Configure switch health monitoring.](#)

Migrate from a switched configuration with direct-attached storage

You can migrate from a switched configuration with direct-attached storage by adding two new shared switches.

Supported switches

The following switches are supported:

- Nexus 9336C-FX2
- Nexus 3232C

The ONTAP and NX-OS versions supported in this procedure are on the Cisco Ethernet Switches page. See [Cisco Ethernet switches](#).

Connection Ports

The switches use the following ports to connect to nodes:

- Nexus 9336C-FX2:
 - Ports 1- 3: Breakout mode (4x10G) Intra-Cluster Ports, int e1/1/1-4, e1/2/1-4, e1/3/1-4
 - Ports 4- 6: Breakout mode (4x25G) Intra-Cluster/HA Ports, int e1/4/1-4, e1/5/1-4, e1/6/1-4
 - Ports 7-34: 40/100GbE Intra-Cluster/HA Ports, int e1/7-34
- Nexus 3232C:
 - Ports 1-30: 10/40/100 GbE
- The switches use the following Inter-Switch Link (ISL) ports:
 - Ports int e1/35-36: Nexus 9336C-FX2
 - Ports e1/31-32: Nexus 3232C

The [Hardware Universe](#) contains information about supported cabling for all cluster switches.

What you'll need

- Make sure you completed the following tasks:
 - Configured some of the ports on Nexus 9336C-FX2 switches to run at 100 GbE.
 - Planned, migrated, and documented 100 GbE connectivity from nodes to Nexus 9336C-FX2 switches.
 - Migrated nondisruptively other Cisco cluster switches from an ONTAP cluster to Cisco Nexus 9336C-FX2 network switches.
- The existing switch network is properly set up and functioning.
- All ports are in the **up** state to ensure nondisruptive operations.
- The Nexus 9336C-FX2 switches are configured and operating under the proper version of NX-OS installed and reference configuration file (RCF) applied.

- The existing network configuration has the following:
 - A redundant and fully functional NetApp cluster using both older Cisco switches.
 - Management connectivity and console access to both the older Cisco switches and the new switches.
 - All cluster LIFs in the **up** state with the cluster LIFs are on their home ports.
 - ISL ports enabled and cabled between the other Cisco switches and between the new switches.

About the examples

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

- The existing Cisco Nexus 3232C cluster switches are *c1* and *c2*.
- The new Nexus 9336C-FX2 switches are *sh1* and *sh2*.
- The nodes are *node1* and *node2*.
- The cluster LIFs are *node1_clus1* and *node1_clus2* on node 1, and *node2_clus1* and *node2_clus2* on node 2 respectively.
- Switch *c2* is replaced by switch *sh2* first and then switch *c1* is replaced by switch *sh1*.

Steps

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

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

Where x is the duration of the maintenance window in hours.

2. Check the administrative and operational status of each cluster port.
3. Verify that all the cluster ports are up with a healthy status:

```
network port show -role cluster
```

Show example

```
cluster1::*> network port show -role cluster
Node: node1

Ignore
Speed (Mbps)  Health
Health
Port  IPspace  Broadcast Domain Link MTU  Admin/Ope  Status
Status
-----
-----
e3a   Cluster  Cluster      up   9000  auto/100000 healthy
false
e3b   Cluster  Cluster      up   9000  auto/100000 healthy
false

Node: node2

Ignore
Speed (Mbps)  Health
Health
Port  IPspace  Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e3a   Cluster  Cluster      up   9000  auto/100000 healthy
false
e3b   Cluster  Cluster      up   9000  auto/100000 healthy
false
4 entries were displayed.
cluster1::*>
```

4. Verify that all the cluster interfaces (LIFs) are on the home port:

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

Show example

```
cluster1::*> network interface show -role cluster
          Logical      Status      Network      Current
Current Is
Vserver  Interface      Admin/Oper  Address/Mask  Node      Port
Home
-----  -----
-----  -----
Cluster
true     node1_clus1  up/up      169.254.3.4/23  node1     e3a
true     node1_clus2  up/up      169.254.3.5/23  node1     e3b
true     node2_clus1  up/up      169.254.3.8/23  node2     e3a
true     node2_clus2  up/up      169.254.3.9/23  node2     e3b
true
4 entries were displayed.
cluster1::*>
```

5. Verify that the cluster displays information for both cluster switches:

```
system cluster-switch show -is-monitoring-enabled-operational true
```

Show example

```
cluster1::*> system cluster-switch show -is-monitoring-enabled  
-operational true
```

Switch	Type	Address	Model
sh1 C9336C	cluster-network	10.233.205.90	N9K-

```
Serial Number: FOCXXXXXXGD  
Is Monitored: true  
Reason: None  
Software Version: Cisco Nexus Operating System (NX-OS) Software,  
Version  
9.3(5)  
Version Source: CDP
```

sh2 C9336C	cluster-network	10.233.205.91	N9K-
---------------	-----------------	---------------	------

```
Serial Number: FOCXXXXXXGS  
Is Monitored: true  
Reason: None  
Software Version: Cisco Nexus Operating System (NX-OS) Software,  
Version  
9.3(5)  
Version Source: CDP  
cluster1::*>
```

6. Disable auto-revert on the cluster LIFs.

Show example

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

7. Shut down the c2 switch.

Show example

```
c2# configure terminal  
Enter configuration commands, one per line. End with CNTL/Z.  
c2(config)# interface ethernet <int range>  
c2(config)# shutdown
```

8. Verify that the cluster LIFs have migrated to the ports hosted on cluster switch sh1:

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

This might take a few seconds.

Show example

```
cluster1::*> network interface show -role cluster  
      Logical      Status      Network      Current  
Current Is  
Vserver Interface Admin/Oper Address/Mask Node      Port  
Home  
-----  
-----  
Cluster  
      node1_clus1 up/up      169.254.3.4/23 node1      e3a  
true  
      node1_clus2 up/up      169.254.3.5/23 node1      e3a  
false  
      node2_clus1 up/up      169.254.3.8/23 node2      e3a  
true  
      node2_clus2 up/up      169.254.3.9/23 node2      e3a  
false  
4 entries were displayed.  
cluster1::*>
```

9. Replace switch c2 with the new switch sh2 and re-cable the new switch.
10. Verify that the ports are back up on sh2. **Note** that the LIFs are still on switch c1.
11. Shut down the c1 switch.

Show example

```
c1# configure terminal  
Enter configuration commands, one per line. End with CNTL/Z.  
c1(config)# interface ethernet <int range>  
c1(config)# shutdown
```

12. Verify that the cluster LIFs have migrated to the ports hosted on cluster switch sh2. This might take a few seconds.

Show example

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

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

Cluster				
node1_clus1	up/up	169.254.3.4/23	node1	e3a
true				
node1_clus2	up/up	169.254.3.5/23	node1	e3a
false				
node2_clus1	up/up	169.254.3.8/23	node2	e3a
true				
node2_clus2	up/up	169.254.3.9/23	node2	e3a
false				

```
4 entries were displayed.  
cluster1::*>
```

13. Replace switch c1 with the new switch sh1 and re-cable the new switch.
14. Verify that the ports are back up on sh1. **Note** that the LIFs are still on switch c2.
15. Enable auto-revert on the cluster LIFs:

Show example

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto  
-revert True
```

16. Verify that the cluster is healthy:

```
cluster show
```

Show example


```
cluster1::*> cluster show
Node           Health  Eligibility  Epsilon
-----
node1          true    true         false
node2          true    true         false
2 entries were displayed.
cluster1::*>
```

Migrate from a switchless configuration with switch-attached storage by reusing the storage switches

You can migrate from a switchless configuration with switch-attached storage by reusing the storage switches.

By reusing the storage switches the storage switches of HA pair 1 become the shared switches as shown in the following figure.

Switch Attached



Steps

1. Verify that the storage configuration of HA pair 1 (and HA pair 2) is correct and error free:

```
system switch ethernet show
```

Show example

```
storage::*> system switch ethernet show
Switch                               Type                               Address
Model
-----
sh1
                                storage-network                    172.17.227.5
C9336C

    Serial Number: FOC221206C2
    Is Monitored: true
    Reason: none
    Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                                9.3(5)
    Version Source: CDP
sh2
                                storage-network                    172.17.227.6
C9336C

    Serial Number: FOC220443LZ
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                                9.3(5)
    Version Source: CDP
2 entries were displayed.
storage::*>
```

2. Verify that the node ports are healthy and operational:

```
storage port show -port-type ENET
```

Show example

```
storage::*> storage port show -port-type ENET
                               Speed
VLAN Node   Port   Type   Mode   (Gb/s)  State  Status
ID
-----
node1
30    e0c    ENET   storage 100    enabled online
30    e0d    ENET   storage 100    enabled online
30    e5a    ENET   storage 100    enabled online
30    e5b    ENET   storage 100    enabled online
node2
30    e0c    ENET   storage 100    enabled online
30    e0d    ENET   storage 100    enabled online
30    e5a    ENET   storage 100    enabled online
30    e5b    ENET   storage 100    enabled online
```

3. Move the HA pair 1, NSM224 path A cables from storage switch A to the shared NS224 storage ports for HA pair 1, path A on storage switch A.
4. Move the cable from HA pair 1, node A, path A to the shared storage port for HA pair 1, node A on storage switch A.
5. Move the cable from HA pair 1, node B, path A to the shared storage port for HA pair 1, node B on storage switch A.
6. Verify the storage attached to HA pair 1, storage switch A is healthy:

```
system health alert show -instance
```

Show example

```
storage::*> system health alert show -instance  
There are no entries matching your query.
```

7. Replace the storage RCF on shared switch A with the shared RCF file. See [Install the RCF on a Cisco Nexus 9336C-FX2 shared switch](#) for further details.
8. Verify the storage attached to HA pair 1, storage switch B is healthy:

```
system health alert show -instance
```

Show example

```
storage::*> system health alert show -instance  
There are no entries matching your query.
```

9. Move the HA pair 1, NSM224 path B cables from storage switch B to the shared NS224 storage ports for HA pair 1, path B to storage switch B.
10. Move the cable from HA pair 1, node A, path B to the shared storage port for HA pair 1, node A, path B on storage switch B.
11. Move the cable from HA pair 1, node B, path B to the shared storage port for HA pair 1, node B, path B on storage switch B.
12. Verify the storage attached to HA pair 1, storage switch B is healthy:

```
system health alert show -instance
```

Show example

```
storage::*> system health alert show -instance  
There are no entries matching your query.
```

13. Replace the storage RCF file on shared switch B with the shared RCF file. See [Install the RCF on a Cisco Nexus 9336C-FX2 shared switch](#) for further details.
14. Verify the storage attached to HA pair 1, storage switch B is healthy:

```
system health alert show -instance
```

Show example

```
storage::*> system health alert show -instance  
There are no entries matching your query.
```

15. Install the ISLs between shared switch A and shared switch B:

Show example

```
sh1# configure  
Enter configuration commands, one per line. End with CNTL/Z.  
sh1 (config)# interface e1/35-36  
sh1 (config-if-range)# no lldp transmit  
sh1 (config-if-range)# no lldp receive  
sh1 (config-if-range)# switchport mode trunk  
sh1 (config-if-range)# no spanning-tree bpduguard enable  
sh1 (config-if-range)# channel-group 101 mode active  
sh1 (config-if-range)# exit  
sh1 (config)# interface port-channel 101  
sh1 (config-if)# switchport mode trunk  
sh1 (config-if)# spanning-tree port type network  
sh1 (config-if)# exit  
sh1 (config)# exit
```

16. Convert HA pair 1 from a switchless cluster to a switched cluster. Use the cluster port assignments defined by the shared RCF. See [Install NX-OS software and Reference Configuration Files \(RCFs\)](#) for further details.
17. Verify that the switched networking configuration is valid:


```
network port show
```

Migrate from a switched cluster with switch-attached storage

You can migrate from a switched cluster with switch-attached storage by reusing the storage switches.

By reusing the storage switches the storage switches of HA pair 1 become the shared switches as shown in the following figure.

Switch Attached



Steps

1. Verify that the storage configuration of HA pair 1 (and HA pair 2) is correct and error free:

```
system switch ethernet show
```

Show example

```
storage::*> system switch ethernet show
Switch                Type                Address            Model
-----
sh1
                        storage-network    172.17.227.5     C9336C
    Serial Number: FOC221206C2
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        9.3(5)
    Version Source: CDP
sh2
                        storage-network    172.17.227.6     C9336C
    Serial Number: FOC220443LZ
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        9.3(5)
    Version Source: CDP
2 entries were displayed.
storage::*>
```

2. Move the HA pair 1, NSM224 path A cables from storage switch A to the NSM224 storage ports for HA pair 1, path A on storage switch A.
3. Move the cable from HA pair 1, node A, path A to the NSM224 storage port for HA pair 1, node A on storage switch A.
4. Move the cable from HA pair 1, node B, path A to the NSM224 storage port for HA pair 1, node B on storage switch A.
5. Verify the storage attached to HA pair 1, storage switch A is healthy:

```
storage port show -port-type ENET
```

Show example

```
storage::*> storage port show -port-type ENET
                               Speed
VLAN
Node   Port   Type   Mode   (Gb/s)   State   Status
ID
-----
node1
30     e0c    ENET   storage 100     enabled online
30     e0d    ENET   storage 100     enabled online
30     e5a    ENET   storage 100     enabled online
30     e5b    ENET   storage 100     enabled online
node2
30     e0c    ENET   storage 100     enabled online
30     e0d    ENET   storage 100     enabled online
30     e5a    ENET   storage 100     enabled online
30     e5b    ENET   storage 100     enabled online
```

6. Replace the storage RCF on shared switch A with the shared RCF file. See [Install the RCF on a Cisco Nexus 9336C-FX2 shared switch](#) for further details.

7. Verify the storage attached to HA pair 1, storage switch A is healthy:

```
system health alert show -instance
```

Show example

```
storage::*> system health alert show -instance
There are no entries matching your query.
```

8. Move the HA pair 1, NSM224 path B cables from storage switch B to the shared NS224 storage ports for HA pair 1, path B to storage switch B.

9. Move the cable from HA pair 1, node A, path B to the shared storage port for HA pair 1, node A, path B on storage switch B.
10. Move the cable from HA pair 1, node B, path B to the shared storage port for HA pair 1, node B, path B on storage switch B.
11. Verify the storage attached to HA pair 1, storage switch B is healthy:

```
system health alert show -instance
```

Show example

```
storage::*> system health alert show -instance  
There are no entries matching your query.
```

12. Replace the storage RCF file on shared switch B with the shared RCF file. See [Install the RCF on a Cisco Nexus 9336C-FX2 shared switch](#) for further details.
13. Verify the storage attached to HA pair 1, storage switch B is healthy:

```
system health alert show -instance
```

Show example

```
storage::*> system health alert show -instance  
There are no entries matching your query.
```

14. Verify the storage configuration of HA pair 1 is correct and error free:

```
system switch ethernet show
```


Show example

```
storage::*> system switch ethernet show
Switch                Type                Address
Model
-----
sh1
                    storage-network    172.17.227.5
C9336C
    Serial Number: FOC221206C2
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                    9.3(5)
    Version Source: CDP
sh2
                    storage-network    172.17.227.6
C9336C
    Serial Number: FOC220443LZ
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                    9.3(5)
    Version Source: CDP
2 entries were displayed.
storage::*>
```

15. Install the ISLs between shared switch A and shared switch B:

Show example

```
sh1# configure
Enter configuration commands, one per line. End with CNTL/Z.
sh1 (config)# interface e1/35-36*
sh1 (config-if-range)# no lldp transmit
sh1 (config-if-range)# no lldp receive
sh1 (config-if-range)# switchport mode trunk
sh1 (config-if-range)# no spanning-tree bpduguard enable
sh1 (config-if-range)# channel-group 101 mode active
sh1 (config-if-range)# exit
sh1 (config)# interface port-channel 101
sh1 (config-if)# switchport mode trunk
sh1 (config-if)# spanning-tree port type network
sh1 (config-if)# exit
sh1 (config)# exit
```

16. Migrate the cluster networking from the existing cluster switches to the shared switches using the switch replacement procedure and the shared RCF. The new shared switch A is "cs1". The new shared switch B is "cs2". See [Replace a Cisco Nexus 9336C-FX2 shared switch](#) and [Install the RCF on a Cisco Nexus 9336C-FX2 shared switch](#) for further details.
17. Verify that the switched networking config is valid:

```
network port show
```

18. Remove the unused cluster switches.
19. Remove the unused storage switches.

Replace a Cisco Nexus 9336C-FX2 shared switch

You can replace a defective Nexus 9336C-FX2 shared switch. This is a nondisruptive procedure (NDU).

What you'll need

Before performing the switch replacement, make sure that:

- In the existing cluster and network infrastructure:
 - The existing cluster is verified as completely functional, with at least one fully connected cluster switch.
 - All cluster ports are **up**.
 - All cluster logical interfaces (LIFs) are **up** and on their home ports.
 - The ONTAP cluster ping-cluster -node node1 command must indicate that basic connectivity and larger than PMTU communication are successful on all paths.
- For the Nexus 9336C-FX2 replacement switch:
 - Management network connectivity on the replacement switch is functional.

- Console access to the replacement switch is in place.
- The node connections are ports 1/1 through 1/34:
- All Inter-Switch Link (ISL) ports are disabled on ports 1/35 and 1/36.
- The desired reference configuration file (RCF) and NX-OS operating system image switch is loaded onto the switch.
- Any previous site customizations, such as STP, SNMP, and SSH, should be copied to the new switch.

About the examples

You must execute the command for migrating a cluster LIF from the node where the cluster LIF is hosted.

The examples in this procedure use the following switch and node nomenclature:

- The names of the existing Nexus 9336C-FX2 switches are *sh1* and *sh2*.
- The name of the new Nexus 9336C-FX2 switches are *newsh1* and *newsh2*.
- The node names are *node1* and *node2*.
- The cluster ports on each node are named *e3a* and *e3b*.
- The cluster LIF names are *node1_clus1* and *node1_clus2* for *node1*, and *node2_clus1* and *node2_clus2* for *node2*.
- The prompt for changes to all cluster nodes is `cluster1::*>`.



The following procedure is based on the following network topology:

Show example topology

```
cluster1::*> network port show -ipSpace Cluster
```

```
Node: node1
```

```
Ignore
```

						Speed(Mbps)	Health
Health	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper	Status
Status							
	e3a	Cluster	Cluster	up	9000	auto/100000	healthy
false							
	e3b	Cluster	Cluster	up	9000	auto/100000	healthy
false							

```
Node: node2
```

```
Ignore
```

						Speed(Mbps)	Health
Health	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper	Status
Status							
	e3a	Cluster	Cluster	up	9000	auto/100000	healthy
false							
	e3b	Cluster	Cluster	up	9000	auto/100000	healthy
false							

```
4 entries were displayed.
```

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current		
Current Is	Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home						
	Cluster	node1_clus1	up/up	169.254.209.69/16	node1	e3a
true						
		node1_clus2	up/up	169.254.49.125/16	node1	e3b
true						

```

node2_clus1 up/up 169.254.47.194/16 node2 e3a
true
node2_clus2 up/up 169.254.19.183/16 node2 e3b
true
4 entries were displayed.

```

cluster1::*> **network device-discovery show -protocol cdp**

Node/	Local	Discovered		
Protocol	Port	Device (LLDP: ChassisID)	Interface	Platform
node2	/cdp			
	e3a	sh1	Eth1/2	N9K-
C9336C				
	e3b	sh2	Eth1/2	N9K-
C9336C				
node1	/cdp			
	e3a	sh1	Eth1/1	N9K-
C9336C				
	e3b	sh2	Eth1/1	N9K-
C9336C				

4 entries were displayed.

sh1# **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 Infrfce	Hldtme	Capability	Platform	Port
node1	Eth1/1	144	H	FAS2980	e3a
node2	Eth1/2	145	H	FAS2980	e3a
sh2	Eth1/35	176	R S I s	N9K-C9336C	
Eth1/35					
sh2 (FDO220329V5)	Eth1/36	176	R S I s	N9K-C9336C	
Eth1/36					

Total entries displayed: 4

sh2# **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 Infrfce	Hldtme	Capability	Platform	Port
ID					

```

node1          Eth1/1          139    H          FAS2980      eb
node2          Eth1/2          124    H          FAS2980      eb
sh1           Eth1/35         178    R S I s    N9K-C9336C
Eth1/35
sh1           Eth1/36         178    R S I s    N9K-C9336C
Eth1/36
Total entries displayed: 4

```

Steps

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=xh
```

Where x is the duration of the maintenance window in hours.

2. Optional: Install the appropriate RCF and image on the switch, newsh2, and make any necessary site preparations.
 - a. If necessary, verify, download, and install the appropriate versions of the RCF and NX-OS software for the new switch. If you have verified that the new switch is correctly set up and does not need updates to the RCF and NX-OS software, continue to [Step 3](#).
 - b. Go to the [NetApp Cluster and Management Network Switches Reference Configuration File Description Page](#) on the NetApp Support Site.
 - c. Click the link for the [Cluster Network and Management Network Compatibility Matrix](#), and then note the required switch software version.
 - d. Click your browser's back arrow to return to the Description page, click CONTINUE, accept the license agreement, and then go to the Download page.
 - e. Follow the steps on the Download page to download the correct RCF and NX-OS files for the version of ONTAP software you are installing.
3. On the new switch, log in as admin and shut down all the ports that will be connected to the node cluster interfaces (ports 1/1 to 1/34).
If the switch that you are replacing is not functional and is powered down, go to [Step 4](#). The LIFs on the cluster nodes should have already failed over to the other cluster port for each node.

Show example

```

newsh2# config
Enter configuration commands, one per line. End with CNTL/Z.
newsh2(config)# interface e1/1-34
newsh2(config-if-range)# shutdown

```

4. Verify that all cluster LIFs have auto-revert enabled.

```
network interface show - vserver Cluster -fields auto-revert
```

Show example

```
cluster1::> network interface show -vserver Cluster -fields auto-revert
```

	Logical	
Vserver	Interface	Auto-revert
-----	-----	-----
Cluster	node1_clus1	true
Cluster	node1_clus2	true
Cluster	node2_clus1	true
Cluster	node2_clus2	true

4 entries were displayed.

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```



```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e3a
Cluster node1_clus2 169.254.49.125 node1 e3b
Cluster node2_clus1 169.254.47.194 node2 e3a
Cluster node2_clus2 169.254.19.183 node2 e3b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

6. Shut down the ISL ports 1/35 and 1/36 on the Nexus 9336C-FX2 switch sh1.

Show example

```

sh1# configure
Enter configuration commands, one per line. End with CNTL/Z.
sh1(config)# interface e1/35-36
sh1(config-if-range)# shutdown

```

7. Remove all the cables from the Nexus 9336C-FX2 sh2 switch, and then connect them to the same ports on the Nexus C9336C-FX2 newsh2 switch.
8. Bring up the ISLs ports 1/35 and 1/36 between the sh1 and newsh2 switches, and then verify the port channel operation status.

Port-Channel should indicate Po1(SU) and Member Ports should indicate Eth1/35(P) and Eth1/36(P).

Show example

This example enables ISL ports 1/35 and 1/36 and displays the port channel summary on switch sh1.

```
sh1# configure
Enter configuration commands, one per line. End with CNTL/Z.
sh1 (config)# int e1/35-36
sh1 (config-if-range)# no shutdown
sh1 (config-if-range)# 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      Eth1/35 (P)  Eth1/36 (P)

sh1 (config-if-range)#
```

9. Verify that port e3b is up on all nodes:

```
network port show ipspace Cluster
```

Show example

The output should be like the following:

```
cluster1::*> network port show -ipSpace Cluster

Node: node1

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper  Speed (Mbps)
Status      Status
-----
-----
e3a         Cluster   Cluster           up   9000  auto/100000
healthy    false
e3b         Cluster   Cluster           up   9000  auto/100000
healthy    false

Node: node2

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper  Speed (Mbps)
Status      Status
-----
-----
e3a         Cluster   Cluster           up   9000  auto/100000
healthy    false
e3b         Cluster   Cluster           up   9000  auto/auto   -
false
4 entries were displayed.
```

10. On the same node you used in the previous step, revert the cluster LIF associated with the port in the previous step by using the network interface revert command.

In this example, LIF node1_clus2 on node1 is successfully reverted if the Home value is true and the port is e3b.

The following commands return LIF node1_clus2 on node1 to home port e3a and displays information about the LIFs on both nodes. Bringing up the first node is successful if the Is Home column is **true** for both cluster interfaces and they show the correct port assignments, in this example e3a and e3b on node1.

Show example

```
cluster1::*> network interface show -vserver Cluster
```

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
e3a	node1_clus1	up/up	169.254.209.69/16	node1
	true			
e3b	node1_clus2	up/up	169.254.49.125/16	node1
	true			
e3a	node2_clus1	up/up	169.254.47.194/16	node2
	true			
e3a	node2_clus2	up/up	169.254.19.183/16	node2
	false			

4 entries were displayed.

11. Display information about the nodes in a cluster:

```
cluster show
```

Show example

This example shows that the node health for node1 and node2 in this cluster is true:

```
cluster1::*> cluster show
```

Node	Health	Eligibility

node1	false	true
node2	true	true

12. Verify that all physical cluster ports are up:

```
network port show ipspace Cluster
```

Show example

```
cluster1::*> network port show -ipSpace Cluster

Node node1
Ignore
Speed (Mbps)
Health Health
Port IPspace Broadcast Domain Link MTU Admin/Oper
Status Status
-----
-----
e3a Cluster Cluster up 9000 auto/100000
healthy false
e3b Cluster Cluster up 9000 auto/100000
healthy false

Node: node2

Ignore
Speed (Mbps)
Health Health
Port IPspace Broadcast Domain Link MTU Admin/Oper
Status Status
-----
-----
e3a Cluster Cluster up 9000 auto/100000
healthy false
e3b Cluster Cluster up 9000 auto/100000
healthy false
4 entries were displayed.
```

13. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```
cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e3a
Cluster node1_clus2 169.254.49.125 node1 e3b
Cluster node2_clus1 169.254.47.194 node2 e3a
Cluster node2_clus2 169.254.19.183 node2 e3b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

14. Confirm the following cluster network configuration:

```
network port show
```

Show example

```
cluster1::*> network port show -ipspace Cluster
```

```
Node: node1
```

```
Ignore
```

```
Speed (Mbps)
```

```
Health Health
```

```
Port IPspace Broadcast Domain Link MTU Admin/Oper
```

```
Status Status
```

```
-----  
-----  
e3a Cluster Cluster up 9000 auto/100000  
healthy false  
e3b Cluster Cluster up 9000 auto/100000  
healthy false
```

```
Node: node2
```

```
Ignore
```

```
Speed (Mbps)
```

```
Health Health
```

```
Port IPspace Broadcast Domain Link MTU Admin/Oper
```

```
Status Status
```

```
-----  
-----  
e3a Cluster Cluster up 9000 auto/100000  
healthy false  
e3b Cluster Cluster up 9000 auto/100000  
healthy false
```

```
4 entries were displayed.
```

```
cluster1::*> network interface show -vserver Cluster
```

```
Logical Status Network Current
```

```
Current Is
```

```
Vserver Interface Admin/Oper Address/Mask Node
```

```
Port Home
```

```
-----  
-----  
Cluster  
node1_clus1 up/up 169.254.209.69/16 node1  
e3a true  
node1_clus2 up/up 169.254.49.125/16 node1  
e3b true  
node2_clus1 up/up 169.254.47.194/16 node2
```



```
e3a      true
          node2_clus2  up/up      169.254.19.183/16  node2
```

```
e3b      true
```

4 entries were displayed.

```
cluster1::> network device-discovery show -protocol cdp
```

```
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface
Platform

-----
node2      /cdp
          e3a   sh1    0/2                N9K-C9336C
          e3b   newsh2                0/2                N9K-
C9336C
node1      /cdp
          e3a   sh1                0/1                N9K-
C9336C
          e3b   newsh2                0/1                N9K-
C9336C
```

4 entries were displayed.

```
sh1# 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
node1          Eth1/1        144     H           FAS2980
e3a
node2          Eth1/2        145     H           FAS2980
e3a
newsh2         Eth1/35       176     R S I s     N9K-C9336C
Eth1/35
newsh2         Eth1/36       176     R S I s     N9K-C9336C
Eth1/36
```

Total entries displayed: 4

```
sh2# 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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e3b	Eth1/1	139	H	FAS2980
node2 eb	Eth1/2	124	H	FAS2980
sh1 Eth1/35	Eth1/35	178	R S I s	N9K-C9336C
sh1 Eth1/36	Eth1/36	178	R S I s	N9K-C9336C

Total entries displayed: 4

15. Move the storage ports from the old switch sh2 to the new switch newsh2.
16. Verify the storage attached to HA pair 1, shared switch newsh2 is healthy.
17. Verify the storage attached to HA pair 2, shared switch newsh2 is healthy:

```
storage port show -port-type ENET
```

Show example

```
storage::*> storage port show -port-type ENET
```

```
Speed
VLAN
Node   Port   Type   Mode   Speed   State   Status
ID      (Gb/s)
-----
node1
30     e3a    ENET   storage 100     enabled online
30     e3b    ENET   storage 0        enabled offline
30     e7a    ENET   storage 0        enabled offline
30     e7b    ENET   storage 100     enabled online
node2
30     e3a    ENET   storage 100     enabled online
30     e3b    ENET   storage 0        enabled offline
30     e7a    ENET   storage 0        enabled offline
30     e7b    ENET   storage 100     enabled online
```

18. Verify that the shelves are correctly cabled:

```
storage shelf port show -fields remote- device,remote-port
```

Show example

```
cluster1::*> storage shelf port show -fields remote-device,remote-  
port  
shelf id remote-port  remote-device  
-----  
3.20  0  Ethernet1/13  sh1  
3.20  1  Ethernet1/13  newsh2  
3.20  2  Ethernet1/14  sh1  
3.20  3  Ethernet1/14  newsh2  
3.30  0  Ethernet1/15  sh1  
3.30  1  Ethernet1/15  newsh2  
3.30  2  Ethernet1/16  sh1  
3.30  3  Ethernet1/16  newsh2  
8 entries were displayed.
```

19. Remove the old switch sh2.
20. Repeat these steps for the switch sh1 and new switch newsh1.
21. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next?

[Configure switch health monitoring.](#)

Monitor switch health

Overview of switch health monitor

The Ethernet switch health monitor (CSHM) is responsible for ensuring the operational health of Cluster and Storage network switches and collecting switch logs for debugging purposes.

Configure switch health monitoring

Configuration overview

The Ethernet switch health monitor (CSHM) is responsible for ensuring the operational health of Cluster and Storage network switches and collecting switch logs for debugging purposes.

- [Configure log collection](#)
- [Optional: Configure SNMPv3](#)

Configure log collection

The Ethernet switch health monitor (CSHM) is responsible for ensuring the operational health of Cluster and Storage network switches and collecting switch logs for debugging purposes. This procedure guides you through the process of setting up collection, requesting detailed **Support** logs, and enabling an hourly collection of **Periodic** data that is collected by AutoSupport.

NOTE: If you enable FIPS mode, you must complete the following:



1. Regenerate SSH keys on the switch using the vendor instructions.
2. Regenerate SSH keys in ONTAP using `debug system regenerate-systemshell-key-pair`
3. Re-run log collection setup routine using the `system switch ethernet log setup-password` command

Before you begin

- The user must have access to the switch `show` commands. If these are not available, create a new user and grant the necessary permissions to the user.
- Switch health monitoring must be enabled for the switch. Verify this by ensuring the `Is Monitored:` field is set to **true** in the output of the `system switch ethernet show` command.
- For log collection with Broadcom and Cisco switches:
 - Users must have network admin privileges.
 - A new user should be created on the switch for each cluster setup with log collection enabled. These switches do not support multiple SSH keys for the same user. Any additional log collection setup

performed overwrites any pre-existing SSH keys for the user.

- For support log collection with NVIDIA switches, the *user* for log collection must be permitted to run the `cl-support` command without having to provide a password. To allow this usage, run the command:

```
echo '<user> ALL = NOPASSWD: /usr/cumulus/bin/cl-support' | sudo EDITOR='tee  
-a' visudo -f /etc/sudoers.d/cumulus
```

Steps

ONTAP 9.14.1 and earlier

1. To set up log collection, run the following command for each switch. You are prompted to enter the switch name, username, and password for log collection.

NOTE: If answering `y` to the user specification prompt, make sure that the user has the necessary permissions as outlined in [Before you begin](#) .

```
system switch ethernet log setup-password
```

```
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
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
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
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
Enter the password: <enter switch password>
```

```
Enter the password again: <enter switch password>
```

2. To request support log collection and enable periodic collection, run the following command. This starts both types of log collection: the detailed Support logs and an hourly collection of Periodic data.

```
system switch ethernet log modify -device <switch-name> -log-request  
true
```

```
cluster1::*> system switch ethernet log modify -device cs1 -log  
-request true
```

```
Do you want to modify the cluster switch log collection  
configuration? {y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

```
cluster1::*> system switch ethernet log modify -device cs2 -log  
-request true
```

```
Do you want to modify the cluster switch log collection  
configuration? {y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

Wait for 10 minutes and then check that the log collection completes:

```
system switch ethernet log show
```

ONTAP 9.15.1 and later

1. To set up log collection, run the following command for each switch. You are prompted to enter the switch name, username, and password for log collection.

NOTE: If answering `y` to the user specification prompt, make sure that the user has the necessary permissions as outlined in [Before you begin](#).

```
system switch ethernet log setup-password
```



```
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
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
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
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
Enter the password: <enter switch password>
```

```
Enter the password again: <enter switch password>
```

2. Enable periodic log collection:

```
system switch ethernet log modify -device <switch-name> -periodic  
-enabled true
```

```
cluster1::*> system switch ethernet log modify -device cs1 -periodic
-enabled true
```

Do you want to modify the cluster switch log collection configuration? {y|n}: [n] **y**

cs1: Periodic log collection has been scheduled to run every hour.

```
cluster1::*> system switch ethernet log modify -device cs2 -periodic
-enabled true
```

Do you want to modify the cluster switch log collection configuration? {y|n}: [n] **y**

cs2: Periodic log collection has been scheduled to run every hour.

```
cluster1::*> system switch ethernet log show
```

	Periodic	Periodic
Support		
Switch	Log Enabled	Log State
Log State		
cs1	true	scheduled
never-run		
cs2	true	scheduled
never-run		

2 entries were displayed.

3. Request support log collection:

```
system switch ethernet log collect-support-log -device <switch-name>
```

```
cluster1::*> system switch ethernet log collect-support-log -device cs1
```

```
cs1: Waiting for the next Ethernet switch polling cycle to begin support collection.
```

```
cluster1::*> system switch ethernet log collect-support-log -device cs2
```

```
cs2: Waiting for the next Ethernet switch polling cycle to begin support collection.
```

```
cluster1::*> *system switch ethernet log show
```

	Periodic	Periodic
Support		
Switch	Log Enabled	Log State
Log State		
cs1	false	halted
initiated		
cs2	true	scheduled
initiated		

2 entries were displayed.

4. To view all details of log collection, including the enablement, status message, previous timestamp and filename of periodic collection, the request status, status message, and previous timestamp and filename of support collection, use the following:

```
system switch ethernet log show -instance
```

```
cluster1::*> system switch ethernet log show -instance

                Switch Name: cs1
    Periodic Log Enabled: true
        Periodic Log Status: Periodic log collection has been
scheduled to run every hour.
    Last Periodic Log Timestamp: 3/11/2024 11:02:59
        Periodic Log Filename: cluster1:/mroot/etc/log/shm-
cluster-info.tgz
    Support Log Requested: false
        Support Log Status: Successfully gathered support logs
- see filename for their location.
    Last Support Log Timestamp: 3/11/2024 11:14:20
        Support Log Filename: cluster1:/mroot/etc/log/shm-
cluster-log.tgz

                Switch Name: cs2
    Periodic Log Enabled: false
        Periodic Log Status: Periodic collection has been
halted.
    Last Periodic Log Timestamp: 3/11/2024 11:05:18
        Periodic Log Filename: cluster1:/mroot/etc/log/shm-
cluster-info.tgz
    Support Log Requested: false
        Support Log Status: Successfully gathered support logs
- see filename for their location.
    Last Support Log Timestamp: 3/11/2024 11:18:54
        Support Log Filename: cluster1:/mroot/etc/log/shm-
cluster-log.tgz
2 entries were displayed.
```



If any error statuses are reported by the log collection feature (visible in the output of `system switch ethernet log show`), see [Troubleshoot log collection](#) for further details.

What's next?

[Configure SNMPv3 \(Optional\)](#).

Optional: Configure SNMPv3 for your switch

SNMP is used to monitor the switches. The Ethernet Switch Health Monitor (CSHM) utilizes SNMP to monitor the health and performance of cluster and storage switches. By default, SNMPv2c is configured automatically through the Reference Configuration File (RCF).

SNMPv3 is more secure than SNMPv2 because it introduces robust security features such as authentication, encryption, and message integrity, which protect against unauthorized access and ensure data confidentiality and integrity during transmission.



SNMPv3 is only supported on ONTAP 9.12.1 and later.

Follow this procedure to configure SNMPv3 for your specific switch, which supports CSHM.

About this task

The following commands are used to configure an SNMPv3 username on **Broadcom**, **Cisco**, and **NVIDIA** switches:

Broadcom switches

Configure an SNMPv3 username NETWORK-OPERATOR on Broadcom BES-53248 switches.

- For **no authentication**:

```
snmp-server user SNMPv3UserNoAuth NETWORK-OPERATOR noauth
```

- For **MD5/SHA authentication**:

```
snmp-server user SNMPv3UserAuth NETWORK-OPERATOR [auth-md5|auth-sha]
```

- For **MD5/SHA authentication with AES/DES encryption**:

```
snmp-server user SNMPv3UserAuthEncrypt NETWORK-OPERATOR [auth-  
md5|auth-sha] [priv-aes128|priv-des]
```

The following command configures an SNMPv3 username on the ONTAP side:

```
security login create -user-or-group-name SNMPv3_USER -application snmp  
-authentication-method usm -remote-switch-ipaddress ADDRESS
```

The following command establishes the SNMPv3 username with CSHM:

```
cluster1::*> system switch ethernet modify -device DEVICE -snmp-version  
SNMPv3 -community-or-username SNMPv3_USER
```

Steps

1. Set up the SNMPv3 user on the switch to use authentication and encryption:

```
show snmp status
```

```
(sw1) (Config)# snmp-server user <username> network-admin auth-md5
<password> priv-aes128 <password>
```

```
(cs1) (Config)# show snmp user snmp
```

Name	Group Name	Auth Meth	Priv Meth	Remote Engine ID
<username>	network-admin	MD5	AES128	8000113d03d8c497710bee

2. Set up the SNMPv3 user on the ONTAP side:

```
security login create -user-or-group-name <username> -application
snmp -authentication-method usm -remote-switch-ipaddress
10.231.80.212
```

```
cluster1::*> security login create -user-or-group-name <username>
-application snmp -authentication-method usm -remote-switch
-ipaddress 10.231.80.212
```

Enter the authoritative entity's EngineID [remote EngineID]:

Which authentication protocol do you want to choose (none, md5, sha, sha2-256)

[none]: **md5**

Enter the authentication protocol password (minimum 8 characters long):

Enter the authentication protocol password again:

Which privacy protocol do you want to choose (none, des, aes128)

[none]: **aes128**

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:

3. Configure CSHM to monitor with the new SNMPv3 user:

```
system switch ethernet show-all -device "sw1" -instance
```

```

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance

Device Name: sw1
IP Address: 10.228.136.24
SNMP Version: SNMPv2c
Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
Community String or SNMPv3 Username: cshml!
Model Number: BES-53248
Switch Network: cluster-network
Software Version: 3.9.0.2
Reason For Not Monitoring: None <---- should
display this if SNMP settings are valid
Source Of Switch Version: CDP/ISDP
Is Monitored ?: true
Serial Number of the Device: QTFCU3826001C
RCF Version: v1.8X2 for

Cluster/HA/RDMA

cluster1::*>
cluster1::*> system switch ethernet modify -device "sw1" -snmp
-version SNMPv3 -community-or-username <username>

```

4. Verify that the serial number to be queried with the newly created SNMPv3 user is the same as detailed in the previous step after the CSHM polling period has completed.

```
system switch ethernet polling-interval show
```



```

cluster1::*> system switch ethernet polling-interval show
Polling Interval (in minutes): 5

cluster1::*> system switch ethernet show-all -device "sw1" -instance
Device Name: sw1
IP Address: 10.228.136.24
SNMP Version: SNMPv3
Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
Community String or SNMPv3 Username: <username>
Model Number: BES-53248
Switch Network: cluster-network
Software Version: 3.9.0.2
Reason For Not Monitoring: None <---- should
display this if SNMP settings are valid
Source Of Switch Version: CDP/ISDP
Is Monitored?: true
Serial Number of the Device: QTFUCU3826001C
RCF Version: v1.8X2 for

Cluster/HA/RDMA

```

Cisco switches

Configure an SNMPv3 username SNMPv3_USER on Cisco 9336C-FX2 switches:

- For no authentication:

```
snmp-server user SNMPv3_USER NoAuth
```

- For MD5/SHA authentication:

```
snmp-server user SNMPv3_USER auth [md5|sha] AUTH-PASSWORD
```

- For MD5/SHA authentication with AES/DES encryption:

```
snmp-server user SNMPv3_USER AuthEncrypt auth [md5|sha] AUTH-
PASSWORD priv aes-128 PRIV-PASSWORD
```

The following command configures an SNMPv3 username on the ONTAP side:

```
security login create -user-or-group-name SNMPv3_USER -application snmp
-authentication-method usm -remote-switch-ipaddress ADDRESS
```

The following command establishes the SNMPv3 username with CSHM:

```
system switch ethernet modify -device DEVICE -snmp-version SNMPv3
-community-or-username SNMPv3_USER
```

Steps

1. Set up the SNMPv3 user on the switch to use authentication and encryption:

```
show snmp user
```

```
(sw1) (Config)# snmp-server user SNMPv3User auth md5 <auth_password>
priv aes-128 <priv_password>
```

```
(sw1) (Config)# show snmp user
```

```
-----
-----
```

SNMP USERS

```
-----
-----
```

User	Auth	Priv(enforce)	Groups
acl_filter			
admin	md5	des(no)	network-admin
SNMPv3User	md5	aes-128(no)	network-operator

```
-----
-----
```

NOTIFICATION TARGET USERS (configured for sending V3 Inform)

```
-----
-----
```

User	Auth	Priv
------	------	------

```
-----
-----
```

```
(sw1) (Config)#
```

2. Set up the SNMPv3 user on the ONTAP side:

```
security login create -user-or-group-name <username> -application  
snmp -authentication-method usm -remote-switch-ipaddress  
10.231.80.212
```

```
cluster1::*> system switch ethernet modify -device "sw1  
(b8:59:9f:09:7c:22)" -is-monitoring-enabled-admin true
```

```
cluster1::*> security login create -user-or-group-name <username>  
-application snmp -authentication-method usm -remote-switch  
-ipaddress 10.231.80.212
```

Enter the authoritative entity's EngineID [remote EngineID]:

Which authentication protocol do you want to choose (none, md5, sha,
sha2-256)

[none]: **md5**

Enter the authentication protocol password (minimum 8 characters
long):

Enter the authentication protocol password again:

Which privacy protocol do you want to choose (none, des, aes128)

[none]: **aes128**

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:

3. Configure CSHM to monitor with the new SNMPv3 user:

```
system switch ethernet show-all -device "sw1" -instance
```

```

cluster1::*> system switch ethernet show-all -device "sw1" -instance

                Device Name: sw1
                IP Address: 10.231.80.212
                SNMP Version: SNMPv2c
                Is Discovered: true
                SNMPv2c Community String or SNMPv3 Username: cshml!
                Model Number: N9K-C9336C-FX2
                Switch Network: cluster-network
                Software Version: Cisco Nexus
Operating System (NX-OS) Software, Version 9.3(7)
                Reason For Not Monitoring: None <---- displays
when SNMP settings are valid
                Source Of Switch Version: CDP/ISDP
                Is Monitored?: true
                Serial Number of the Device: QTFCU3826001C
                RCF Version: v1.8X2 for

Cluster/HA/RDMA

cluster1::*>
cluster1::*> system switch ethernet modify -device "sw1" -snmp
-version SNMPv3 -community-or-username <username>
cluster1::*>

```

4. Verify that the serial number to be queried with the newly created SNMPv3 user is the same as detailed in the previous step after the CSHM polling period has completed.

```

system switch ethernet polling-interval show

```

```

cluster1::*> system switch ethernet polling-interval show
                Polling Interval (in minutes): 5

cluster1::*> system switch ethernet show-all -device "sw1" -instance

                Device Name: sw1
                IP Address: 10.231.80.212
                SNMP Version: SNMPv3
                Is Discovered: true
                SNMPv2c Community String or SNMPv3 Username: SNMPv3User
                Model Number: N9K-C9336C-FX2
                Switch Network: cluster-network
                Software Version: Cisco Nexus
                Operating System (NX-OS) Software, Version 9.3(7)
                Reason For Not Monitoring: None <---- displays
when SNMP settings are valid
                Source Of Switch Version: CDP/ISDP
                Is Monitored ?: true
                Serial Number of the Device: QTFCU3826001C
                RCF Version: v1.8X2 for
                Cluster/HA/RDMA

cluster1::*>

```

NVIDIA - CLI 5.4

Configure an SNMPv3 username SNMPv3_USER on NVIDIA SN2100 switches running CLI 5.4:

- For **no authentication**:

```
nv set service snmp-server username SNMPv3_USER auth-none
```

- For **MD5/SHA authentication**:

```
nv set service snmp-server username SNMPv3_USER [auth-md5|auth-sha]
AUTH-PASSWORD
```

- For **MD5/SHA authentication with AES/DES encryption**:

```
nv set service snmp-server username SNMPv3_USER [auth-md5|auth-sha]
AUTH-PASSWORD [encrypt-aes|encrypt-des] PRIV-PASSWORD
```

The following command configures an SNMPv3 username on the ONTAP side:

```
security login create -user-or-group-name SNMPv3_USER -application snmp
-authentication-method usm -remote-switch-ipaddress ADDRESS
```

The following command establishes the SNMPv3 username with CSHM:

```
system switch ethernet modify -device DEVICE -snmp-version SNMPv3
-community-or-username SNMPv3_USER
```

Steps

1. Set up the SNMPv3 user on the switch to use authentication and encryption:

```
net show snmp status
```

```
cumulus@sw1:~$ net show snmp status
Simple Network Management Protocol (SNMP) Daemon.
-----
Current Status          active (running)
Reload Status           enabled
Listening IP Addresses  all vrf mgmt
Main snmpd PID          4318
Version 1 and 2c Community String  Configured
Version 3 Usernames     Not Configured
-----

cumulus@sw1:~$
cumulus@sw1:~$ net add snmp-server username SNMPv3User auth-md5
<password> encrypt-aes <password>
cumulus@sw1:~$ net commit
--- /etc/snmp/snmpd.conf      2020-08-02 21:09:34.686949282 +0000
+++ /run/nclu/snmp/snmpd.conf 2020-08-11 00:13:51.826126655 +0000
@@ -1,26 +1,28 @@
# Auto-generated config file: do not edit. #
agentaddress udp:@mgmt:161
agentxperms 777 777 snmp snmp
agentxsocket /var/agentx/master
createuser _snmptrapusernameX
+createuser SNMPv3User MD5 <password> AES <password>
ifmib_max_num_ifaces 500
iquerysecname _snmptrapusernameX
master agentx
monitor -r 60 -o laNames -o laErrorMessage "laTable" laErrorFlag != 0
pass -p 10 1.3.6.1.2.1.1.1 /usr/share/snmp/sysDescr_pass.py
pass_persist 1.2.840.10006.300.43
```

```

/usr/share/snmp/ieee8023_lag_pp.py
  pass_persist 1.3.6.1.2.1.17 /usr/share/snmp/bridge_pp.py
  pass_persist 1.3.6.1.2.1.31.1.1.1.18
/usr/share/snmp/snmpifAlias_pp.py
  pass_persist 1.3.6.1.2.1.47 /usr/share/snmp/entity_pp.py
  pass_persist 1.3.6.1.2.1.99 /usr/share/snmp/entity_sensor_pp.py
  pass_persist 1.3.6.1.4.1.40310.1 /usr/share/snmp/resq_pp.py
  pass_persist 1.3.6.1.4.1.40310.2
/usr/share/snmp/cl_drop_cntrs_pp.py
  pass_persist 1.3.6.1.4.1.40310.3 /usr/share/snmp/cl_poe_pp.py
  pass_persist 1.3.6.1.4.1.40310.4 /usr/share/snmp/bgpun_pp.py
  pass_persist 1.3.6.1.4.1.40310.5 /usr/share/snmp/cumulus-status.py
  pass_persist 1.3.6.1.4.1.40310.6 /usr/share/snmp/cumulus-sensor.py
  pass_persist 1.3.6.1.4.1.40310.7 /usr/share/snmp/vrf_bgpun_pp.py
+rocommunity cshml! default
  rouser _snmptrapusernameX
+rouser SNMPv3User priv
  sysobjectid 1.3.6.1.4.1.40310
  sysservices 72
-rocommunity cshml! default

```

net add/del commands since the last "net commit"

User	Timestamp	Command
SNMPv3User	2020-08-11 00:13:51.826987	net add snmp-server username SNMPv3User auth-md5 <password> encrypt-aes <password>

```

cumulus@sw1:~$
cumulus@sw1:~$ net show snmp status
Simple Network Management Protocol (SNMP) Daemon.
-----
Current Status          active (running)
Reload Status           enabled
Listening IP Addresses  all vrf mgmt
Main snmpd PID          24253
Version 1 and 2c Community String Configured
Version 3 Usernames     Configured    <---- Configured
here
-----
cumulus@sw1:~$

```

2. Set up the SNMPv3 user on the ONTAP side:

```
security login create -user-or-group-name SNMPv3User -application  
snmp -authentication-method usm -remote-switch-ipaddress  
10.231.80.212
```

```
cluster1::*> security login create -user-or-group-name SNMPv3User  
-application snmp -authentication-method usm -remote-switch  
-ipaddress 10.231.80.212
```

Enter the authoritative entity's EngineID [remote EngineID]:

Which authentication protocol do you want to choose (none, md5, sha,
sha2-256)

[none]: **md5**

Enter the authentication protocol password (minimum 8 characters
long):

Enter the authentication protocol password again:

Which privacy protocol do you want to choose (none, des, aes128)

[none]: **aes128**

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:

3. Configure CSHM to monitor with the new SNMPv3 user:

```
system switch ethernet show-all -device "sw1 (b8:59:9f:09:7c:22)"  
-instance
```



```

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance
                                     Device Name: sw1
(b8:59:9f:09:7c:22)
                                     IP Address: 10.231.80.212
                                     SNMP Version: SNMPv2c
                                     Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
      Community String or SNMPv3 Username: cshml!
      Model Number: MSN2100-CB2FC
      Switch Network: cluster-network
      Software Version: Cumulus Linux
version 5.4.0 running on Mellanox Technologies Ltd. MSN2100
      Reason For Not Monitoring: None
      Source Of Switch Version: LLDP
      Is Monitored ?: true
      Serial Number of the Device: MT2110X06399 <----
serial number to check
      RCF Version: MSN2100-RCF-v1.9X6-
Cluster-LLDP Aug-18-2022

cluster1::*>
cluster1::*> system switch ethernet modify -device "sw1
(b8:59:9f:09:7c:22)" -snmp-version SNMPv3 -community-or-username
SNMPv3User

```

4. Verify that the serial number to be queried with the newly created SNMPv3 user is the same as detailed in the previous step after the CSHM polling period has completed.

```
system switch ethernet polling-interval show
```

```

cluster1::*> system switch ethernet polling-interval show
      Polling Interval (in minutes): 5

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance
      Device Name: sw1
      IP Address: 10.231.80.212
      SNMP Version: SNMPv3
      Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
      Community String or SNMPv3 Username: SNMPv3User
      Model Number: MSN2100-CB2FC
      Switch Network: cluster-network
      Software Version: Cumulus Linux
version 5.4.0 running on Mellanox Technologies Ltd. MSN2100
      Reason For Not Monitoring: None
      Source Of Switch Version: LLDP
      Is Monitored?: true
      Serial Number of the Device: MT2110X06399 <----
serial number to check
      RCF Version: MSN2100-RCF-v1.9X6-
Cluster-LLDP Aug-18-2022

```

NVIDIA - CLI 5.11

Configure an SNMPv3 username SNMPv3_USER on NVIDIA SN2100 switches running CLI 5.11:

- For no authentication:

```
nv set system snmp-server username SNMPv3_USER auth-none
```

- For MD5/SHA authentication:

```
nv set system snmp-server username SNMPv3_USER [auth-md5|auth-sha]
AUTH-PASSWORD
```

- For MD5/SHA authentication with AES/DES encryption:

```
nv set system snmp-server username SNMPv3_USER [auth-md5|auth-sha]
AUTH-PASSWORD [encrypt-aes|encrypt-des] PRIV-PASSWORD
```

The following command configures an SNMPv3 username on the ONTAP side:

```
security login create -user-or-group-name SNMPv3_USER -application snmp
-authentication-method usm -remote-switch-ipaddress ADDRESS
```

The following command establishes the SNMPv3 username with CSHM:

```
system switch ethernet modify -device DEVICE -snmp-version SNMPv3
-community-or-username SNMPv3_USER
```

Steps

1. Set up the SNMPv3 user on the switch to use authentication and encryption:

```
nv show system snmp-server
```

```
cumulus@sw1:~$ nv show system snmp-server
                                     applied
-----
[username]          SNMPv3_USER
[username]          limiteduser1
[username]          testuserauth
[username]          testuserauthaes
[username]          testusernoauth
trap-link-up
  check-frequency   60
trap-link-down
  check-frequency   60
[listening-address] all
[readonly-community] $nvsec$94d69b56e921aec1790844eb53e772bf
state               enabled
cumulus@sw1:~$
```

2. Set up the SNMPv3 user on the ONTAP side:

```
security login create -user-or-group-name SNMPv3User -application
snmp -authentication-method usm -remote-switch-ipaddress
10.231.80.212
```

```
cluster1::*> security login create -user-or-group-name SNMPv3User
-application snmp -authentication-method usm -remote-switch
-ipaddress 10.231.80.212
```

Enter the authoritative entity's EngineID [remote EngineID]:

Which authentication protocol do you want to choose (none, md5, sha, sha2-256)

[none]: **md5**

Enter the authentication protocol password (minimum 8 characters long):

Enter the authentication protocol password again:

Which privacy protocol do you want to choose (none, des, aes128)

[none]: **aes128**

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:

3. Configure CSHM to monitor with the new SNMPv3 user:

```
system switch ethernet show-all -device "sw1 (b8:59:9f:09:7c:22)"
-instance
```

```

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance
                                     Device Name: sw1
(b8:59:9f:09:7c:22)
                                     IP Address: 10.231.80.212
                                     SNMP Version: SNMPv2c
                                     Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
      Community String or SNMPv3 Username: cshml!
                                     Model Number: MSN2100-CB2FC
                                     Switch Network: cluster-network
                                     Software Version: Cumulus Linux
version 5.11.0 running on Mellanox Technologies Ltd. MSN2100
      Reason For Not Monitoring: None
      Source Of Switch Version: LLDP
      Is Monitored ?: true
      Serial Number of the Device: MT2110X06399 <----
serial number to check
                                     RCF Version: MSN2100-RCF-v1.9X6-
Cluster-LLDP Aug-18-2022

cluster1::*>
cluster1::*> system switch ethernet modify -device "sw1
(b8:59:9f:09:7c:22)" -snmp-version SNMPv3 -community-or-username
SNMPv3User

```

4. Verify that the serial number to be queried with the newly created SNMPv3 user is the same as detailed in the previous step after the CSHM polling period has completed.

```
system switch ethernet polling-interval show
```

```

cluster1::*> system switch ethernet polling-interval show
Polling Interval (in minutes): 5

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance
Device Name: sw1
IP Address: 10.231.80.212
SNMP Version: SNMPv3
Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
Community String or SNMPv3 Username: SNMPv3User
Model Number: MSN2100-CB2FC
Switch Network: cluster-network
Software Version: Cumulus Linux
version 5.11.0 running on Mellanox Technologies Ltd. MSN2100
Reason For Not Monitoring: None
Source Of Switch Version: LLDP
Is Monitored?: true
Serial Number of the Device: MT2110X06399 <----
serial number to check
RCF Version: MSN2100-RCF-v1.9X6-
Cluster-LLDP Aug-18-2022

```

Check switch health

Health check overview

Health monitors proactively monitor certain critical conditions in your cluster and raise alerts if they detect a fault or risk.

To view the currently raised Ethernet switch health monitor alerts, run the command: `system health alert show -monitor ethernet-switch`

To view the available Ethernet switch health monitor alerts, run the command: `system health alert definition show -monitor ethernet-switch`

Manage the monitoring of Ethernet Switches

In most cases, Ethernet switches are automatically discovered by ONTAP and monitored by CSHM. The Reference Configuration File (RCF) applied to the switch, among other things, enables the Cisco Discovery Protocol (CDP) and/or the Link Layer Discovery Protocol (LLDP). However, you might need to manually add a switch that is not discovered or remove a switch that is no longer in use. You can also stop active

monitoring while retaining the switch in the configuration, such as during maintenance.

Create a switch entry so that ONTAP can monitor it

Use the `system switch ethernet create` command to manually configure and enable monitoring for a specified Ethernet switch. This is helpful if ONTAP does not add the switch automatically, or if you previously removed the switch and want to re-add it.

```
system switch ethernet create -device DeviceName -address 1.2.3.4 -snmp
-version SNMPv2c -community-or-username cshml! -model NX3132V -type
cluster-network
```

A typical example is adding a switch named [DeviceName], with IP address 1.2.3.4, and SNMPv2c credentials set to **cshml!**. Use `-type storage-network` instead of `-type cluster-network` if you are configuring a storage switch.

Disable monitoring without deleting the switch

If you want to pause or stop monitoring for a certain switch, but still retain it for future monitoring, modify its `is-monitoring-enabled-admin` parameter instead of deleting it.

For example:

```
system switch ethernet modify -device DeviceName -is-monitoring-enabled
-admin false
```

This lets you preserve switch details and configuration without generating new alerts or re-discoveries.

Remove a switch you no longer need

Use `system switch ethernet delete` to delete a switch that has been disconnected or is no longer required:

```
system switch ethernet delete -device DeviceName
```

By default, this command succeeds only if ONTAP does not currently detect the switch through CDP or LLDP. To remove a discovered switch, use the `-force` parameter:

```
system switch ethernet delete -device DeviceName -force
```

When `-force` is used, the switch might be re-added automatically if ONTAP detects it again.

Verify the monitoring of Ethernet switches

The Ethernet switch health monitor (CSHM) automatically attempts to monitor the switches that it discovers; however, monitoring might not happen automatically if the

switches are not configured correctly. You should verify that the health monitor is properly configured to monitor your switches.

Confirm monitoring of the connected Ethernet switches

To confirm that the connected Ethernet switches are being monitored, run:

```
system switch ethernet show
```

If the `Model` column displays **OTHER** or the `IS Monitored` field displays **false**, then ONTAP cannot monitor the switch. A value of **OTHER** typically indicates that ONTAP does not support that switch for health monitoring.

The `IS Monitored` field is set to **false** for the reason specified in the `Reason` field.



If a switch is not listed in the command output, ONTAP likely has not discovered it. Confirm that the switch is cabled correctly. If necessary, you can add the switch manually. See [Manage the monitoring of Ethernet Switches](#) for further details.

Confirm firmware and RCF versions are up to date

Make sure that the switch is running the latest supported firmware and that a compatible reference configuration file (RCF) has been applied. More information is available on the [NetApp Support Downloads page](#).

By default, the health monitor uses SNMPv2c with the community string **csbm1!** for monitoring, but SNMPv3 can also be configured.

If you need to change the default SNMPv2c community string, make sure that the desired SNMPv2c community string has been configured on the switch.

```
system switch ethernet modify -device SwitchA -snmp-version SNMPv2c  
-community-or-username newCommunity!
```



See [Optional: Configure SNMPv3](#) for details on configuring SNMPv3 for use.

Confirm management network connection

Verify that the switch's management port is connected to the management network.

A correct management port connection is required for ONTAP to perform SNMP queries and log collection.

Troubleshoot alerts

Alerts are raised if a fault, risk, or critical condition is detected for an Ethernet switch in your cluster.

If there are raised alerts, the system health status reports a degraded status for the cluster. The alerts raised include the information that you need to respond to degraded system health.

To view the available Ethernet switch health monitor alerts, run the command: `system health alert definition show -monitor ethernet-switch`

See the Knowledge Base article [Switch Health Monitor Alert Resolution Guide](#) for advanced resolution details of alerts.

Log collection

Log collection overview

With log collection set up, you can enable an hourly collection of periodic data that is collected by AutoSupport, and request detailed support logs.

See [Configure log collection](#) for further details.

Troubleshoot log collection

If you encounter any of the following error statuses reported by the log collection feature (visible in the output of the `system switch ethernet log show` command), try the corresponding debug steps:

Log collection error status	Resolution
RSA keys not present	Regenerate ONTAP SSH keys.
Switch password error	Verify credentials, test SSH connectivity, and regenerate ONTAP SSH keys. Review the switch documentation or contact NetApp support for instructions.
ECDSA keys not present for FIPS	If FIPS mode is enabled, ECDSA keys need to be generated on the switch before retrying.
Pre-existing log found	Remove the previous log collection file on the switch.
Switch dump log error	Ensure the switch user has log collection permissions. Refer to the prerequisites above.



If the resolution details do not work, contact NetApp support.

End-of-availability switches

End-of-availability

The following switches are no longer available for purchase, but are still supported.

- [Cisco Nexus 3232C](#)
- [Cisco Nexus 3132Q-V](#)
- [Cisco Nexus 92300YC](#)
- [NetApp CN1610](#)

End-of-sale and End-of-life announcements

- [End-of-Sale and End-of-Life Announcement for the Cisco Nexus 3232C](#)
- [End-of-Sale and End-of-Life Announcement for the Cisco Nexus 31108PC-V, 31108TC-V and Nexus 3132Q-V](#)
- [End-of-Sale and End-of-Life Announcement for the Cisco N9K-C93120TX, N9K-C92300YC](#)
- [End-of-Sale and End-of-Life Announcement for the Cisco Nexus 5500 Series Switches](#)
- [End of Availability: NetApp CN1610 Cluster Interconnect SKUs](#)

Cisco Nexus 3232C

Overview

Overview of installation and configuration for Cisco Nexus 3232c switches

Cisco Nexus 3232C switches can be used as cluster switches in your AFF or FAS cluster. Cluster switches allow you to build ONTAP clusters with more than two nodes.

Initial configuration overview

To initially configure a Cisco Nexus 3232c switch on systems running ONTAP, follow these steps:

1. [Complete Cisco Nexus 3232C cabling worksheet](#). The sample cabling worksheet provides examples of recommended port assignments from the switches to the controllers. The blank worksheet provides a template that you can use in setting up your cluster.
2. [Install a Cisco Nexus 3232C cluster switch in a NetApp cabinet](#). Install the Cisco Nexus 3232C cluster switch and pass-through panel in a NetApp cabinet with the standard brackets that are included with the switch.
3. [Configure the 3232C cluster switch](#). Set up and configure the Cisco Nexus 3232C switch.
4. [Prepare to install NX-OS software and Reference Configuration File](#). Prepare to install the NX-OS software and the Reference Configuration File (RCF).
5. [Install the NX-OS software](#). Install the NX-OS software on the Nexus 3232C cluster switch.
6. [Install the Reference Configuration File \(RCF\)](#). 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.

Additional information

Before you begin installation or maintenance, be sure to review the following:

- [Configuration requirements](#)
- [Required documentation](#)
- [Smart Call Home requirements](#)

Configuration requirements for Cisco Nexus 3232C switches

For Cisco Nexus 3232C switch installation and maintenance, be sure to review configuration and network requirements.

Configuration requirements

To configure your cluster, you need the appropriate number and type of cables and cable connectors for your switches. Depending on the type of switch you are initially configuring, you need to connect to the switch console port with the included console cable; you also need to provide specific network information.

Network requirements

You need the following network information for all switch configurations:

- IP subnet for management network traffic
- Host names and IP addresses for each of the storage system controllers and all applicable switches
- Most storage system controllers are managed through the e0M interface by connecting to the Ethernet service port (wrench icon). On AFF A800 and AFF A700 systems, the e0M interface uses a dedicated Ethernet port.

Refer to the [Hardware Universe](#) for latest information.

Documentation requirements for Cisco Nexus 3232C switches

For Cisco Nexus 3232C switch installation and maintenance, be sure to review all recommended documentation.

Switch documentation

To set up the Cisco Nexus 3232C switches, you need the following documentation from the [Cisco Nexus 3000 Series Switches Support](#) page.

Document title	Description
<i>Nexus 3000 Series Hardware Installation Guide</i>	Provides detailed information about site requirements, switch hardware details, and installation options.
<i>Cisco Nexus 3000 Series Switch Software Configuration Guides</i> (choose the guide for the NX-OS release installed on your switches)	Provides initial switch configuration information that you need before you can configure the switch for ONTAP operation.

Document title	Description
<i>Cisco Nexus 3000 Series NX-OS Software Upgrade and Downgrade Guide</i> (choose the guide for the NX-OS release installed on your switches)	Provides information on how to downgrade the switch to ONTAP supported switch software, if necessary.
<i>Cisco Nexus 3000 Series NX-OS Command Reference Master Index</i>	Provides links to the various command references provided by Cisco.
<i>Cisco Nexus 3000 MIBs Reference</i>	Describes the Management Information Base (MIB) files for the Nexus 3000 switches.
<i>Nexus 3000 Series NX-OS System Message Reference</i>	Describes the system messages for Cisco Nexus 3000 series switches, those that are informational, and others that might help diagnose problems with links, internal hardware, or the system software.
<i>Cisco Nexus 3000 Series NX-OS Release Notes</i> (choose the notes for the NX-OS release installed on your switches)	Describes the features, bugs, and limitations for the Cisco Nexus 3000 Series.
Regulatory, Compliance, and Safety Information for the Cisco Nexus 6000, Cisco Nexus 5000 Series, Cisco Nexus 3000 Series, and Cisco Nexus 2000 Series	Provides international agency compliance, safety, and statutory information for the Nexus 3000 series switches.

ONTAP systems documentation

To set up an ONTAP system, you need the following documents for your version of the operating system from the [ONTAP 9 Documentation Center](#).

Name	Description
Controller-specific <i>Installation and Setup Instructions</i>	Describes how to install NetApp hardware.
ONTAP documentation	Provides detailed information about all aspects of the ONTAP releases.
Hardware Universe	Provides NetApp hardware configuration and compatibility information.

Rail kit and cabinet documentation

To install a 3232C Cisco switch in a NetApp cabinet, see the following hardware documentation.

Name	Description
42U System Cabinet, Deep Guide	Describes the FRUs associated with the 42U system cabinet, and provides maintenance and FRU replacement instructions.
Install a Cisco Nexus 3232C switch in a NetApp Cabinet	Describes how to install a Cisco Nexus 3232C switch in a four-post NetApp cabinet.

Smart Call Home requirements

To use Smart Call Home feature, review the following guidelines.

Smart Call Home monitors the hardware and software components on your network. When a critical system configuration occurs, it generates an email-based notification and raises an alert to all the recipients that are configured in your destination profile. To use Smart Call Home, you must configure a cluster network switch to communicate using email with the Smart Call Home system. In addition, you can optionally set up your cluster network switch to take advantage of Cisco's embedded Smart Call Home support feature.

Before you can use Smart Call Home, be aware of the following considerations:

- An email server must be in place.
- The switch must have IP connectivity to the email server.
- The contact name (SNMP server contact), phone number, and street address information must be configured. This is required to determine the origin of messages received.
- A CCO ID must be associated with an appropriate Cisco SMARTnet Service contract for your company.
- Cisco SMARTnet Service must be in place for the device to be registered.

The [Cisco support site](#) contains information about the commands to configure Smart Call Home.

Install hardware

Complete Cisco Nexus 3232C cabling worksheet

If you want to document the supported platforms, download a PDF of this page and complete the cabling worksheet.

The sample cabling worksheet provides examples of recommended port assignments from the switches to the controllers. The blank worksheet provides a template that you can use in setting up your cluster.

Each switch can be configured as a single 100GbE, 40GbE port or 4 x 10GbE ports.

Sample cabling worksheet

The sample port definition on each pair of switches is as follows:

Cluster switch A		Cluster switch B	
Switch port	Node and port usage	Switch port	Node and port usage

Cluster switch A		Cluster switch B	
1	4x10GbE/4x25GbE or 40/100GbE node	1	4x10GbE/4x25GbE or 40/100GbE node
2	4x10GbE/4x25GbE or 40/100GbE node	2	4x10GbE/4x25GbE or 40/100GbE node
3	4x10GbE/4x25GbE or 40/100GbE node	3	4x10GbE/4x25GbE or 40/100GbE node
4	4x10GbE/4x25GbE or 40/100GbE node	4	4x10GbE/4x25GbE or 40/100GbE node
5	4x10GbE/4x25GbE or 40/100GbE node	5	4x10GbE/4x25GbE or 40/100GbE node
6	4x10GbE/4x25GbE or 40/100GbE node	6	4x10GbE/4x25GbE or 40/100GbE node
7	4x10GbE/4x25GbE or 40/100GbE node	7	4x10GbE/4x25GbE or 40/100GbE node
8	4x10GbE/4x25GbE or 40/100GbE node	8	4x10GbE/4x25GbE or 40/100GbE node
9	4x10GbE/4x25GbE or 40/100GbE node	9	4x10GbE/4x25GbE or 40/100GbE node
10	4x10GbE/4x25GbE or 40/100GbE node	10	4x10GbE/4x25GbE or 40/100GbE node
11	4x10GbE/4x25GbE or 40/100GbE node	11	4x10GbE/4x25GbE or 40/100GbE node
12	4x10GbE/4x25GbE or 40/100GbE node	12	4x10GbE/4x25GbE or 40/100GbE node
13	4x10GbE/4x25GbE or 40/100GbE node	13	4x10GbE/4x25GbE or 40/100GbE node
14	4x10GbE/4x25GbE or 40/100GbE node	14	4x10GbE/4x25GbE or 40/100GbE node
15	4x10GbE/4x25GbE or 40/100GbE node	15	4x10GbE/4x25GbE or 40/100GbE node

Cluster switch A		Cluster switch B	
16	4x10GbE/4x25GbE or 40/100GbE node	16	4x10GbE/4x25GbE or 40/100GbE node
17	4x10GbE/4x25GbE or 40/100GbE node	17	4x10GbE/4x25GbE or 40/100GbE node
18	4x10GbE/4x25GbE or 40/100GbE node	18	4x10GbE/4x25GbE or 40/100GbE node
19	40G/100GbE node 19	19	40G/100GbE node 19
20	40G/100GbE node 20	20	40G/100GbE node 20
21	40G/100GbE node 21	21	40G/100GbE node 21
22	40G/100GbE node 22	22	40G/100GbE node 22
23	40G/100GbE node 23	23	40G/100GbE node 23
24	40G/100GbE node 24	24	40G/100GbE node 24
25 through 30	Reserved	25 through 30	Reserved
31	100GbE ISL to switch B port 31	31	100GbE ISL to switch A port 31
32	100GbE ISL to switch B port 32	32	100GbE ISL to switch A port 32

Blank cabling worksheet

You can use the blank cabling worksheet to document the platforms that are supported as nodes in a cluster. The *Supported Cluster Connections* section of the [Hardware Universe](#) defines the cluster ports used by the platform.

Cluster switch A		Cluster switch B	
Switch port	Node/port usage	Switch port	Node/port usage
1		1	
2		2	
3		3	

Cluster switch A		Cluster switch B	
4		4	
5		5	
6		6	
7		7	
8		8	
9		9	
10		10	
11		11	
12		12	
13		13	
14		14	
15		15	
16		16	
17		17	
18		18	
19		19	
20		20	
21		21	
22		22	
23		23	
24		24	
25 through 30	Reserved	25 through 30	Reserved

Cluster switch A		Cluster switch B	
31	100GbE ISL to switch B port 31	31	100GbE ISL to switch A port 31
32	100GbE ISL to switch B port 32	32	100GbE ISL to switch A port 32

Configure the 3232C cluster switch

Follow this procedure to set up and configure the Cisco Nexus 3232C switch.

Before you begin

- Access to an HTTP, FTP or TFTP server at the installation site to download the applicable NX-OS and reference configuration file (RCF) releases.
- Applicable NX-OS version, downloaded from the [Cisco software download](#) page.
- Required cluster network and management network switch documentation.

See [Required documentation](#) for more information.

- Required controller documentation and ONTAP documentation.

[NetApp documentation](#)

- Applicable licenses, network and configuration information, and cables.
- Completed cabling worksheets.
- Applicable NetApp cluster network and management network RCFs, downloaded from the NetApp Support Site at mysupport.netapp.com for the switches that you receive. All Cisco cluster network and management network switches arrive with the standard Cisco factory-default configuration. These switches also have the current version of the NX-OS software, but do not have the RCFs loaded.

Steps


1. Rack the cluster network and management network switches and controllers.


If you are installing your...	Then...
Cisco Nexus 3232C in a NetApp system cabinet	See the <i>Installing a Cisco Nexus 3232C cluster switch and pass-through panel in a NetApp cabinet</i> guide for instructions to install the switch in a NetApp cabinet.
Equipment in a Telco rack	See the procedures provided in the switch hardware installation guides and the NetApp installation and setup instructions.

2. Cable the cluster network and management network switches to the controllers using the completed cabling worksheets.
3. Power on the cluster network and management network switches and controllers.
4. Perform an initial configuration of the cluster network switches.

Provide applicable responses to the following initial setup questions when you first boot the switch. Your

site's security policy defines the responses and services to enable.

Prompt	Response
Abort Auto Provisioning and continue with normal setup? (yes/no)	Respond with yes . The default is no.
Do you want to enforce secure password standard? (yes/no)	Respond with yes . The default is yes.
Enter the password for admin.	The default password is "admin"; you must create a new, strong password. A weak password can be rejected.
Would you like to enter the basic configuration dialog? (yes/no)	Respond with yes at the initial configuration of the switch.
Create another login account? (yes/no)	Your answer depends on your site's policies on alternate administrators. The default is no .
Configure read-only SNMP community string? (yes/no)	Respond with no . The default is no.
Configure read-write SNMP community string? (yes/no)	Respond with no . The default is no.
Enter the switch name.	The switch name is limited to 63 alphanumeric characters.
Continue with Out-of-band (mgmt0) management configuration? (yes/no)	Respond with yes (the default) at that prompt. At the mgmt0 IPv4 address: prompt, enter your IP address: ip_address.
Configure the default-gateway? (yes/no)	Respond with yes . At the IPv4 address of the default-gateway: prompt, enter your default_gateway.
Configure advanced IP options? (yes/no)	Respond with no . The default is no.
Enable the telnet service? (yes/no)	Respond with no . The default is no.
Enabled SSH service? (yes/no)	Respond with yes . The default is yes. <div style="display: flex; align-items: center;">  <div> <p>SSH is recommended when using Ethernet Switch Health Monitor (CSHM) for its log collection features. SSHv2 is also recommended for enhanced security.</p> </div> </div>

Prompt	Response
Enter the type of SSH key you want to generate (dsa/rsa/rsa1).	The default is rsa .
Enter the number of key bits (1024-2048).	Enter the number of key bits from 1024-2048.
Configure the NTP server? (yes/no)	Respond with no . The default is no.
Configure default interface layer (L3/L2):	Respond with L2 . The default is L2.
Configure default switch port interface state (shut/noshut):	Respond with noshut . The default is noshut.
Configure CoPP system profile (strict/moderate/lenient/dense):	Respond with strict . The default is strict.
Would you like to edit the configuration? (yes/no)	You should see the new configuration at this point. Review and make any necessary changes to the configuration you just entered. Respond with no at the prompt if you are satisfied with the configuration. Respond with yes if you want to edit your configuration settings.
Use this configuration and save it? (yes/no)	Respond with yes to save the configuration. This automatically updates the kickstart and system images. <div style="display: flex; align-items: center;">  <p>If you do not save the configuration at this stage, none of the changes will be in effect the next time you reboot the switch.</p> </div>

5. Verify the configuration choices you made in the display that appears at the end of the setup, and make sure that you save the configuration.
6. Check the version on the cluster network switches, and if necessary, download the NetApp-supported version of the software to the switches from the [Cisco software download](#) page.

What's next?

[Prepare to install NX-OS and RCF.](#)

Install a Cisco Nexus 3232C cluster switch in a NetApp cabinet

Depending on your configuration, you might need to install the Cisco Nexus 3232C cluster switch and pass-through panel in a NetApp cabinet with the standard brackets that are included with the switch.

Before you begin

- The initial preparation requirements, kit contents, and safety precautions in the [Cisco Nexus 3000 Series Hardware Installation Guide](#).
- For each switch, the eight 10-32 or 12-24 screws and clip nuts to mount the brackets and slider rails to the front and rear cabinet posts.
- Cisco standard rail kit to install the switch in a NetApp cabinet.



The jumper cords are not included with the pass-through kit and should be included with your switches. If they were not shipped with the switches, you can order them from NetApp (part number X1558A-R6).

Steps


1. Install the pass-through blanking panel in the NetApp cabinet.

The pass-through panel kit is available from NetApp (part number X8784-R6).

The NetApp pass-through panel kit contains the following hardware:


- One pass-through blanking panel
- Four 10-32 x .75 screws
- Four 10-32 clip nuts
 - a. Determine the vertical location of the switches and blanking panel in the cabinet.

In this procedure, the blanking panel will be installed in U40.
 - b. Install two clip nuts on each side in the appropriate square holes for front cabinet rails.
 - c. Center the panel vertically to prevent intrusion into adjacent rack space, and then tighten the screws.
 - d. Insert the female connectors of both 48-inch jumper cords from the rear of the panel and through the brush assembly.




(1) Female connector of the jumper cord.

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.
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 cabinet RU41 and 42.

3. Install the slider rails in the cabinet.
 - 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 cabinet posts.

b. Repeat step 4a for the right side rear post.


c. Repeat steps 4a and 4b at the RU41 locations on the cabinet.

4. Install the switch in the cabinet.



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) As the chassis is pushed toward the rear posts, align the two rear rack-mount guides with the slider rails.

(2) Gently slide the switch until the front rack-mount brackets are flush with the front posts.

b. Attach the switch to the cabinet.



(1) With one person holding the front of the chassis level, the other person should fully tighten the four rear screws to the cabinet 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.



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.

- 5. When the switches are installed, connect the jumper cords to the switch power inlets.
- 6. Connect the male plugs of both jumper cords to the closest available PDU outlets.



To maintain redundancy, the two cords must be connected to different PDUs.

- 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.

Review cabling and configuration considerations

Before configuring your Cisco 3232C switch, review the following considerations.

Support for NVIDIA CX6, CX6-DX, and CX7 Ethernet ports

If connecting a switch port to an ONTAP controller using NVIDIA ConnectX-6 (CX6), ConnectX-6 Dx (CX6-DX), or ConnectX-7 (CX7) NIC ports, you must hard-code the switch port speed.

```
(cs1)(config)# interface Ethernet1/19
For 100GbE speed:
(cs1)(config-if)# speed 100000
For 40GbE speed:
(cs1)(config-if)# speed 40000
(cs1)(config-if)# no negotiate auto
(cs1)(config-if)# exit
(cs1)(config)# exit
Save the changes:
(cs1)# copy running-config startup-config
```

See the [Hardware Universe](#) for more information on switch ports.

Configure software

Prepare to install NX-OS software and Reference Configuration File (RCF)

Before you install the NX-OS software and the Reference Configuration File (RCF), follow this procedure.

About the examples

The examples in this procedure use two nodes. These nodes use two 10GbE cluster interconnect ports e0a and e0b.

See the [Hardware Universe](#) to verify the correct cluster ports on your platforms.



The command outputs might vary depending on different releases of ONTAP.

Switch and node nomenclature

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.

About this task

The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

Steps

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=x h
```

where `x` is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (*>) appears.

3. Display how many cluster interconnect interfaces are configured in each node for each cluster interconnect switch:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
```

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	
-----	-----	-----	-----	

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.
```

4. Check the administrative or operational status of each cluster interface.
 - a. Display the network port attributes:

```
network port show -ipSpace Cluster
```

Show example

```
cluster1::*> network port show -ipSpace Cluster

Node: cluster1-02

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

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

Show example

```
cluster1::*> network interface show -vserver Cluster
```

Current Vserver Port	Logical Current Interface Home	Is	Status Admin/Oper	Network Address/Mask	Node	
Cluster	cluster1-01	e0a	true	up/up	169.254.209.69/16	
cluster1-01	cluster1-01	e0b	true	up/up	169.254.49.125/16	
cluster1-01	cluster1-02	e0a	true	up/up	169.254.47.194/16	
cluster1-02	cluster1-02	e0b	true	up/up	169.254.19.183/16	

4 entries were displayed.

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	Date	LIF
Loss		
-----	-----	-----
-----	-----	-----
cluster1-01		
3/5/2022 19:21:18 -06:00	cluster1-01_clus2	cluster1-02_clus1
none		
3/5/2022 19:21:20 -06:00	cluster1-01_clus2	cluster1-02_clus2
none		
.		
.		
cluster1-02		
3/5/2022 19:21:18 -06:00	cluster1-02_clus2	cluster1-01_clus1
none		
3/5/2022 19:21:20 -06:00	cluster1-02_clus2	cluster1-01_clus2
none		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is 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)

```

- Verify that the auto-revert command is enabled on all cluster LIFs:
network interface show -vserver Cluster -fields auto-revert

Show example

```

cluster1::*> network interface show -vserver Cluster -fields auto-
revert

```

Vserver	Logical Interface	Auto-revert
Cluster	cluster1-01_clus1	true
	cluster1-01_clus2	true
	cluster1-02_clus1	true
	cluster1-02_clus2	true

4 entries were displayed.

What's next?

[Install NX-OS software](#)

Install the NX-OS software

You can use this procedure to install the NX-OS software on the Nexus 3232C cluster switch.

Review requirements

Before you begin

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).
- [Cisco Ethernet switch page](#). Consult the switch compatibility table for the supported ONTAP and NX-OS versions.
- [Cisco Nexus 3000 Series Switches](#). Refer to the appropriate software and upgrade guides available on the Cisco web site for complete documentation on the Cisco switch upgrade and downgrade procedures.

Install the software

The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

Be sure to complete the procedure in [Prepare to install NX-OS and RCF](#), and then follow the steps below.

Steps

1. Connect the cluster switch to the management network.
2. Use the `ping` command to verify connectivity to the server hosting the NX-OS software and the RCF.

Show example

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.
```

3. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

Show example

```
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::*>
```

4. 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
```

Show example

```
cluster1::*> network port show -role cluster

Node: cluster1-01

Ignore

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

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

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
```

```
Health Health Speed (Mbps)
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
```

Show example

```
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
```

Show example

```
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: FOCXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
  9.3(5)
  Version Source: CDP

cs2                                         cluster-network                    10.233.205.91
N3K-C3232C
  Serial Number: FOCXXXXXXGS
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
  9.3(5)
  Version Source: CDP
cluster1::*>
```

5. Disable auto-revert on the cluster LIFs. The cluster LIFs fail over to the partner cluster switch and remain there as you perform the upgrade procedure on the targeted switch:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

6. Copy the NX-OS software and EPLD images to the Nexus 3232C switch.

Show example

```
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.
```

7. Verify the running version of the NX-OS software:

```
show version
```

Show example

```
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 FOCXXXXXXGD

  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#
```

8. Install the NX-OS image.

Installing the image file causes it to be loaded every time the switch is rebooted.

Show example

```
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#
```

9. Verify the new version of NX-OS software after the switch has rebooted:

```
show version
```


Show example

```
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 FOCXXXXXXGS

  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#
```

10. Upgrade the EPLD image and reboot the switch.

Show example

```
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.
```

```
cs2#
```

11. If you are upgrading to NX-OS version 9.3(11), you must upgrade the EPLD golden image and reboot the switch once again. Otherwise, skip to step 12.

See [EPLD Upgrade Release Notes, Release 9.3\(11\)](#) for further details.

Show example

```
cs2# install epld bootflash:n9000-epld.9.3.11.img module 1 golden
Digital signature verification is successful
Compatibility check:
Module          Type          Upgradable    Impact        Reason
-----
-----
          1          SUP          Yes          Disruptive    Module
Upgradable

Retrieving EPLD versions.... Please wait.
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 : MI FPGA [Programming] : 100.00% (    64 of    64 sect)
Module 1 : IO FPGA [Programming] : 100.00% (    64 of    64 sect)
Module 1 EPLD upgrade is successful.
Module          Type          Upgrade-Result
-----
-----
          1          SUP          Success

EPLDs upgraded.

Module 1 EPLD upgrade is successful.
cs2#
```

12. After the switch reboot, log in to verify that the new version of EPLD loaded successfully.

Show example

```
cs2# show version module 1 epld
```

EPLD Device	Version
MI FPGA	0x12
IO FPGA	0x12

13. Verify the health of cluster ports on the cluster.

a. Verify that cluster ports are up and healthy across all nodes in the cluster:

```
network port show -role cluster
```

Show example

```
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
```

```
Health Health Speed (Mbps)
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.

```
network device-discovery show -protocol cdp
```

Show example

```
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: FOCXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                  9.3(5)
  Version Source: CDP

cs2              cluster-network  10.233.205.91
```



```

N3K-C3232C
  Serial Number: FOCXXXXXXGS
    Is Monitored: true
      Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                9.3(5)
  Version Source: CDP

2 entries were displayed.

```

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.

```

14. Verify that the cluster is healthy:

```
cluster show
```

Show example

```

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::*>

```

15. Repeat steps 6 to 14 on switch cs1.

16. Enable auto-revert on the cluster LIFs.

```
network interface modify -vserver Cluster -lif * -auto-revert true
```

17. Verify that the cluster LIFs have reverted to their home port:

```
network interface show -role cluster
```

Show example

```
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::*>
```

If any cluster LIFs have not returned to their home ports, revert them manually from the local node:

```
network interface revert -vserver Cluster -lif <lif_name>
```

What's next?

[Install RCF config file.](#)

Install the Reference Configuration File (RCF)

Follow this procedure to 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. See the Knowledge Base article [How to clear](#)

[configuration on a Cisco interconnect switch while retaining remote connectivity](#) for further information when upgrading your RCF.

Review requirements

Before you begin

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).
- The current Reference Configuration File (RCF).
- A console connection to the switch, required when installing the RCF.
- [Cisco Ethernet switch page](#) Consult the switch compatibility table for the supported ONTAP and RCF versions. Note that there can be command dependencies between the command syntax in the RCF and that found in versions of NX-OS.
- [Cisco Nexus 3000 Series Switches](#). Refer to the appropriate software and upgrade guides available on the Cisco web site for complete documentation on the Cisco switch upgrade and downgrade procedures.

Install the file

About the examples

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.

About this task

The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure migrates all of the cluster LIFs to the operational partner switch while performing the steps on the target switch.

Be sure to complete the procedure in [Prepare to install NX-OS and RCF](#), and then follow the steps below.

Steps

1. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

Show example

```
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::*>
```

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
```

Show example

```
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
```

```
Health Health Speed (Mbps)
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
```

Show example

```
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
```

Show example

```
cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                               Type                Address
Model
-----
cs1                                   cluster-network     10.233.205.92
NX3232C
  Serial Number: FOXXXXXXXXGS
  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: FOXXXXXXXXGD
  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.
```

3. Disable auto-revert on the cluster LIFs.

Show example

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto
-revert false
```

4. On cluster switch cs2, shut down the ports connected to the cluster ports of the nodes.

Show example

```
cs2(config)# interface eth1/1/1-2,eth1/7-8  
cs2(config-if-range)# shutdown
```

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
```

Show example

```
cluster1::*> network interface show -role cluster
```

Current Is	Logical	Status	Network	Current
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::*>
```

6. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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::*>
```

7. If you have not already done so, save a copy of the current switch configuration by copying the output of the following command to a text file:

```
show running-config
```

8. Clean the configuration on switch cs2 and reboot the switch.



When updating or applying a new RCF, you must erase the switch settings and perform basic configuration. You must be connected to the switch serial console port to set up the switch again.

- a. Clean the configuration:

Show example

```
(cs2)# write erase

Warning: This command will erase the startup-configuration.

Do you wish to proceed anyway? (y/n) [n] y
```

- b. Reboot the switch:

Show example

```
(cs2)# reload

Are you sure you would like to reset the system? (y/n) y
```

9. Perform a basic setup of the switch. See [Configure the 3232C cluster switch](#) for details.

10. 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.

Show example

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)...
```

11. 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.

Show example

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
```

12. Examine the banner output from the `show banner motd` command. You must read and follow the instructions under **Important Notes** to make sure the proper configuration and operation of the switch.

Show example

```
cs2# show banner motd

*****
*****
* NetApp 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
*****
*****
```



When applying the RCF for the first time, the **ERROR: Failed to write VSH commands** message is expected and can be ignored.

13. 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.

14. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.

15. 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
```

16. 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
```

17. Apply the same RCF and save the running configuration for a second time.

Show example

```
cs2# copy Nexus_3232C_RCF_v1.6-Cluster-HA-Breakout.txt running-  
config echo-commands  
cs2# copy running-config startup-config  
[#####] 100% Copy complete
```

18. 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
```

Show example

```
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

Health Health Speed (Mbps)
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).

Show example

```
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: FOXXXXXXXXGD
  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: FOXXXXXXXXGS
    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.

```

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.

```



It can take up to 5 minutes for the cluster nodes to report as healthy.

19. On cluster switch cs1, shut down the ports connected to the cluster ports of the nodes.

Show example

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

```

20. 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

```

Show example

```
cluster1::*> network interface show -role cluster
          Logical          Status      Network          Current
Current Is
Vserver   Interface              Admin/Oper  Address/Mask     Node
Port      Home
-----
Cluster
cluster1-01 cluster1-01_clus1 up/up      169.254.3.4/23
          e0d             false
cluster1-01 cluster1-01_clus2 up/up      169.254.3.5/23
          e0d             true
cluster1-02 cluster1-02_clus1 up/up      169.254.3.8/23
          e0d             false
cluster1-02 cluster1-02_clus2 up/up      169.254.3.9/23
          e0d             true
cluster1-03 cluster1-03_clus1 up/up      169.254.1.3/23
          e0b             false
cluster1-03 cluster1-03_clus2 up/up      169.254.1.1/23
          e0b             true
cluster1-04 cluster1-04_clus1 up/up      169.254.1.6/23
          e0b             false
cluster1-04 cluster1-04_clus2 up/up      169.254.1.7/23
          e0b             true
8 entries were displayed.
cluster1::*>
```

21. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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::*>
```

22. Repeat Steps 7 to 17 on switch cs1.
23. Enable auto-revert on the cluster LIFs.

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto
-revert true
```

24. 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
```

25. Verify that the switch ports connected to the cluster ports are up.

Show example

```
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) --
.
.
```

26. Verify that the ISL between cs1 and cs2 is functional:

```
show port-channel summary
```

Show example

```
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#
```

27. Verify that the cluster LIFs have reverted to their home port:

```
network interface show -role cluster
```

Show example

```
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::*>
```

If any cluster LIFS have not returned to their home ports, revert them manually:

```
network interface revert -vserver vservice_name -lif lif_name
```

28. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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::*>
```

29. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	Date	LIF
Loss		
-----	-----	-----
-----	-----	-----
cluster1-01		
3/5/2022 19:21:18 -06:00	cluster1-01_clus2	cluster1-02_clus1
none		
3/5/2022 19:21:20 -06:00	cluster1-01_clus2	cluster1-02_clus2
none		
.		
.		
cluster1-02		
3/5/2022 19:21:18 -06:00	cluster1-02_clus2	cluster1-01_clus1
none		
3/5/2022 19:21:20 -06:00	cluster1-02_clus2	cluster1-01_clus2
none		
.		
.		
cluster1-03		
.		
.		
.		
.		
cluster1-04		
.		
.		
.		
.		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```
cluster1::*> cluster ping-cluster -node local
Host is 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)
```

What's next?

[Verify SSH configuration.](#)

Verify your SSH configuration

If you are using the Ethernet Switch Health Monitor (CSHM) and log collection features, verify that SSH and SSH keys are enabled on the cluster switches.

Steps

1. Verify that SSH is enabled:

```
(switch) show ssh server  
ssh version 2 is enabled
```

2. Verify that the SSH keys are enabled:

```
show ssh key
```

Show example

```
(switch)# show ssh key

rsa Keys generated:Fri Jun 28 02:16:00 2024

ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQGDINrD52Q586wTGJjFABjBlFaA23EpDrZ2sDCew
l7nwl1oC6HBejxluIObAH8hrW8kR+gj0ZAfPpNeLGTg3APj/yIPTBoIZZxbWRShywAM5
PqyxWwRb7kp9Zt1YHzVuHYpSO82KUDowKrL6lox/YtpKoZUDZjrZjAp8hTv3JZsPgQ==

bitcount:1024
fingerprint:
SHA256:aHwhpzo7+YCDsrp3isJv2uVGz+mjMMokqdMeXVVXfdo

could not retrieve dsa key information

ecdsa Keys generated:Fri Jun 28 02:30:56 2024

ecdsa-sha2-nistp521
AAAAE2VjZHNhLXNoYTItbmlzdHA1MjEAAAABmlzdHA1MjEAAACFBABJ+ZX5SFKhS57e
vkE273e0VoqZi4/32dt+f14fBuKv80MjMsmLfjKtCWylwgVt1Zi+C5TIBbugpzez529z
kFSF0ADb8JaGCoaAYe2HvWR/f6QLbKbqVIewCdqWgxzrIY5BPP5GBdxQJMBiOwEdnHg1
u/9Pzh/Vz9cHDcCW9qGE780QHA==

bitcount:521
fingerprint:
SHA256:TFGe2hXn6QIpcs/vyHzftHJ7Dceg0vQaULYRALZeHwQ

(switch)# show feature | include scpServer
scpServer          1          enabled
(switch)# show feature | include ssh
sshServer          1          enabled
(switch)#
```



When enabling FIPS, you must change the bitcount to 256 on the switch using the command `ssh key ecdsa 256 force`. See [Configure network security using FIPS](#) for more details.

What's next?

[Configure switch health monitoring.](#)

Migrate switches

CN1610 switch migration

Migrate from CN1610 switches to Nexus 3232C switches workflow

Follow these workflow steps to migrate your CN1610 switches to Cisco Nexus 3232C switches.

1

Migration requirements

Review the example switch information for the migration process.

2

Prepare for migration

Prepare your CN1610 switches for migration to Nexus 3232C switches.

3

Replace cluster switch CL2

Replace cluster switch CL2 with the new Nexus 3232C switch C2.

4

Replace cluster switch CL1

Replace cluster switch CL1 with the new Nexus 3232C switch C1.

5

Complete your migration

Complete your migration to the new Nexus 3232C switches.

Migration requirements

Cisco Nexus 3232C switches can be used as cluster switches in your AFF or FAS cluster. Cluster switches allow you to build ONTAP clusters with more than two nodes.



The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

If necessary, refer to the following for more information:

- [NetApp CN1601 and CN1610 description page](#)
- [Cisco Ethernet Switch description page](#)
- [Hardware Universe](#)

CN1610 migrate requirements

The cluster switches support the following node connections:

- NetApp CN1610: ports 0/1 through 0/12 (10 GbE)
- Cisco Nexus 3232C: ports e1/1-30 (40 or 100 or 4x10GbE)

The cluster switches use the following inter-switch link (ISL) ports.

- NetApp CN1610: ports 0/13 through 0/16 (10 GbE)
- Cisco Nexus 3232C: ports 1/31-32 (100GbE)



You must use 4x10G breakout cables on the Cisco Nexus 3232C cluster switch.

The following table shows the cabling connections that are required at each stage as you make the transition from NetApp CN1610 switches to Cisco Nexus 3232C cluster switches:

Stage	Description	Required cables
Initial	CN1610 to CN1610 (SFP+ to SFP+)	4 SFP+ optical fiber or copper direct-attach cables
Transition	CN1610 to 3232C (QSFP to SFP+)	1 QSFP and 4 SFP+ optical fiber or copper breakout cables
Final	3232C to 3232C (QSFP to QSFP)	2 QSFP optical fiber or copper direct-attach cables

You must have downloaded the applicable reference configuration files (RCFs). The number of 10 GbE and 40/100 GbE ports are defined in the RCFs available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.

The ONTAP and NX-OS versions that are supported in this procedure are listed on the [Cisco Ethernet Switches](#) page.

The ONTAP and FASTPATH versions that are supported in this procedure are listed on the [NetApp CN1601 and CN1610 Switches](#) page.

About the examples used

The examples in the migration procedure use four nodes:

- Two nodes use four 10 GbE cluster interconnect ports: **e0a**, **e0b**, **e0c**, and **e0d**.
- The other two nodes use two 40 GbE cluster interconnect fiber cables: **e4a** and **e4e**.

The [Hardware Universe](#) has information about the cluster fiber cables for your platforms.

The examples in this procedure use the following switch and node nomenclature:

- The nodes are **n1**, **n2**, **n3**, and **n4**.
- The command outputs might vary depending on different releases of ONTAP software.
- The CN1610 switches to be replaced are **CL1** and **CL2**.
- The Nexus 3232C switches to replace the CN1610 switches are **C1** and **C2**.
- **n1_clus1** is the first cluster logical interface (LIF) that is connected to cluster switch 1 (CL1 or C1) for node **n1**.
- **n1_clus2** is the first cluster LIF that is connected to cluster switch 2 (CL2 or C2) for node **n1**.

- **n1_clus3** is the second LIF that is connected to cluster switch 2 (CL2 or C2) for node **n1**.
- **n1_clus4** is the second LIF that is connected to cluster switch 1 (CL1 or C1) for node **n1**.
- The number of 10 GbE and 40/100 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.

What's next?

[Prepare for migration.](#)

Prepare for migration from CN1610 switches 3232C switches

Follow these steps to prepare your CN1610 switch migration to a Cisco Nexus 3232C switch.

Steps

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=xh
```

x is the duration of the maintenance window in hours.



The message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Display information about the devices in your configuration:

```
network device-discovery show
```

Show example

The following example displays how many cluster interconnect interfaces have been configured in each node for each cluster interconnect switch:

```
cluster::> network device-discovery show
```

Node	Local Port	Discovered Device	Interface	Platform
n1	/cdp			
	e0a	CL1	0/1	CN1610
	e0b	CL2	0/1	CN1610
	e0c	CL2	0/2	CN1610
	e0d	CL1	0/2	CN1610
n2	/cdp			
	e0a	CL1	0/3	CN1610
	e0b	CL2	0/3	CN1610
	e0c	CL2	0/4	CN1610
	e0d	CL1	0/4	CN1610

```
8 entries were displayed.
```

3. Determine the administrative or operational status for each cluster interface.

a. Display the cluster network port attributes:

```
network port show -role cluster
```

Show example

```
cluster::*> network port show -role cluster
(network port show)

Node: n1

Port  IPspace  Broadcast  Link  MTU  Speed (Mbps)  Health  Ignore
Status  Domain    Status      Admin/Open  Status  Health
-----  -
e0a   cluster  cluster    up     9000  auto/10000    -
e0b   cluster  cluster    up     9000  auto/10000    -
e0c   cluster  cluster    up     9000  auto/10000    -    -
e0d   cluster  cluster    up     9000  auto/10000    -    -

Node: n2

Port  IPspace  Broadcast  Link  MTU  Speed (Mbps)  Health  Ignore
Status  Domain    Status      Admin/Open  Status  Health
-----  -
e0a   cluster  cluster    up     9000  auto/10000    -
e0b   cluster  cluster    up     9000  auto/10000    -
e0c   cluster  cluster    up     9000  auto/10000    -
e0d   cluster  cluster    up     9000  auto/10000    -

8 entries were displayed.
```

b. Display information about the logical interfaces:

```
network interface show -role cluster
```


Show example

```
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
true     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

8 entries were displayed.
```

c. Display information about the discovered cluster switches:

```
system cluster-switch show
```

Show example

The following example displays the cluster switches that are known to the cluster along with their management IP addresses:

```
cluster::> system cluster-switch show
Switch                               Type           Address
Model
-----
CL1                                   cluster-network 10.10.1.101
CN1610
    Serial Number: 01234567
    Is Monitored: true
    Reason:
    Software Version: 1.2.0.7
    Version Source: ISDP
CL2                                   cluster-network 10.10.1.102
CN1610
    Serial Number: 01234568
    Is Monitored: true
    Reason:
    Software Version: 1.2.0.7
    Version Source: ISDP

2 entries displayed.
```

4. Verify that the appropriate RCF and image are installed on the new 3232C switches as necessary for your requirements, and make any essential site customizations.

You should prepare both switches at this time. If you need to upgrade the RCF and image, you must complete the following procedure:

- a. See the [Cisco Ethernet Switch](#) page on the NetApp Support Site.
 - b. Note your switch and the required software versions in the table on that page.
 - c. Download the appropriate version of the RCF.
 - d. Select **CONTINUE** on the **Description** page, accept the license agreement, and then follow the instructions on the **Download** page to download the RCF.
 - e. Download the appropriate version of the image software at [Cisco® Cluster and Management Network Switch Reference Configuration File Download](#).
5. Migrate the LIFs associated with the second CN1610 switch that you plan to replace:

```
network interface migrate -vserver vserver-name -lif lif-name -source-node
source-node-name destination-node destination-node-name -destination-port
destination-port-name
```

Show example

You must migrate each LIF 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
```

6. Verify the cluster's health:

```
network interface show -role cluster
```

Show example

```
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
true    n1_clus1  up/up      10.10.0.1/24  n1       e0a
false  n1_clus2  up/up      10.10.0.2/24  n1       e0a
false  n1_clus3  up/up      10.10.0.3/24  n1       e0d
true   n1_clus4  up/up      10.10.0.4/24  n1       e0d
true   n2_clus1  up/up      10.10.0.5/24  n2       e0a
false  n2_clus2  up/up      10.10.0.6/24  n2       e0a
false  n2_clus3  up/up      10.10.0.7/24  n2       e0d
true   n2_clus4  up/up      10.10.0.8/24  n2       e0d

8 entries were displayed.
```

What's next?

[Replace CN1610 switch CL2 with 3232C switch C2.](#)

Replace CN1610 switch CL2 with 3232C switch C2

Follow these steps to replace your CN1610 (CL2) switch with a Nexus 3232C (C2) switch.

Steps

1. Shut down the cluster interconnect ports that are physically connected to switch CL2:

```
network port modify -node node-name -port port-name -up-admin false
```

Show example

The following example shows the four cluster interconnect ports being shut down for node n1 and node n2:

```
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
```

2. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
n1		
3/5/2022 19:21:18 -06:00	n1_clus2	n2-clus1
none		
3/5/2022 19:21:20 -06:00	n1_clus2	n2_clus2
none		
n2		
3/5/2022 19:21:18 -06:00	n2_clus2	n1_clus1
none		
3/5/2022 19:21:20 -06:00	n2_clus2	n1_clus2
none		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is 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
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
Cluster Vserver Id = 4294967293 Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 16 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 10.10.0.2 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 10.10.0.3 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 10.10.0.4 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

Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
4 paths up, 0 paths down (udp check)

```

3. Shut down the ISL ports 13 through 16 on the active CN1610 switch CL1 using the appropriate command.

For more information on Cisco commands, see the guides listed in the [Cisco Nexus 3000 Series NX-OS](#)

Command References.

Show example

The following example shows ISL ports 13 through 16 being shut down on the CN1610 switch CL1:

```
(CL1) # configure
(CL1) (Config) # interface 0/13-0/16
(CL1) (Interface 0/13-0/16) # shutdown
(CL1) (Interface 0/13-0/16) # exit
(CL1) (Config) # exit
(CL1) #
```

4. Build a temporary ISL between CL1 and C2:

For more information on Cisco commands, see the guides listed in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Show example

The following example shows a temporary ISL being built between CL1 (ports 13-16) and C2 (ports e1/24/1-4) using the Cisco switchport mode trunk command:

```
C2# configure
C2(config) # interface port-channel 2
C2(config-if) # switchport mode trunk
C2(config-if) # spanning-tree port type network
C2(config-if) # mtu 9216
C2(config-if) # interface breakout module 1 port 24 map 10g-4x
C2(config) # interface e1/24/1-4
C2(config-if-range) # switchport mode trunk
C2(config-if-range) # mtu 9216
C2(config-if-range) # channel-group 2 mode active
C2(config-if-range) # exit
C2(config-if) # exit
```

5. Remove the cables that are attached to the CN1610 switch CL2 on all the nodes.

Using supported cabling, you must reconnect the disconnected ports on all the nodes to the Nexus 3232C switch C2.

6. Remove four ISL cables from ports 13 to 16 on the CN1610 switch CL1.

You must attach the appropriate Cisco QSFP28 to SFP+ breakout cables connecting port 1/24 on the new Cisco 3232C switch C2 to ports 13 to 16 on the existing CN1610 switch CL1.



When reconnecting any cables to the new Cisco 3232C switch, the cables used must be either optical fiber or Cisco twinax cables.

7. Make the ISL dynamic by configuring the ISL interface 3/1 on the active CN1610 switch to disable the static mode.

This configuration matches with the ISL configuration on the 3232C switch C2 when the ISLs are brought up on both switches.

For more information on Cisco commands, see the guides listed in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Show example

The following example shows the ISL interface 3/1 being configured to make the ISL dynamic:

```
(CL1) # configure
(CL1) (Config) # interface 3/1
(CL1) (Interface 3/1) # no port-channel static
(CL1) (Interface 3/1) # exit
(CL1) (Config) # exit
(CL1) #
```

8. Bring up ISLs 13 through 16 on the active CN1610 switch CL1.

For more information on Cisco commands, see the guides listed in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Show example

The following example shows ISL ports 13 through 16 being brought up on the port-channel interface 3/1:

```
(CL1) # configure
(CL1) (Config) # interface 0/13-0/16,3/1
(CL1) (Interface 0/13-0/16,3/1) # no shutdown
(CL1) (Interface 0/13-0/16,3/1) # exit
(CL1) (Config) # exit
(CL1) #
```

9. Verify that the ISLs are up on the CN1610 switch CL1.

The "Link State" should be Up, "Type" should be Dynamic, and the "Port Active" column should be True for ports 0/13 to 0/16.

Show example

The following example shows the ISLs being verified as up on the CN1610 switch CL1:

```
(CL1)# show port-channel 3/1
Local Interface..... 3/1
Channel Name..... ISL-LAG
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr      Device/      Port      Port
Ports   Timeout     Speed     Active
-----
0/13    actor/long   10 Gb Full  True
        partner/long
0/14    actor/long   10 Gb Full  True
        partner/long
0/15    actor/long   10 Gb Full  True
        partner/long
0/16    actor/long   10 Gb Full  True
        partner/long
```

10. Verify that the ISLs are up on the 3232C switch 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](#).

Ports Eth1/24/1 through Eth1/24/4 should indicate (P), meaning that all four ISL ports are up in the port channel. Eth1/31 and Eth1/32 should indicate (D) as they are not connected.

Show example

The following example shows the ISLs being verified as up on the 3232C switch C2:

```
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 (D)  Eth1/32 (D)
2      Po2 (SU)       Eth       LACP      Eth1/24/1 (P) Eth1/24/2 (P)
Eth1/24/3 (P)
                                   Eth1/24/4 (P)
```

11. Bring up all of the cluster interconnect ports that are connected to the 3232C switch C2 on all of the nodes:

```
network port modify -node node-name -port port-name -up-admin true
```

Show example

The following example shows how to bring up the cluster interconnect ports connected to the 3232C switch C2:

```
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
```

12. Revert all of the migrated cluster interconnect LIFs that are connected to C2 on all of the nodes:

```
network interface revert -vserver cluster -lif lif-name
```

Show example

```
cluster::*> network interface revert -vserver cluster -lif n1_clus2
cluster::*> network interface revert -vserver cluster -lif n1_clus3
cluster::*> network interface revert -vserver cluster -lif n2_clus2
cluster::*> network interface revert -vserver cluster -lif n2_clus3
```

13. Verify that all of the cluster interconnect ports are reverted to their home ports:

```
network interface show -role cluster
```

Show example

The following example shows that the LIFs on clus2 are reverted to their home ports; the LIFs are successfully reverted if the ports in the "Current Port" column have a status of `true` in the "Is Home" column. If the "Is Home" value is `false`, then the LIF is not 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
true   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

8 entries were displayed.
```

14. Verify that all of the cluster ports are connected:

```
network port show -role cluster
```

Show example

The following example shows the output verifying all of the cluster interconnects are up:

```
cluster::*> network port show -role cluster
(network port show)

Node: n1

Port  IPspace  Broadcast  Link  MTU  Speed (Mbps)  Health  Ignore
Status  Domain    Link       MTU   Admin/Open   Status   Health
-----  -
e0a   cluster  cluster    up    9000  auto/10000    -
e0b   cluster  cluster    up    9000  auto/10000    -
e0c   cluster  cluster    up    9000  auto/10000    -
e0d   cluster  cluster    up    9000  auto/10000    -

Node: n2

Port  IPspace  Broadcast  Link  MTU  Speed (Mbps)  Health  Ignore
Status  Domain    Link       MTU   Admin/Open   Status   Health
-----  -
e0a   cluster  cluster    up    9000  auto/10000    -
e0b   cluster  cluster    up    9000  auto/10000    -
e0c   cluster  cluster    up    9000  auto/10000    -
e0d   cluster  cluster    up    9000  auto/10000    -

8 entries were displayed.
```

15. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
n1		
3/5/2022 19:21:18 -06:00	n1_clus2	n2-clus1
none		
3/5/2022 19:21:20 -06:00	n1_clus2	n2_clus2
none		
n2		
3/5/2022 19:21:18 -06:00	n2_clus2	n1_clus1
none		
3/5/2022 19:21:20 -06:00	n2_clus2	n1_clus2
none		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is 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
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
Cluster Vserver Id = 4294967293 Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 16 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 10.10.0.2 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 10.10.0.3 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 10.10.0.4 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

Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
4 paths up, 0 paths down (udp check)

```

16. Migrate the LIFs that are associated with the first CN1610 switch CL1:

```

network interface migrate -vserver cluster -lif lif-name -source-node node-name

```

Show example

You must migrate each cluster LIF individually to the appropriate cluster ports hosted on cluster switch C2 as shown in the following example:

```
cluster::*> network interface migrate -vserver cluster -lif n1_clus1
             -source-node n1
             -destination-node n1 -destination-port e0b
cluster::*> network interface migrate -vserver cluster -lif n1_clus4
             -source-node n1
             -destination-node n1 -destination-port e0c
cluster::*> network interface migrate -vserver cluster -lif n2_clus1
             -source-node n2
             -destination-node n2 -destination-port e0b
cluster::*> network interface migrate -vserver cluster -lif n2_clus4
             -source-node n2
             -destination-node n2 -destination-port e0c
```

What's next?

[Replace CN1610 switch CL1 with 3232C switch C1.](#)

Replace CN1610 switch CL1 with 3232C switch C1

Follow these steps to replace your CN1610 (CL1) switch with a Nexus 3232C (C1) switch.

Steps

1. Verify the cluster's status:

```
network interface show -role cluster
```


Show example

The following example shows that the required cluster LIFs have been migrated to the appropriate cluster ports hosted on cluster switch C2:

```
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
false  n1_clus1  up/up      10.10.0.1/24  n1       e0b
true   n1_clus2  up/up      10.10.0.2/24  n1       e0b
true   n1_clus3  up/up      10.10.0.3/24  n1       e0c
false  n1_clus4  up/up      10.10.0.4/24  n1       e0c
false  n2_clus1  up/up      10.10.0.5/24  n2       e0b
true   n2_clus2  up/up      10.10.0.6/24  n2       e0b
true   n2_clus3  up/up      10.10.0.7/24  n2       e0c
false  n2_clus4  up/up      10.10.0.8/24  n2       e0c

8 entries were displayed.
```

2. Shut down the node ports that are connected to CL1 on all of the nodes:

```
network port modify -node node-name -port port-name -up-admin false
```

Show example

The following example shows specific ports being shut down on nodes n1 and n2:

```
cluster::*> network port modify -node n1 -port e0a -up-admin false
cluster::*> network port modify -node n1 -port e0d -up-admin false
cluster::*> network port modify -node n2 -port e0a -up-admin false
cluster::*> network port modify -node n2 -port e0d -up-admin false
```

3. Shut down the ISL ports 24, 31, and 32 on the active 3232C switch C2.

For more information on Cisco commands, see the guides listed in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Show example

The following example shows ISLs 24, 31, and 32 being shut down on the active 3232C switch C2:

```
C2# configure
C2(config)# interface ethernet 1/24/1-4
C2(config-if-range)# shutdown
C2(config-if-range)# exit
C2(config)# interface ethernet 1/31-32
C2(config-if-range)# shutdown
C2(config-if-range)# exit
C2(config)# exit
C2#
```

4. Remove the cables that are attached to the CN1610 switch CL1 on all of the nodes.

Using the appropriate cabling, you must reconnect the disconnected ports on all the nodes to the Nexus 3232C switch C1.

5. Remove the QSFP28 cables from Nexus 3232C C2 port e1/24.

You must connect ports e1/31 and e1/32 on C1 to ports e1/31 and e1/32 on C2 using supported Cisco QSFP28 optical fiber or direct-attach cables.

6. Restore the configuration on port 24 and remove the temporary port-channel 2 on C2:

For more information on Cisco commands, see the guides listed in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Show example

The following example shows the running-configuration file being copied to the startup-configuration file:

```
C2# configure
C2(config)# no interface breakout module 1 port 24 map 10g-4x
C2(config)# no interface port-channel 2
C2(config-if)# interface e1/24
C2(config-if)# description 100GbE/40GbE Node Port
C2(config-if)# spanning-tree port type edge
Edge port type (portfast) should only be enabled on ports connected
to a single
host. Connecting hubs, concentrators, switches, bridges, etc... to
this
interface when edge port type (portfast) is enabled, can cause
temporary bridging loops.
Use with CAUTION

Edge Port Type (Portfast) has been configured on Ethernet 1/24 but
will only
have effect when the interface is in a non-trunking mode.

C2(config-if)# spanning-tree bpduguard enable
C2(config-if)# mtu 9216
C2(config-if-range)# exit
C2(config)# exit
C2# copy running-config startup-config
[] 100%
Copy Complete.
```

7. Bring up ISL ports 31 and 32 on C2, the active 3232C switch.

For more information on Cisco commands, see the guides listed in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Show example

The following example shows ISLs 31 and 32 being brought upon the 3232C switch C2:

```
C2# configure
C2(config)# interface ethernet 1/31-32
C2(config-if-range)# no shutdown
C2(config-if-range)# exit
C2(config)# exit
C2# copy running-config startup-config
[] 100%
Copy Complete.
```

8. Verify that the ISL connections are up on the 3232C switch C2.

For more information on Cisco commands, see the guides listed in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Show example

The following example shows the ISL connections being verified. Ports Eth1/31 and Eth1/32 indicate (P), meaning that both the ISL ports are up in the port-channel:

```
C1# show port-channel summary
Flags:  D - Down          P - Up in port-channel (members)
        I - Individual   H - Hot-standby (LACP only)
        s - Suspended    r - Module-removed
        S - Switched     R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met

-----
-----
Group Port-          Type      Protocol  Member Ports
Channel
-----
-----
1      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)
```

9. Bring up all of the cluster interconnect ports connected to the new 3232C switch C1 on all of the nodes:

```
network port modify -node node-name -port port-name -up-admin true
```

Show example

The following example shows all of the cluster interconnect ports connected to the new 3232C switch C1 being brought up:

```
cluster::*> network port modify -node n1 -port e0a -up-admin true
cluster::*> network port modify -node n1 -port e0d -up-admin true
cluster::*> network port modify -node n2 -port e0a -up-admin true
cluster::*> network port modify -node n2 -port e0d -up-admin true
```

10. Verify the status of the cluster node port:

```
network port show -role cluster
```

Show example

The following example shows output that verifies that the cluster interconnect ports on nodes n1 and n2 on the new 3232C switch C1 are up:

```
cluster::*> network port show -role cluster
(network port show)

Node: n1

Port  IPspace  Broadcast  Link  MTU  Speed (Mbps)  Health  Ignore
Status  Domain    Status      Admin/Open  Status  Health
-----  -----  -----  -----  -----  -----  -----  -----
e0a   cluster  cluster    up     9000  auto/10000    -        -
e0b   cluster  cluster    up     9000  auto/10000    -        -
e0c   cluster  cluster    up     9000  auto/10000    -        -
e0d   cluster  cluster    up     9000  auto/10000    -        -

Node: n2

Port  IPspace  Broadcast  Link  MTU  Speed (Mbps)  Health  Ignore
Status  Domain    Status      Admin/Open  Status  Health
-----  -----  -----  -----  -----  -----  -----  -----
e0a   cluster  cluster    up     9000  auto/10000    -        -
e0b   cluster  cluster    up     9000  auto/10000    -        -
e0c   cluster  cluster    up     9000  auto/10000    -        -
e0d   cluster  cluster    up     9000  auto/10000    -        -

8 entries were displayed.
```

What's next?

[Complete your migration.](#)

Complete your migration from CN1610 switches to 3232C switches

Complete the following steps to finalize the CN1610 switches migration to Nexus 3232C switches.

Steps

1. Revert all of the migrated cluster interconnect LIFs that were originally connected to C1 on all of the nodes:

```
network interface revert -server cluster -lif lif-name
```

Show example

You must migrate each LIF individually as shown in the following example:

```
cluster::*> network interface revert -vserver cluster -lif n1_clus1
cluster::*> network interface revert -vserver cluster -lif n1_clus4
cluster::*> network interface revert -vserver cluster -lif n2_clus1
cluster::*> network interface revert -vserver cluster -lif n2_clus4
```

2. Verify that the interface is now home:

```
network interface show -role cluster
```


Show example

The following example shows the status of cluster interconnect interfaces is up and "Is Home" for nodes n1 and n2:

```
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
true      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

8 entries were displayed.
```

3. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
n1		
3/5/2022 19:21:18 -06:00	n1_clus2	n2-clus1
none		
3/5/2022 19:21:20 -06:00	n1_clus2	n2_clus2
none		
n2		
3/5/2022 19:21:18 -06:00	n2_clus2	n1_clus1
none		
3/5/2022 19:21:20 -06:00	n2_clus2	n1_clus2
none		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is 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
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
Cluster Vserver Id = 4294967293 Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 16 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 10.10.0.2 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 10.10.0.3 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 10.10.0.4 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

Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
3 paths up, 0 paths down (udp check)

```

4. Expand the cluster by adding nodes to the Nexus 3232C cluster switches.
5. Display the information about the devices in your configuration:
 - ° network device-discovery show

- ° network port show -role cluster
- ° network interface show -role cluster
- ° system cluster-switch show

Show example

The following examples show nodes n3 and n4 with 40 GbE cluster ports connected to ports e1/7 and e1/8, respectively, on both the Nexus 3232C cluster switches. Both nodes are joined to the cluster. The 40 GbE cluster interconnect ports used are e4a and e4e.

```
cluster::*> network device-discovery show
```

Node	Local Port	Discovered 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	Port	IPspace	Domain	Link	MTU	Speed (Mbps)	Health	Status
	e0a	cluster	cluster	up	9000	auto/10000	-	-
	e0b	cluster	cluster	up	9000	auto/10000	-	-
	e0c	cluster	cluster	up	9000	auto/10000	-	-
	e0d	cluster	cluster	up	9000	auto/10000	-	-

Node: n2

		Broadcast			Speed (Mbps)	Health	
Ignore							
Port	IPspace	Domain	Link	MTU	Admin/Open	Status	
Health	Status						
-----	-----	-----	-----	-----	-----	-----	-----
e0a	cluster	cluster	up	9000	auto/10000	-	
e0b	cluster	cluster	up	9000	auto/10000	-	
e0c	cluster	cluster	up	9000	auto/10000	-	
e0d	cluster	cluster	up	9000	auto/10000	-	-

Node: n3

		Broadcast			Speed (Mbps)	Health	
Ignore							
Port	IPspace	Domain	Link	MTU	Admin/Open	Status	
Health	Status						
-----	-----	-----	-----	-----	-----	-----	-----
e4a	cluster	cluster	up	9000	auto/40000	-	
e4e	cluster	cluster	up	9000	auto/40000	-	-

Node: n4

		Broadcast			Speed (Mbps)	Health	
Ignore							
Port	IPspace	Domain	Link	MTU	Admin/Open	Status	
Health	Status						
-----	-----	-----	-----	-----	-----	-----	-----
e4a	cluster	cluster	up	9000	auto/40000	-	
e4e	cluster	cluster	up	9000	auto/40000	-	

12 entries were displayed.

cluster::*> **network interface show -role cluster**

(network interface show)

	Logical	Status	Network	Current	Current
Is					
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	-----
Cluster					
true	n1_clus1	up/up	10.10.0.1/24	n1	e0a
	n1_clus2	up/up	10.10.0.2/24	n1	e0b


```

    Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                        7.0(3)I6(1)
    Version Source: CDP
CL1                      cluster-network  10.10.1.101    CN1610

    Serial Number: 01234567
    Is Monitored: true
    Reason:
    Software Version: 1.2.0.7
    Version Source: ISDP
CL2                      cluster-network  10.10.1.102
CN1610

    Serial Number: 01234568
    Is Monitored: true
    Reason:
    Software Version: 1.2.0.7
    Version Source: ISDP 4 entries were displayed.

```

6. Remove the replaced CN1610 switches if they are not automatically removed:

```
system cluster-switch delete -device switch-name
```

Show example

You must delete both devices individually as shown in the following example:

```

cluster::> system cluster-switch delete -device CL1
cluster::> system cluster-switch delete -device CL2

```

7. Verify that the proper cluster switches are monitored:

```
system cluster-switch show
```


Show example

The following example shows cluster switches C1 and C2 are being monitored:

```
cluster::> system cluster-switch show
```

Switch Model	Type	Address
C1 NX3232C	cluster-network	10.10.1.103
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 NX3232C	cluster-network	10.10.1.104
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.

8. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next?

[Configure switch health monitoring.](#)

Nexus 5596 switch migration

Migrate from Nexus 5596 switches to Nexus 3232C switches workflow

Follow these workflow steps to migrate your Cisco Nexus 5596 switches to Cisco Nexus 3232C switches.

1

Migration requirements

Review the example switch information for the migration process.

2

Prepare for migration

Prepare your Nexus 5596 switches for migration to Nexus 3232C switches.

3

Configure your ports

Configure your ports for migration to the new Nexus 3232C switches.

4

Complete your migration

Complete you migration to the new Nexus 3232C switches.

Migration requirements

Cisco Nexus 3232C switches can be used as cluster switches in your AFF or FAS cluster. Cluster switches allow you to build ONTAP clusters with more than two nodes.



The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

For more information, see:

- [Cisco Ethernet Switch description page](#)
- [Hardware Universe](#)

Cisco Nexus 5596 requirements

The cluster switches use the following ports for connections to nodes:

- Nexus 5596: ports e1/1-40 (10 GbE)
- Nexus 3232C: ports e1/1-30 (10/40/100 GbE)

The cluster switches use the following Inter-Switch Link (ISL) ports:

- Nexus 5596: ports e1/41-48 (10 GbE)
- Nexus 3232C: ports e1/31-32 (40/100 GbE)

The [Hardware Universe](#) contains information about supported cabling to Nexus 3232C switches:

- Nodes with 10 GbE cluster connections require QSFP to SFP+ optical fiber breakout cables or QSFP to SFP+ copper breakout cables.
- Nodes with 40/100 GbE cluster connections require supported QSFP/QSFP28 optical modules with fiber cables or QSFP/QSFP28 copper direct-attach cables.

The cluster switches use the appropriate ISL cabling:

- Beginning: Nexus 5596 (SFP+ to SFP+)
 - 8x SFP+ fiber or copper direct-attach cables
- Interim: Nexus 5596 to Nexus 3232C (QSFP to 4xSFP+ break-out)
 - 1x QSFP to SFP+ fiber break-out or copper break-out cables
- Final: Nexus 3232C to Nexus 3232C (QSFP28 to QSFP28)
 - 2x QSFP28 fiber or copper direct-attach cables
- On Nexus 3232C switches, you can operate QSFP/QSFP28 ports in either 40/100 Gigabit Ethernet or 4 x10 Gigabit Ethernet modes.

By default, there are 32 ports in the 40/100 Gigabit Ethernet mode. These 40 Gigabit Ethernet ports are numbered in a 2-tuple naming convention. For example, the second 40 Gigabit Ethernet port is numbered as 1/2.

The process of changing the configuration from 40 Gigabit Ethernet to 10 Gigabit Ethernet is called *breakout* and the process of changing the configuration from 10 Gigabit Ethernet to 40 Gigabit Ethernet is called *breakin*.

When you break out a 40/100 Gigabit Ethernet port into 10 Gigabit Ethernet ports, the resulting ports are numbered using a 3-tuple naming convention. For example, the break-out ports of the second 40/100 Gigabit Ethernet port are numbered as 1/2/1, 1/2/2, 1/2/3, and 1/2/4.

- On the left side of Nexus 3232C switches are 2 SFP+ ports, called 1/33 and 1/34.
- You have configured some of the ports on Nexus 3232C switches to run at 10 GbE or 40/100 GbE.



You can break out the first six ports into 4x10 GbE mode by using the `interface breakout module 1 port 1-6 map 10g-4x` command. Similarly, you can regroup the first six QSFP+ ports from breakout configuration by using the `no interface breakout module 1 port 1-6 map 10g-4x` command.

- You have done the planning, migration, and read the required documentation on 10 GbE and 40/100 GbE connectivity from nodes to Nexus 3232C cluster switches.
- The ONTAP and NX-OS versions supported in this procedure are on the [Cisco Ethernet Switches page](#).

About the examples used

The examples in this procedure describe replacing Cisco Nexus 5596 switches with Cisco Nexus 3232C switches. You can use these steps (with modifications) for other older Cisco switches (for example, 3132Q-V).

The procedure also uses the following switch and node nomenclature:

- The command outputs might vary depending on different releases of ONTAP.
- The Nexus 5596 switches to be replaced are **CL1** and **CL2**.
- The Nexus 3232C switches to replace the Nexus 5596 switches are **C1** and **C2**.
- **n1_clus1** is the first cluster logical interface (LIF) connected to cluster switch 1 (CL1 or C1) for node **n1**.
- **n1_clus2** is the first cluster LIF connected to cluster switch 2 (CL2 or C2) for node **n1**.
- **n1_clus3** is the second LIF connected to cluster switch 2 (CL2 or C2) for node **n1**.

- **n1_clus4** is the second LIF connected to cluster switch 1 (CL1 or C1) for node **n1**.
- The number of 10 GbE and 40/100 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.
- The nodes are **n1**, **n2**, **n3**, and **n4**.

The examples in this procedure use four nodes:

- Two nodes use four 10 GbE cluster interconnect ports: **e0a**, **e0b**, **e0c**, and **e0d**.
- The other two nodes use two 40 GbE cluster interconnect ports: **e4a** and **e4e**.

The [Hardware Universe](#) lists the actual cluster ports on your platforms.

Scenarios covered

This procedure covers the following scenarios:

- The cluster starts with two nodes connected and functioning in a two Nexus 5596 cluster switches.
- The cluster switch CL2 to be replaced by C2 (steps 1 to 19):
 - Traffic on all cluster ports and LIFs on all nodes connected to CL2 are migrated onto the first cluster ports and LIFs connected to CL1.
 - Disconnect cabling from all cluster ports on all nodes connected to CL2, and then use supported break-out cabling to reconnect the ports to new cluster switch C2.
 - Disconnect cabling between ISL ports between CL1 and CL2, and then use supported break-out cabling to reconnect the ports from CL1 to C2.
 - Traffic on all cluster ports and LIFs connected to C2 on all nodes is reverted.
- The cluster switch CL2 to be replaced by C2.
 - Traffic on all cluster ports or LIFs on all nodes connected to CL1 are migrated onto the second cluster ports or LIFs connected to C2.
 - Disconnect cabling from all cluster port on all nodes connected to CL1 and reconnect, using supported break-out cabling, to new cluster switch C1.
 - Disconnect cabling between ISL ports between CL1 and C2, and reconnect using supported cabling, from C1 to C2.
 - Traffic on all cluster ports or LIFs connected to C1 on all nodes is reverted.
- Two FAS9000 nodes have been added to cluster with examples showing cluster details.

What's next?

[Prepare for migration.](#)

Prepare for migration from Nexus 5596 switches to Nexus 3232C switches

Follow these steps to prepare your Cisco Nexus 5596 switches for migration to Cisco Nexus 3232C switches.

Steps

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all - message MAINT=xh
```

x is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Display information about the devices in your configuration:

```
network device-discovery show
```

Show example

The following example shows how many cluster interconnect interfaces have been configured in each node for each cluster interconnect switch:

```
cluster::> network device-discovery show
```

Node	Local Port	Discovered Device	Interface	Platform
n1	/cdp			
	e0a	CL1	Ethernet1/1	N5K-C5596UP
	e0b	CL2	Ethernet1/1	N5K-C5596UP
	e0c	CL2	Ethernet1/2	N5K-C5596UP
	e0d	CL1	Ethernet1/2	N5K-C5596UP
n2	/cdp			
	e0a	CL1	Ethernet1/3	N5K-C5596UP
	e0b	CL2	Ethernet1/3	N5K-C5596UP
	e0c	CL2	Ethernet1/4	N5K-C5596UP
	e0d	CL1	Ethernet1/4	N5K-C5596UP

8 entries were displayed.

3. Determine the administrative or operational status for each cluster interface.

a. Display the network port attributes:

```
network port show -role cluster
```

Show example

The following example displays the network port attributes on nodes n1 and n2:

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper  Speed (Mbps)
Status      Status
-----
-----
e0a         Cluster    Cluster          up   9000 auto/10000  -
-
e0b         Cluster    Cluster          up   9000 auto/10000  -
-
e0c         Cluster    Cluster          up   9000 auto/10000  -
-
e0d         Cluster    Cluster          up   9000 auto/10000  -
-

Node: n2

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper  Speed (Mbps)
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  -
-

8 entries were displayed.
```

b. Display information about the logical interfaces:

```
network interface show -role cluster
```

Show example

The following example displays the general information about all of the LIFs on the cluster, including their current ports:

```
cluster::*> network interface show -role cluster
(network interface show)
Current Is Logical Status Network Current
Vserver Is Interface Admin/Oper Address/Mask Node
Port Home
-----
Cluster
e0a true n1_clus1 up/up 10.10.0.1/24 n1
e0b true n1_clus2 up/up 10.10.0.2/24 n1
e0c true n1_clus3 up/up 10.10.0.3/24 n1
e0d true n1_clus4 up/up 10.10.0.4/24 n1
e0a true n2_clus1 up/up 10.10.0.5/24 n2
e0b true n2_clus2 up/up 10.10.0.6/24 n2
e0c true n2_clus3 up/up 10.10.0.7/24 n2
e0d true n2_clus4 up/up 10.10.0.8/24 n2
8 entries were displayed.
```

c. Display information about the discovered cluster switches:

```
system cluster-switch show
```

Show example

The following example shows the active cluster switches:

```
cluster::*> system cluster-switch show

Switch                               Type                               Address
Model
-----
CL1                                   cluster-network                    10.10.1.101
NX5596
  Serial Number: 01234567
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                               7.1(1)N1(1)
  Version Source: CDP
CL2                                   cluster-network                    10.10.1.102
NX5596
  Serial Number: 01234568
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                               7.1(1)N1(1)
  Version Source: CDP

2 entries were displayed.
```

4. Verify that the appropriate RCF and image are installed on the new 3232C switches as necessary for your requirements, and make the essential site customizations, such as users and passwords, network addresses, and other customizations.



You must prepare both switches at this time.

If you need to upgrade the RCF and image, you must complete the following steps:

- a. Go to the *Cisco Ethernet Switches* page on the NetApp Support Site.

[Cisco Ethernet Switches](#)

- b. Note your switch and the required software versions in the table on that page.
- c. Download the appropriate version of the RCF.
- d. Select **CONTINUE** on the **Description** page, accept the license agreement, and then follow the

instructions on the **Download** page to download the RCF.

- e. Download the appropriate version of the image software.

See the *ONTAP 8.x or later Cluster and Management Network Switch Reference Configuration Files Download* page, and then select the appropriate version.

To find the correct version, see the *ONTAP 8.x or later Cluster Network Switch Download* page.

5. Migrate the LIFs associated with the second Nexus 5596 switch to be replaced:

```
network interface migrate -vserver vservice-name -lif lif-name -source-node
source-node-name - destination-node node-name -destination-port destination-
port-name
```

Show example

The following example shows the LIFs being migrated for nodes n1 and n2; LIF migration must be done on all of the nodes:

```
cluster::*> network interface migrate -vserver Cluster -lif n1_clus2
-source-node n1 -
destination-node n1 -destination-port e0a
cluster::*> network interface migrate -vserver Cluster -lif n1_clus3
-source-node n1 -
destination-node n1 -destination-port e0d
cluster::*> network interface migrate -vserver Cluster -lif n2_clus2
-source-node n2 -
destination-node n2 -destination-port e0a
cluster::*> network interface migrate -vserver Cluster -lif n2_clus3
-source-node n2 -
destination-node n2 -destination-port e0d
```

6. Verify the cluster's health:

```
network interface show -role cluster
```

Show example

The following example shows the current status of each 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
e0a      n1_clus1  up/up      10.10.0.1/24  n1
true
e0a      n1_clus2  up/up      10.10.0.2/24  n1
false
e0d      n1_clus3  up/up      10.10.0.3/24  n1
false
e0d      n1_clus4  up/up      10.10.0.4/24  n1
true
e0a      n2_clus1  up/up      10.10.0.5/24  n2
true
e0a      n2_clus2  up/up      10.10.0.6/24  n2
false
e0d      n2_clus3  up/up      10.10.0.7/24  n2
false
e0d      n2_clus4  up/up      10.10.0.8/24  n2
true
8 entries were displayed.
```

What's next?

[Configure your ports.](#)

Configure your ports for migration from Nexus 5596 switches to Nexus 3232C switches

Follow these steps to configure your ports for migration from the Nexus 5596 switches to the new Nexus 3232C switches.

Steps

1. Shut down the cluster interconnect ports that are physically connected to switch CL2:

```
network port modify -node node-name -port port-name -up-admin false
```

Show example

The following commands shut down the specified ports on n1 and n2, but the ports must be 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
```

2. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
n1		
3/5/2022 19:21:18 -06:00	n1_clus2	n2-clus1
none		
3/5/2022 19:21:20 -06:00	n1_clus2	n2_clus2
none		
n2		
3/5/2022 19:21:18 -06:00	n2_clus2	n1_clus1
none		
3/5/2022 19:21:20 -06:00	n2_clus2	n1_clus2
none		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is 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
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
Cluster Vserver Id = 4294967293 Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 16 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 10.10.0.2 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 10.10.0.3 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 10.10.0.4 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

Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
4 paths up, 0 paths down (udp check)

```

3. Shut down ISLs 41 through 48 on CL1, the active Nexus 5596 switch, using the Cisco `shutdown` command.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-](#)

Show example

The following example shows ISLs 41 through 48 being shut down on the Nexus 5596 switch CL1:

```
(CL1) # configure
(CL1) (Config) # interface e1/41-48
(CL1) (config-if-range) # shutdown
(CL1) (config-if-range) # exit
(CL1) (Config) # exit
(CL1) #
```

4. Build a temporary ISL between CL1 and C2 using the appropriate Cisco commands.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-OS Command References.](#)

Show example

The following example shows a temporary ISL being set up between CL1 and C2:

```
C2# configure
C2(config) # interface port-channel 2
C2(config-if) # switchport mode trunk
C2(config-if) # spanning-tree port type network
C2(config-if) # mtu 9216
C2(config-if) # interface breakout module 1 port 24 map 10g-4x
C2(config) # interface e1/24/1-4
C2(config-if-range) # switchport mode trunk
C2(config-if-range) # mtu 9216
C2(config-if-range) # channel-group 2 mode active
C2(config-if-range) # exit
C2(config-if) # exit
```

5. On all nodes, remove all cables attached to the Nexus 5596 switch CL2.

With supported cabling, reconnect disconnected ports on all nodes to the Nexus 3232C switch C2.

6. Remove all the cables from the Nexus 5596 switch CL2.

Attach the appropriate Cisco QSFP to SFP+ break-out cables connecting port 1/24 on the new Cisco 3232C switch, C2, to ports 45 to 48 on existing Nexus 5596, CL1.

7. Bring up ISLs ports 45 through 48 on the active Nexus 5596 switch CL1.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-](#)

Show example

The following example shows ISLs ports 45 through 48 being brought up:

```
(CL1) # configure
(CL1) (Config) # interface e1/45-48
(CL1) (config-if-range) # no shutdown
(CL1) (config-if-range) # exit
(CL1) (Config) # exit
(CL1) #
```

8. Verify that the ISLs are up on the Nexus 5596 switch CL1.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Show example

The following example shows Ports eth1/45 through eth1/48 indicating (P), meaning 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      Eth1/41 (D)  Eth1/42 (D)
Eth1/43 (D)
                                     Eth1/44 (D)  Eth1/45 (P)
Eth1/46 (P)
                                     Eth1/47 (P)  Eth1/48 (P)
```

9. Verify that interfaces eth1/45-48 already have `channel-group 1 mode active` in their running configuration.
10. On all nodes, bring up all the cluster interconnect ports connected to the 3232C switch C2:

```
network port modify -node node-name -port port-name -up-admin true
```

Show example

The following example shows the specified ports being brought up on nodes n1 and n2:

```
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
```

11. On all nodes, revert all of the migrated cluster interconnect LIFs connected to C2:

```
network interface revert -vserver Cluster -lif lif-name
```

Show example

The following example shows the migrated cluster LIFs being reverted to their home ports:

```
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
```

12. Verify all the cluster interconnect ports are now reverted to their home:

```
network interface show -role cluster
```


Show example

The following example shows that the LIFs on clus2 reverted to their home ports and shows that the LIFs are successfully reverted if the ports in the Current Port column have a status of `true` in the `Is Home` column. If the `Is Home` value is `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
e0a      n1_clus1   up/up    10.10.0.1/24   n1
      true
e0b      n1_clus2   up/up    10.10.0.2/24   n1
      true
e0c      n1_clus3   up/up    10.10.0.3/24   n1
      true
e0d      n1_clus4   up/up    10.10.0.4/24   n1
      true
e0a      n2_clus1   up/up    10.10.0.5/24   n2
      true
e0b      n2_clus2   up/up    10.10.0.6/24   n2
      true
e0c      n2_clus3   up/up    10.10.0.7/24   n2
      true
e0d      n2_clus4   up/up    10.10.0.8/24   n2
      true
8 entries were displayed.
```

13. Verify that the clustered ports are connected:

```
network port show -role cluster
```

Show example

The following example shows the result of the previous `network port modify` command, verifying that all the cluster interconnects are up:

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Speed(Mbps) Health
Status
-----
-----
e0a      Cluster      Cluster      up   9000  auto/10000  -
-
e0b      Cluster      Cluster      up   9000  auto/10000  -
-
e0c      Cluster      Cluster      up   9000  auto/10000  -
-
e0d      Cluster      Cluster      up   9000  auto/10000  -
-

Node: n2

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Speed(Mbps) Health
Status
-----
-----
e0a      Cluster      Cluster      up   9000  auto/10000  -
-
e0b      Cluster      Cluster      up   9000  auto/10000  -
-
e0c      Cluster      Cluster      up   9000  auto/10000  -
-
e0d      Cluster      Cluster      up   9000  auto/10000  -
-

8 entries were displayed.
```

14. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
n1		
3/5/2022 19:21:18 -06:00	n1_clus2	n2-clus1
none		
3/5/2022 19:21:20 -06:00	n1_clus2	n2_clus2
none		
n2		
3/5/2022 19:21:18 -06:00	n2_clus2	n1_clus1
none		
3/5/2022 19:21:20 -06:00	n2_clus2	n1_clus2
none		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is 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
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
Cluster Vserver Id = 4294967293 Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 16 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 10.10.0.2 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 10.10.0.3 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 10.10.0.4 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

Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
4 paths up, 0 paths down (udp check)

```

- On each node in the cluster, migrate the interfaces associated with the first Nexus 5596 switch, CL1, to be replaced:

```
network interface migrate -vserver vserver-name -lif lif-name -source-node
```

source-node-name

-destination-node destination-node-name -destination-port destination-port-name

Show example

The following example shows the ports or LIFs being migrated on nodes n1 and n2:

```
cluster::*> network interface migrate -vserver Cluster -lif n1_clus1
-source-node n1 -
destination-node n1 -destination-port e0b
cluster::*> network interface migrate -vserver Cluster -lif n1_clus4
-source-node n1 -
destination-node n1 -destination-port e0c
cluster::*> network interface migrate -vserver Cluster -lif n2_clus1
-source-node n2 -
destination-node n2 -destination-port e0b
cluster::*> network interface migrate -vserver Cluster -lif n2_clus4
-source-node n2 -
destination-node n2 -destination-port e0c
```

16. Verify the cluster's status:

```
network interface show
```

Show example

The following example shows that the required cluster LIFs have been migrated to appropriate cluster ports hosted on cluster switch, C2:

```
cluster::*> network interface show
```

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
e0b	n1_clus1	up/up	10.10.0.1/24	n1
e0b	false			
e0b	n1_clus2	up/up	10.10.0.2/24	n1
e0b	true			
e0c	n1_clus3	up/up	10.10.0.3/24	n1
e0c	true			
e0c	n1_clus4	up/up	10.10.0.4/24	n1
e0c	false			
e0b	n2_clus1	up/up	10.10.0.5/24	n2
e0b	false			
e0b	n2_clus2	up/up	10.10.0.6/24	n2
e0b	true			
e0c	n2_clus3	up/up	10.10.0.7/24	n2
e0c	true			
e0c	n2_clus4	up/up	10.10.0.8/24	n2
e0c	false			

8 entries were displayed.

17. On all the nodes, shut down the node ports that are connected to CL1:

```
network port modify -node node-name -port port-name -up-admin false
```

Show example

The following example shows the specified ports being shut down on nodes n1 and n2:

```
cluster::*> network port modify -node n1 -port e0a -up-admin false
cluster::*> network port modify -node n1 -port e0d -up-admin false
cluster::*> network port modify -node n2 -port e0a -up-admin false
cluster::*> network port modify -node n2 -port e0d -up-admin false
```

18. Shut down ISL 24, 31 and 32 on the active 3232C switch C2.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Show example

The following example shows ISLs being shutdown:

```
C2# configure
C2(Config)# interface e1/24/1-4
C2(config-if-range)# shutdown
C2(config-if-range)# exit
C2(config)# interface 1/31-32
C2(config-if-range)# shutdown
C2(config-if-range)# exit
C2(config-if)# exit
C2#
```

19. On all nodes, remove all cables attached to the Nexus 5596 switch CL1.

With supported cabling, reconnect disconnected ports on all nodes to the Nexus 3232C switch C1.

20. Remove the QSFP breakout cable from Nexus 3232C C2 ports e1/24.

Connect ports e1/31 and e1/32 on C1 to ports e1/31 and e1/32 on C2 using supported Cisco QSFP optical fiber or direct-attach cables.

21. Restore the configuration on port 24 and remove the temporary Port Channel 2 on C2.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Show example

The following example shows the configuration on port m24 being restored using the appropriate Cisco commands:

```
C2# configure
C2(config)# no interface breakout module 1 port 24 map 10g-4x
C2(config)# no interface port-channel 2
C2(config-if)# int e1/24
C2(config-if)# description 40GbE Node Port
C2(config-if)# spanning-tree port type edge
C2(config-if)# spanning-tree bpduguard enable
C2(config-if)# mtu 9216
C2(config-if-range)# exit
C2(config)# exit
C2# copy running-config startup-config
[] 100%
Copy Complete.
```

22. Bring up ISL ports 31 and 32 on C2, the active 3232C switch, by entering the following Cisco command: `no shutdown`

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Show example

The following example shows the Cisco commands `switchname configure` brought up on the 3232C switch C2:

```
C2# configure
C2(config)# interface ethernet 1/31-32
C2(config-if-range)# no shutdown
```

23. Verify that the ISL connections are up on the 3232C switch C2.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Ports eth1/31 and eth1/32 should indicate (P) meaning that both ISL ports up in the port-channel

Show example

```
C1# show port-channel summary
Flags: D - Down          P - Up in port-channel (members)
      I - Individual    H - Hot-standby (LACP only)
      s - Suspended     r - Module-removed
      S - Switched      R - Routed
      U - Up (port-channel)
      M - Not in use. Min-links not met

-----
-----
Group Port-          Type   Protocol  Member Ports
Channel
-----
-----
1      Po1(SU)        Eth    LACP      Eth1/31(P)  Eth1/32(P)
```

24. On all nodes, bring up all the cluster interconnect ports connected to the new 3232C switch C1:

```
network port modify
```

Show example

The following example shows all the cluster interconnect ports being brought up for n1 and n2 on the 3232C switch C1:

```
cluster::*> network port modify -node n1 -port e0a -up-admin true
cluster::*> network port modify -node n1 -port e0d -up-admin true
cluster::*> network port modify -node n2 -port e0a -up-admin true
cluster::*> network port modify -node n2 -port e0d -up-admin true
```

25. Verify the status of the cluster node port:

```
network port show
```

Show example

The following example shows verifies that all cluster interconnect ports on all nodes on the new 3232C switch C1 are up:

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e0a       Cluster      Cluster      up   9000  auto/10000  -
-
e0b       Cluster      Cluster      up   9000  auto/10000  -
-
e0c       Cluster      Cluster      up   9000  auto/10000  -
-
e0d       Cluster      Cluster      up   9000  auto/10000  -
-

Node: n2

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e0a       Cluster      Cluster      up   9000  auto/10000  -
-
e0b       Cluster      Cluster      up   9000  auto/10000  -
-
e0c       Cluster      Cluster      up   9000  auto/10000  -
-
e0d       Cluster      Cluster      up   9000  auto/10000  -
-

8 entries were displayed.
```

26. On all nodes, revert the specific cluster LIFs to their home ports:

```
network interface revert -server Cluster -lif lif-name
```

Show example

The following example shows the specific cluster LIFs being reverted to their home ports on nodes n1 and n2:

```
cluster::*> network interface revert -vserver Cluster -lif n1_clus1  
cluster::*> network interface revert -vserver Cluster -lif n1_clus4  
cluster::*> network interface revert -vserver Cluster -lif n2_clus1  
cluster::*> network interface revert -vserver Cluster -lif n2_clus4
```

27. Verify that the interface is home:

```
network interface show -role cluster
```

Show example

The following example shows the status of cluster interconnect interfaces are up and Is Home for n1 and n2:

```
cluster::*> network interface show -role cluster
(network interface show)
Current Is Logical Status Network Current
Vserver Interface Admin/Oper Address/Mask Node
Port Home
-----
Cluster
e0a true n1_clus1 up/up 10.10.0.1/24 n1
e0b true n1_clus2 up/up 10.10.0.2/24 n1
e0c true n1_clus3 up/up 10.10.0.3/24 n1
e0d true n1_clus4 up/up 10.10.0.4/24 n1
e0a true n2_clus1 up/up 10.10.0.5/24 n2
e0b true n2_clus2 up/up 10.10.0.6/24 n2
e0c true n2_clus3 up/up 10.10.0.7/24 n2
e0d true n2_clus4 up/up 10.10.0.8/24 n2
8 entries were displayed.
```

28. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
n1		
3/5/2022 19:21:18 -06:00	n1_clus2	n2-clus1
none		
3/5/2022 19:21:20 -06:00	n1_clus2	n2_clus2
none		
n2		
3/5/2022 19:21:18 -06:00	n2_clus2	n1_clus1
none		
3/5/2022 19:21:20 -06:00	n2_clus2	n1_clus2
none		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is 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
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
Cluster Vserver Id = 4294967293 Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 16 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 10.10.0.2 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 10.10.0.3 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 10.10.0.4 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

Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
4 paths up, 0 paths down (udp check)

```

29. Expand the cluster by adding nodes to the Nexus 3232C cluster switches.

The following examples show nodes n3 and n4 have 40 GbE cluster ports connected to ports e1/7 and e1/8 respectively on both the Nexus 3232C cluster switches, and both nodes have joined the cluster. The

40 GbE cluster interconnect ports used are e4a and e4e.

Display the information about the devices in your configuration:

- ° network device-discovery show
- ° network port show -role cluster
- ° network interface show -role cluster
- ° system cluster-switch show

Show example

```
cluster::> network device-discovery show
      Local   Discovered
Node   Port     Device           Interface         Platform
-----
n1     /cdp
      e0a     C1               Ethernet1/1/1     N3K-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

Health
Port      IPspace      Broadcast Domain Link MTU   Admin/Oper   Status
Status
-----
e0a      Cluster      Cluster      up   9000 auto/10000 -
-
e0b      Cluster      Cluster      up   9000 auto/10000 -
-
e0c      Cluster      Cluster      up   9000 auto/10000 -
-
e0d      Cluster      Cluster      up   9000 auto/10000 -
```

```

-
Node: n2

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Speed(Mbps) Health
Status
-----
-----
e0a      Cluster      Cluster      up    9000  auto/10000  -
-
e0b      Cluster      Cluster      up    9000  auto/10000  -
-
e0c      Cluster      Cluster      up    9000  auto/10000  -
-
e0d      Cluster      Cluster      up    9000  auto/10000  -
-

```

```

Node: n3

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Speed(Mbps) Health
Status
-----
-----
e4a      Cluster      Cluster      up    9000  auto/40000  -
-
e4e      Cluster      Cluster      up    9000  auto/40000  -
-

```

```

Node: n4

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Speed(Mbps) Health
Status
-----
-----
e4a      Cluster      Cluster      up    9000  auto/40000  -
-
e4e      Cluster      Cluster      up    9000  auto/40000  -

```

-
12 entries were displayed.

+

```
cluster::*> network interface show -role cluster
(network interface show)
      Logical      Status      Network      Current
Current Is
Vserver  Interface  Admin/Oper  Address/Mask  Node
Port    Home
-----
Cluster
e0a      true      n1_clus1   up/up         10.10.0.1/24  n1
e0b      true      n1_clus2   up/up         10.10.0.2/24  n1
e0c      true      n1_clus3   up/up         10.10.0.3/24  n1
e0d      true      n1_clus4   up/up         10.10.0.4/24  n1
e0a      true      n2_clus1   up/up         10.10.0.5/24  n2
e0b      true      n2_clus2   up/up         10.10.0.6/24  n2
e0c      true      n2_clus3   up/up         10.10.0.7/24  n2
e0d      true      n2_clus4   up/up         10.10.0.8/24  n2
e4a      true      n3_clus1   up/up         10.10.0.9/24  n3
e4e      true      n3_clus2   up/up         10.10.0.10/24 n3
e4a      true      n4_clus1   up/up         10.10.0.11/24 n4
e4e      true      n4_clus2   up/up         10.10.0.12/24 n4
12 entries were displayed.
```

+

```
cluster::*> system cluster-switch show
```

Switch Model	Type	Address
C1 NX3232C	cluster-network	10.10.1.103
Serial Number: FOX000001		
Is Monitored: true		
Reason:		
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version		
7.0(3)I4(1)		
Version Source: CDP		
C2 NX3232C	cluster-network	10.10.1.104
Serial Number: FOX000002		
Is Monitored: true		
Reason:		
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version		
7.0(3)I4(1)		
Version Source: CDP		
CL1 NX5596	cluster-network	10.10.1.101
Serial Number: 01234567		
Is Monitored: true		
Reason:		
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version		
7.1(1)N1(1)		
Version Source: CDP		
CL2 NX5596	cluster-network	10.10.1.102
Serial Number: 01234568		
Is Monitored: true		
Reason:		
Software Version: Cisco Nexus Operating System (NX-OS) Software, Version		
7.1(1)N1(1)		
Version Source: CDP		

```
4 entries were displayed.
```

30. Remove the replaced Nexus 5596 by using the `system cluster-switch delete` command, if it is not automatically removed:

```
system cluster-switch delete -device switch-name
```

Show example

```
cluster::> system cluster-switch delete -device CL1  
cluster::> system cluster-switch delete -device CL2
```

What's next?

[Complete your migration.](#)

Complete your migration from 5596 switches to 3232C switches

Complete the following steps to finalize the Nexus 5596 switches migration to Nexus 3232C switches.

Steps

1. Verify that the proper cluster switches are monitored:

```
system cluster-switch show
```

Show example

```
cluster::> system cluster-switch show

Switch                               Type                               Address
Model
-----
C1                                   cluster-network                   10.10.1.103
NX3232C
  Serial Number: FOX000001
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                    7.0(3)I4(1)
  Version Source: CDP

C2                                   cluster-network                   10.10.1.104
NX3232C
  Serial Number: FOX000002
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                    7.0(3)I4(1)
  Version Source: CDP

2 entries were displayed.
```

2. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next?

[Configure switch health monitoring.](#)

Migrate from two-node switchless clusters

Migrate from a two-node switchless cluster workflow

Follow these workflow steps to migrate from a two-node switchless cluster to a cluster with Cisco Nexus 3232C cluster switches.

1

Migration requirements

Review the example switch information for the migration process.

2

Prepare for migration

Prepare your two-node switchless cluster for migration to a two-node switched cluster.

3

Configure your ports

Configure your two-node switchless cluster for migration to a two-node switched cluster.

4

Complete your migration

Complete your migration to a two-node switched cluster.

Migration requirements

If you have a two-node switchless cluster, you can migrate to a two-node switched cluster that includes Cisco Nexus 3232C cluster network switches. This is a nondisruptive procedure.

Before you begin

Verify the following installations and connections:

- Ports are available for node connections. The cluster switches use the Inter-Switch Link (ISL) ports e1/31-32.
- You have 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.
- The configurations are properly set up and functioning.

The two nodes must be connected and functioning in a two-node switchless cluster setting.

- All cluster ports are in the **up** state.
- The Cisco Nexus 3232C cluster switch are supported.
- The existing cluster network configuration has 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

About the examples used

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**.

The examples in this procedure use two nodes, each using two 40 GbE cluster interconnect ports **e4a** and **e4e**. The [Hardware Universe](#) has details about the cluster ports on your platforms.

- **n1_clus1** is the first cluster logical interface (LIF) to be connected to cluster switch **C1** for node **n1**.
- **n1_clus2** is the first cluster LIF to be connected to cluster switch **C2** for node **n1**.
- **n2_clus1** is the first cluster LIF to be connected to cluster switch **C1** for node **n2**.
- **n2_clus2** is the second cluster LIF to be connected to cluster switch **C2** for node **n2**.
- The number of 10 GbE and 40/100 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.



The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

What's next?

[Prepare for migration.](#)

Prepare for migration from two-node switchless clusters to two-node switched clusters

Follow these steps to prepare your two-node switchless cluster to migrate to a two-node switched cluster that includes Cisco Nexus 3232C cluster network switches.

Steps

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all - message MAINT=xh
```

x is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Determine the administrative or operational status for each cluster interface:
 - a. Display the network port attributes:

```
network port show -role cluster
```


Show example

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
-----
e4a         Cluster    Cluster    up    9000 auto/40000 -
e4e         Cluster    Cluster    up    9000 auto/40000 -
-
Node: n2

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
-----
e4a         Cluster    Cluster    up    9000 auto/40000 -
e4e         Cluster    Cluster    up    9000 auto/40000 -
4 entries were displayed.
```

b. Display information about the logical interfaces and their designated home nodes:

```
network interface show -role cluster
```

Show example

```
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
e4a      true
          n1_clus2    up/up      10.10.0.2/24    n1
e4e      true
          n2_clus1    up/up      10.10.0.3/24    n2
e4a      true
          n2_clus2    up/up      10.10.0.4/24    n2
e4e      true

4 entries were displayed.
```

- c. Verify that switchless cluster detection is enabled using the advanced privilege command:

```
network options detect-switchless-cluster show`
```

Show example

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
```

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 *Cisco Ethernet Switches* page on the NetApp Support Site.

[Cisco Ethernet Switches](#)

- b. Note your switch and the required software versions in the table on that page.

- c. Download the appropriate version of RCF.

- d. Select **CONTINUE** on the **Description** page, accept the license agreement, and then follow the instructions on the **Download** page to download the RCF.
- e. Download the appropriate version of the image software.

[Cisco Cluster and Management Network Switch Reference Configuration File download page](#)

4. Select **CONTINUE** on the **Description** page, accept the license agreement, and then follow the instructions on the **Download** page to download the RCF.
5. 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 following list in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Show example

The following example shows ports 1 through 30 being disabled on Nexus 3232C cluster switches C1 and C2 using a configuration supported in RCF NX3232_RCF_v1.0_24p10g_24p100g.txt:

```
C1# copy running-config startup-config
[] 100% Copy complete.
C1# configure
C1(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C1(config-if-range)# shutdown
C1(config-if-range)# exit
C1(config)# exit
C2# copy running-config startup-config
[] 100% Copy complete.
C2# configure
C2(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C2(config-if-range)# shutdown
C2(config-if-range)# exit
C2(config)# exit
```

6. Connect ports 1/31 and 1/32 on C1 to the same ports on C2 using supported cabling.
7. Verify that the ISL ports are operational on C1 and C2:

```
show port-channel summary
```

For more information on Cisco commands, see the following list in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Show example

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)
```

8. Display the list of neighboring devices on the switch.

For more information on Cisco commands, see the following list in the [Cisco Nexus 3000 Series NX-OS Command References](#).

Show example

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
```

9. Display the cluster port connectivity on each node:

```
network device-discovery show
```

Show example

The following example shows the cluster port connectivity displayed for a two-node switchless cluster configuration:

```
cluster::*> network device-discovery show
```

Node	Local Port	Discovered Device	Interface	Platform
n1	/cdp			
	e4a	n2	e4a	FAS9000
	e4e	n2	e4e	FAS9000
n2	/cdp			
	e4a	n1	e4a	FAS9000
	e4e	n1	e4e	FAS9000

What's next?

[Configure your ports.](#)

Configure your ports for migration from a two-node switchless cluster to a two-node switched cluster

Follow these steps to configure your ports for migration from a two-node switchless cluster to a two-node switched cluster on Nexus 3232C switches.

Steps

1. Migrate the n1_clus1 and n2_clus1 LIFs to the physical ports of their destination nodes:

```
network interface migrate -vserver vserver-name -lif lif-name source-node  
source-node-name -destination-port destination-port-name
```

Show example

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 e4e  
cluster::*> network interface migrate -vserver cluster -lif n2_clus1  
-source-node n2  
-destination-node n2 -destination-port e4e
```

2. Verify the cluster interfaces have successfully migrated:

```
network interface show -role cluster
```

Show example

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
e4e      false
      n1_clus2    up/up      10.10.0.2/24    n1
e4e      true
      n2_clus1    up/up      10.10.0.3/24    n2
e4e      false
      n2_clus2    up/up      10.10.0.4/24    n2
e4e      true
4 entries were displayed.
```

3. Shut down cluster ports for the n1_clus1 and n2_clus1 LIFs, which were migrated in step 9:

```
network port modify -node node-name -port port-name -up-admin false
```

Show example

You must execute the command for each port as shown in the following example:

```
cluster::*> network port modify -node n1 -port e4a -up-admin false
cluster::*> network port modify -node n2 -port e4a -up-admin false
```

4. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
n1		
3/5/2022 19:21:18 -06:00	n1_clus2	n2-clus1
none		
3/5/2022 19:21:20 -06:00	n1_clus2	n2_clus2
none		
n2		
3/5/2022 19:21:18 -06:00	n2_clus2	n1_clus1
none		
3/5/2022 19:21:20 -06:00	n2_clus2	n1_clus2
none		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```



```

cluster1::*> cluster ping-cluster -node local
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1          e4a    10.10.0.1
Cluster n1_clus2 n1          e4e    10.10.0.2
Cluster n2_clus1 n2          e4a    10.10.0.3
Cluster n2_clus2 n2          e4e    10.10.0.4
Local = 10.10.0.1 10.10.0.2
Remote = 10.10.0.3 10.10.0.4
Cluster Vserver Id = 4294967293 Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s) .....
Detected 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)

```

5. Disconnect the cable from e4a on node n1.

You can refer to the running configuration and connect the first 40 GbE port on the switch C1 (port 1/7 in this example) to e4a on n1 using cabling supported for Nexus 3232C switches.

6. Disconnect the cable from e4a on node n2.

You can refer to the running configuration and connect e4a to the next available 40 GbE port on C1, port 1/8, using supported cabling.

7. 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](#).

Show example

The following example shows ports 1 through 30 being enabled on Nexus 3232C cluster switches C1 and C2 using the configuration supported in RCF NX3232_RCF_v1.0_24p10g_26p100g.txt:

```
C1# configure
C1(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C1(config-if-range)# no shutdown
C1(config-if-range)# exit
C1(config)# exit
```

8. Enable the first cluster port, e4a, on each node:

```
network port modify -node node-name -port port-name -up-admin true
```

Show example

```
cluster::*> network port modify -node n1 -port e4a -up-admin true
cluster::*> network port modify -node n2 -port e4a -up-admin true
```

9. Verify that the clusters are up on both nodes:

```
network port show -role cluster
```

Show example

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e4a       Cluster      Cluster      up    9000 auto/40000  -
e4e       Cluster      Cluster      up    9000 auto/40000  -
-

Node: n2

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e4a       Cluster      Cluster      up    9000 auto/40000  -
e4e       Cluster      Cluster      up    9000 auto/40000  -

4 entries were displayed.
```

10. For each node, revert all of the migrated cluster interconnect LIFs:

```
network interface revert -vserver cluster -lif lif-name
```

Show example

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
```

11. Verify that all the LIFs 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.

Show example

```
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
e4a      n1_clus1   up/up     10.10.0.1/24  n1
true
e4e      n1_clus2   up/up     10.10.0.2/24  n1
true
e4a      n2_clus1   up/up     10.10.0.3/24  n2
true
e4e      n2_clus2   up/up     10.10.0.4/24  n2
true
4 entries were displayed.
```

12. Display the cluster port connectivity on each node:

```
network device-discovery show
```

Show example

```
cluster::*> network device-discovery show
      Local   Discovered
Node   Port   Device   Interface   Platform
-----
n1     /cdp
      e4a   C1       Ethernet1/7  N3K-C3232C
      e4e   n2       e4e         FAS9000
n2     /cdp
      e4a   C1       Ethernet1/8  N3K-C3232C
      e4e   n1       e4e         FAS9000
```

13. Migrate clus2 to port e4a 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
```

Show example

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 e4a
cluster::*> network interface migrate -vserver cluster -lif n2_clus2
-source-node n2
-destination-node n2 -destination-port e4a
```

14. Shut down cluster ports clus2 LIF on both nodes:

```
network port modify
```

Show example

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 e4e -up-admin false
cluster::*> network port modify -node n2 -port e4e -up-admin false
```

15. Verify the cluster LIF status:

```
network interface show
```

Show example

```
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
e4a     true
      n1_clus2    up/up      10.10.0.2/24    n1
e4a     false
      n2_clus1    up/up      10.10.0.3/24    n2
e4a     true
      n2_clus2    up/up      10.10.0.4/24    n2
e4a     false
4 entries were displayed.
```

16. Disconnect the cable from e4e on node n1.

You can refer to the running configuration and connect the first 40 GbE port on switch C2 (port 1/7 in this example) to e4e on node n1, using the appropriate cabling for the Nexus 3232C switch model.

17. Disconnect the cable from e4e on node n2.

You can refer to the running configuration and connect e4e to the next available 40 GbE port on C2, port 1/8, using the appropriate cabling for the Nexus 3232C switch model.

18. Enable all node-facing ports on C2.

Show example

The following example shows ports 1 through 30 being enabled on Nexus 3132Q-V cluster switches C1 and C2 using a configuration supported in RCF NX3232C_RCF_v1.0_24p10g_26p100g.txt:

```
C2# configure
C2(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-
4,e1/7-30
C2(config-if-range)# no shutdown
C2(config-if-range)# exit
C2(config)# exit
```

19. Enable the second cluster port, e4e, on each node:

```
network port modify
```

Show example

The following example shows the second cluster port e4e being brought up on each node:

```
cluster::*> network port modify -node n1 -port e4e -up-admin true
cluster::*> *network port modify -node n2 -port e4e -up-admin true*s
```

20. For each node, revert all of the migrated cluster interconnect LIFs:

```
network interface revert
```

Show example

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
```

What's next?

[Complete your migration.](#)

Complete your migration from a two-node switchless cluster to a two-node switched cluster

Complete the following steps to finalize the two-node switchless cluster migration to a two-node switched cluster on Nexus 3232C switches.

Steps

1. 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.

Show example

```
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
e4a      true
          n1_clus2  up/up    10.10.0.2/24   n1
e4e      true
          n2_clus1  up/up    10.10.0.3/24   n2
e4a      true
          n2_clus2  up/up    10.10.0.4/24   n2
e4e      true
4 entries were displayed.
```

2. Verify that all of the cluster interconnect ports are in the up state:

```
network port show -role cluster
```

3. Display the cluster switch port numbers through which each cluster port is connected to each node:

```
network device-discovery show
```

Show example

```
cluster::*> network device-discovery show
          Local   Discovered
Node      Port   Device           Interface           Platform
-----
-----
n1        /cdp
          e4a    C1               Ethernet1/7         N3K-C3232C
          e4e    C2               Ethernet1/7         N3K-C3232C
n2        /cdp
          e4a    C1               Ethernet1/8         N3K-C3232C
          e4e    C2               Ethernet1/8         N3K-C3232C
```


4. Display discovered and monitored cluster switches:

```
system cluster-switch show
```

Show example

```
cluster::*> system cluster-switch show
```

Switch Model	Type	Address
C1 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	cluster-network	10.10.1.101
C2 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	cluster-network	10.10.1.102

Version Source: CDP 2 entries were displayed.

5. Verify that switchless cluster detection changed the switchless cluster option to disabled:

```
network options switchless-cluster show
```

6. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
n1		
3/5/2022 19:21:18 -06:00	n1_clus2	n2-clus1
none		
3/5/2022 19:21:20 -06:00	n1_clus2	n2_clus2
none		
n2		
3/5/2022 19:21:18 -06:00	n2_clus2	n1_clus1
none		
3/5/2022 19:21:20 -06:00	n2_clus2	n1_clus2
none		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1          e4a    10.10.0.1
Cluster n1_clus2 n1          e4e    10.10.0.2
Cluster n2_clus1 n2          e4a    10.10.0.3
Cluster n2_clus2 n2          e4e    10.10.0.4
Local = 10.10.0.1 10.10.0.2
Remote = 10.10.0.3 10.10.0.4
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s) .....
Detected 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)

```

7. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next?

[Configure switch health monitoring.](#)

Replace switches

Replace a Cisco Nexus 3232C cluster switch

Follow these steps to replace a defective Cisco Nexus 3232C switch in a cluster. This is a non-disruptive procedure.

Review requirements

What you'll need

Make sure that the existing cluster and network configuration has the following characteristics:

- The Nexus 3232C cluster infrastructure are 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) are in the **up** state and are not migrated.

The replacement Cisco Nexus 3232C switch has the following characteristics:

- Management network connectivity is functional.
- Console access to the replacement switch is in place.
- The appropriate RCF and NX-OS operating system image is loaded onto the switch.
- Initial customization of the switch is complete.

For more information

See the following:

- [Cisco Ethernet Switch description page](#)
- [Hardware Universe](#)

Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when replacing your switch:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport before and after maintenance to disable case creation for the duration of the maintenance. See this Knowledge Base article [SU92: How to suppress automatic case creation during scheduled maintenance windows](#) for further details.
- Enable session logging for any CLI sessions. For instructions on how to enable session logging, review the "Logging Session Output" section in this Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

Replace the switch

About this task

This replacement 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.

About the examples

This replacement procedure replaces the second Nexus 3232C cluster switch CL2 with the new 3232C switch C2.

The examples in this procedure use the following switch and node nomenclature:

- The four nodes are n1, n2, n3, and n4.
- n1_clus1 is the first cluster logical interface (LIF) connected to cluster switch C1 for node n1.
- n1_clus2 is the first cluster LIF connected to cluster switch CL2 or C2 for node n1.
- n1_clus3 is the second LIF connected to cluster switch C2 for node n1.-
- n1_clus4 is the second LIF connected to cluster switch CL1, for node n1.

The number of 10 GbE and 40/100 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.

The examples in this replacement procedure use four nodes. Two of the nodes use four 10 GB cluster interconnect ports: e0a, e0b, e0c, and e0d. The other two nodes use two 40 GB cluster interconnect ports: e4a and e4e. See the [Hardware Universe](#) to verify the correct cluster ports for your platform.

Step 1: Display and migrate the cluster ports to switch

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all - message MAINT=xh
```

x is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Display information about the devices in your configuration:

```
network device-discovery show
```

Show example

```
cluster::> network device-discovery show
```

Node	Local Port	Discovered Device	Interface	Platform
n1	/cdp			
	e0a	CL1	Ethernet1/1/1	N3K-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

3. Determine the administrative or operational status for each cluster interface.

a. Display the network port attributes:

```
network port show -role cluster
```

Show example

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status Status Speed (Mbps)
-----
e0a Cluster Cluster up 9000 auto/10000 -
e0b Cluster Cluster up 9000 auto/10000 -
e0c Cluster Cluster up 9000 auto/10000 -
e0d Cluster Cluster up 9000 auto/10000 -
-

Node: n2

Ignore

Health Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status Status Speed (Mbps)
-----
e0a Cluster Cluster up 9000 auto/10000 -
e0b Cluster Cluster up 9000 auto/10000 -
e0c Cluster Cluster up 9000 auto/10000 -
e0d Cluster Cluster up 9000 auto/10000 -
-

Node: n3

Ignore

Health Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status Status Speed (Mbps)
-----
e4a Cluster Cluster up 9000 auto/40000 -
-
e4e Cluster Cluster up 9000 auto/40000 -
```

```
-  
  
Node: n4  
  
Ignore  
  
Health Health Speed (Mbps)  
Port IPspace Broadcast Domain Link MTU Admin/Oper  
Status Status  
-----  
-----  
e4a Cluster Cluster up 9000 auto/40000 -  
e4e Cluster Cluster up 9000 auto/40000 -
```

b. Display information about the logical interfaces (LIFs):

```
network interface show -role cluster
```


Show example

```
cluster::*> network interface show -role cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
e0a	n1_clus1	up/up	10.10.0.1/24	n1
	true			
e0b	n1_clus2	up/up	10.10.0.2/24	n1
	true			
e0c	n1_clus3	up/up	10.10.0.3/24	n1
	true			
e0d	n1_clus4	up/up	10.10.0.4/24	n1
	true			
e0a	n2_clus1	up/up	10.10.0.5/24	n2
	true			
e0b	n2_clus2	up/up	10.10.0.6/24	n2
	true			
e0c	n2_clus3	up/up	10.10.0.7/24	n2
	true			
e0d	n2_clus4	up/up	10.10.0.8/24	n2
	true			
e0a	n3_clus1	up/up	10.10.0.9/24	n3
	true			
e0e	n3_clus2	up/up	10.10.0.10/24	n3
	true			
e0a	n4_clus1	up/up	10.10.0.11/24	n4
	true			
e0e	n4_clus2	up/up	10.10.0.12/24	n4
	true			

c. Display the discovered cluster switches:

```
system cluster-switch show
```

Show example

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: None
    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: None
    Software Version: Cisco Nexus Operating System (NX-OS)
    Software, Version 7.0(3)I6(1)
    Version Source: CDP
```

4. Verify that the appropriate RCF and image are installed on the new Nexus 3232C switch and make any necessary site customizations.

a. Go to the NetApp Support Site.

mysupport.netapp.com

b. Go to the **Cisco Ethernet Switches** page and note the required software versions in the table.

[Cisco Ethernet Switches](#)

c. Download the appropriate version of the RCF.

d. Click **CONTINUE** on the **Description** page, accept the license agreement, and then navigate to the **Download** page.

e. Download the correct version of the image software from the **Cisco® Cluster and Management Network Switch Reference Configuration File Download** page.

[Cisco® Cluster and Management Network Switch Reference Configuration File Download](#)

5. Migrate the cluster LIFs to the physical node ports connected to the replacement switch C2:

```
network interface migrate -vserver vserver-name -lif lif-name -source-node
node-name -destination-node node-name -destination-port port-name
```

Show example

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 -destination-
node n4 -destination-port e4a
```

6. Verify the status of the cluster ports and their home designations:

```
network interface show -role cluster
```

Show example

```
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
e0a      n1_clus1   up/up     10.10.0.1/24  n1
true
e0a      n1_clus2   up/up     10.10.0.2/24  n1
false
e0d      n1_clus3   up/up     10.10.0.3/24  n1
false
e0d      n1_clus4   up/up     10.10.0.4/24  n1
true
e0a      n2_clus1   up/up     10.10.0.5/24  n2
true
e0a      n2_clus2   up/up     10.10.0.6/24  n2
false
e0d      n2_clus3   up/up     10.10.0.7/24  n2
false
e0d      n2_clus4   up/up     10.10.0.8/24  n2
true
e4a      n3_clus1   up/up     10.10.0.9/24  n3
true
e4a      n3_clus2   up/up     10.10.0.10/24  n3
false
e4a      n4_clus1   up/up     10.10.0.11/24  n4
true
e4a      n4_clus2   up/up     10.10.0.12/24  n4
false
```

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
```

Show example

The following example shows the cluster interconnect ports are shut down on all nodes:

```
cluster::*> network port modify -node n1 -port e0b -up-admin false
cluster::*> network port modify -node n1 -port e0c -up-admin false
cluster::*> network port modify -node n2 -port e0b -up-admin false
cluster::*> network port modify -node n2 -port e0c -up-admin false
cluster::*> network port modify -node n3 -port e4e -up-admin false
cluster::*> network port modify -node n4 -port e4e -up-admin false
```

8. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
-----	-----	-----
n1		
3/5/2022 19:21:18 -06:00	n1_clus2	n2-clus1
none		
3/5/2022 19:21:20 -06:00	n1_clus2	n2_clus2
none		
.		
.		
n2		
3/5/2022 19:21:18 -06:00	n2_clus2	n1_clus1
none		
3/5/2022 19:21:20 -06:00	n2_clus2	n1_clus2
none		
.		
.		
n3		
.		
.		
.n4		
.		
.		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1      e0a    10.10.0.1
Cluster n1_clus2 n1      e0b    10.10.0.2
Cluster n1_clus3 n1      e0c    10.10.0.3
Cluster n1_clus4 n1      e0d    10.10.0.4
Cluster n2_clus1 n2      e0a    10.10.0.5
Cluster n2_clus2 n2      e0b    10.10.0.6
Cluster n2_clus3 n2      e0c    10.10.0.7
Cluster n2_clus4 n2      e0d    10.10.0.8
Cluster n3_clus1 n4      e0a    10.10.0.9
Cluster n3_clus2 n3      e0e    10.10.0.10
Cluster n4_clus1 n4      e0a    10.10.0.11
Cluster n4_clus2 n4      e0e    10.10.0.12
Local = 10.10.0.1 10.10.0.2 10.10.0.3 10.10.0.4
Remote = 10.10.0.5 10.10.0.6 10.10.0.7 10.10.0.8 10.10.0.9 10.10.0.10
10.10.0.11
10.10.0.12 Cluster Vserver Id = 4294967293 Ping status:
....
Basic connectivity succeeds on 32 path(s)
Basic connectivity fails on 0 path(s) .....
Detected 9000 byte MTU on 32 path(s):
    Local 10.10.0.1 to Remote 10.10.0.5
    Local 10.10.0.1 to Remote 10.10.0.6
    Local 10.10.0.1 to Remote 10.10.0.7
    Local 10.10.0.1 to Remote 10.10.0.8
    Local 10.10.0.1 to Remote 10.10.0.9
    Local 10.10.0.1 to Remote 10.10.0.10
    Local 10.10.0.1 to Remote 10.10.0.11
    Local 10.10.0.1 to Remote 10.10.0.12
    Local 10.10.0.2 to Remote 10.10.0.5
    Local 10.10.0.2 to Remote 10.10.0.6
    Local 10.10.0.2 to Remote 10.10.0.7
    Local 10.10.0.2 to Remote 10.10.0.8
    Local 10.10.0.2 to Remote 10.10.0.9
    Local 10.10.0.2 to Remote 10.10.0.10
    Local 10.10.0.2 to Remote 10.10.0.11
    Local 10.10.0.2 to Remote 10.10.0.12
    Local 10.10.0.3 to Remote 10.10.0.5
    Local 10.10.0.3 to Remote 10.10.0.6
    Local 10.10.0.3 to Remote 10.10.0.7
    Local 10.10.0.3 to Remote 10.10.0.8
    Local 10.10.0.3 to Remote 10.10.0.9
    Local 10.10.0.3 to Remote 10.10.0.10

```

```
Local 10.10.0.3 to Remote 10.10.0.11
Local 10.10.0.3 to Remote 10.10.0.12
Local 10.10.0.4 to Remote 10.10.0.5
Local 10.10.0.4 to Remote 10.10.0.6
Local 10.10.0.4 to Remote 10.10.0.7
Local 10.10.0.4 to Remote 10.10.0.8
Local 10.10.0.4 to Remote 10.10.0.9
Local 10.10.0.4 to Remote 10.10.0.10
Local 10.10.0.4 to Remote 10.10.0.11
Local 10.10.0.4 to Remote 10.10.0.12
Larger than PMTU communication succeeds on 32 path(s) RPC status:
8 paths up, 0 paths down (tcp check)
8 paths up, 0 paths down (udp check)
```

Step 2: Migrate ISLs to switch CL1 and C2

1. 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](#).

Show example

```
(CL1)# configure
(CL1) (Config)# interface e1/31-32
(CL1) (config-if-range)# shutdown
(CL1) (config-if-range)# exit
(CL1) (Config)# exit
(CL1)#
```

2. Remove all the cables attached to the cluster switch CL2 and reconnect them to the replacement switch C2 for all the nodes.
3. 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.
4. 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](#).

Show example

```
(CL1) # configure
(CL1) (Config) # interface e1/31-32
(CL1) (config-if-range) # no shutdown
(CL1) (config-if-range) # exit
(CL1) (Config) # exit
(CL1) #
```

5. 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:

Show example

```
CL1# show port-channel summary
Flags: D - Down          P - Up in port-channel (members)
      I - Individual    H - Hot-standby (LACP only)
      s - Suspended     r - Module-removed
      S - Switched      R - Routed
      U - Up (port-channel)
      M - Not in use. Min-links not met

-----
-----
Group Port-          Type   Protocol  Member Ports
Channel
-----
-----
1      Po1 (SU)       Eth    LACP      Eth1/31 (P)  Eth1/32 (P)
```

6. 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](#).

Show example

Ports Eth1/31 and Eth1/32 should indicate (P), which means that both ISL ports are up in the port-channel.

```
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)
```

7. 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
```

Show example

```
cluster::*> network port modify -node n1 -port e0b -up-admin true
cluster::*> network port modify -node n1 -port e0c -up-admin true
cluster::*> network port modify -node n2 -port e0b -up-admin true
cluster::*> network port modify -node n2 -port e0c -up-admin true
cluster::*> network port modify -node n3 -port e4e -up-admin true
cluster::*> network port modify -node n4 -port e4e -up-admin true
```

Step 3: Revert all LIFs to originally assigned ports

1. Revert all the migrated cluster interconnect LIFs on all the nodes:

```
network interface revert -vserver cluster -lif lif-name
```

Show example

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
```

2. Verify that the cluster interconnect ports are now reverted to their home:

```
network interface show
```

Show example

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
e0a      true      n1_clus1   up/up         10.10.0.1/24   n1
e0b      true      n1_clus2   up/up         10.10.0.2/24   n1
e0c      true      n1_clus3   up/up         10.10.0.3/24   n1
e0d      true      n1_clus4   up/up         10.10.0.4/24   n1
e0a      true      n2_clus1   up/up         10.10.0.5/24   n2
e0b      true      n2_clus2   up/up         10.10.0.6/24   n2
e0c      true      n2_clus3   up/up         10.10.0.7/24   n2
e0d      true      n2_clus4   up/up         10.10.0.8/24   n2
e4a      true      n3_clus1   up/up         10.10.0.9/24   n3
e4e      true      n3_clus2   up/up         10.10.0.10/24  n3
e4a      true      n4_clus1   up/up         10.10.0.11/24  n4
e4e      true      n4_clus2   up/up         10.10.0.12/24  n4
```

3. Verify that the cluster ports are connected:

```
network port show -role cluster
```

Show example

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Speed(Mbps) Health
-----
-----
e0a      Cluster      Cluster      up    9000  auto/10000  -
e0b      Cluster      Cluster      up    9000  auto/10000  -
e0c      Cluster      Cluster      up    9000  auto/10000  -
e0d      Cluster      Cluster      up    9000  auto/10000  -
-

Node: n2

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Speed(Mbps) Health
-----
-----
e0a      Cluster      Cluster      up    9000  auto/10000  -
e0b      Cluster      Cluster      up    9000  auto/10000  -
e0c      Cluster      Cluster      up    9000  auto/10000  -
e0d      Cluster      Cluster      up    9000  auto/10000  -
-

Node: n3

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Speed(Mbps) Health
-----
-----
e4a      Cluster      Cluster      up    9000  auto/40000  -
e4e      Cluster      Cluster      up    9000  auto/40000  -
-

Node: n4
```

Ignore

Health						Speed (Mbps)	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	-
-							

4. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
-----	-----	-----
n1		
3/5/2022 19:21:18 -06:00	n1_clus2	n2-clus1
none		
3/5/2022 19:21:20 -06:00	n1_clus2	n2_clus2
none		
.		
.		
n2		
3/5/2022 19:21:18 -06:00	n2_clus2	n1_clus1
none		
3/5/2022 19:21:20 -06:00	n2_clus2	n1_clus2
none		
.		
.		
n3		
.		
.		
.n4		
.		
.		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1      e0a    10.10.0.1
Cluster n1_clus2 n1      e0b    10.10.0.2
Cluster n1_clus3 n1      e0c    10.10.0.3
Cluster n1_clus4 n1      e0d    10.10.0.4
Cluster n2_clus1 n2      e0a    10.10.0.5
Cluster n2_clus2 n2      e0b    10.10.0.6
Cluster n2_clus3 n2      e0c    10.10.0.7
Cluster n2_clus4 n2      e0d    10.10.0.8
Cluster n3_clus1 n4      e0a    10.10.0.9
Cluster n3_clus2 n3      e0e    10.10.0.10
Cluster n4_clus1 n4      e0a    10.10.0.11
Cluster n4_clus2 n4      e0e    10.10.0.12
Local = 10.10.0.1 10.10.0.2 10.10.0.3 10.10.0.4
Remote = 10.10.0.5 10.10.0.6 10.10.0.7 10.10.0.8 10.10.0.9 10.10.0.10
10.10.0.11
10.10.0.12 Cluster Vserver Id = 4294967293 Ping status:
....
Basic connectivity succeeds on 32 path(s)
Basic connectivity fails on 0 path(s) .....
Detected 9000 byte MTU on 32 path(s):
    Local 10.10.0.1 to Remote 10.10.0.5
    Local 10.10.0.1 to Remote 10.10.0.6
    Local 10.10.0.1 to Remote 10.10.0.7
    Local 10.10.0.1 to Remote 10.10.0.8
    Local 10.10.0.1 to Remote 10.10.0.9
    Local 10.10.0.1 to Remote 10.10.0.10
    Local 10.10.0.1 to Remote 10.10.0.11
    Local 10.10.0.1 to Remote 10.10.0.12
    Local 10.10.0.2 to Remote 10.10.0.5
    Local 10.10.0.2 to Remote 10.10.0.6
    Local 10.10.0.2 to Remote 10.10.0.7
    Local 10.10.0.2 to Remote 10.10.0.8
    Local 10.10.0.2 to Remote 10.10.0.9
    Local 10.10.0.2 to Remote 10.10.0.10
    Local 10.10.0.2 to Remote 10.10.0.11
    Local 10.10.0.2 to Remote 10.10.0.12
    Local 10.10.0.3 to Remote 10.10.0.5
    Local 10.10.0.3 to Remote 10.10.0.6
    Local 10.10.0.3 to Remote 10.10.0.7
    Local 10.10.0.3 to Remote 10.10.0.8
    Local 10.10.0.3 to Remote 10.10.0.9
    Local 10.10.0.3 to Remote 10.10.0.10

```



```
Local 10.10.0.3 to Remote 10.10.0.11
Local 10.10.0.3 to Remote 10.10.0.12
Local 10.10.0.4 to Remote 10.10.0.5
Local 10.10.0.4 to Remote 10.10.0.6
Local 10.10.0.4 to Remote 10.10.0.7
Local 10.10.0.4 to Remote 10.10.0.8
Local 10.10.0.4 to Remote 10.10.0.9
Local 10.10.0.4 to Remote 10.10.0.10
Local 10.10.0.4 to Remote 10.10.0.11
Local 10.10.0.4 to Remote 10.10.0.12
```

```
Larger than PMTU communication succeeds on 32 path(s) RPC status:
```

```
8 paths up, 0 paths down (tcp check)
```

```
8 paths up, 0 paths down (udp check)
```

Step 4: Verify all ports and LIF are correctly migrated

1. 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

Show example

```
cluster::> network device-discovery show
```

Node	Local Port	Discovered 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

```
cluster::*> network port show -role cluster
```

```
(network port show)
```

```
Node: n1
```

```
Ignore
```

Health	Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Speed (Mbps)	Health	Status
--------	------	---------	-----------	--------	------	-----	------------	--------------	--------	--------

	e0a	Cluster	Cluster		up	9000	auto/10000		-	
	e0b	Cluster	Cluster		up	9000	auto/10000		-	
	e0c	Cluster	Cluster		up	9000	auto/10000		-	
	e0d	Cluster	Cluster		up	9000	auto/10000		-	

```
Node: n2
```

```
Ignore
```

Health	Speed (Mbps)	Health
--------	--------------	--------

```

Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e0a       Cluster      Cluster      up    9000  auto/10000  -
e0b       Cluster      Cluster      up    9000  auto/10000  -
e0c       Cluster      Cluster      up    9000  auto/10000  -
e0d       Cluster      Cluster      up    9000  auto/10000  -

```

Node: 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  -

```

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

```

```
cluster::*> system cluster-switch show
```

```

Switch                               Type                Address
Model
-----
CL1                                   cluster-network    10.10.1.101
NX3232C
    Serial Number: FOX000001
    Is Monitored: true
    Reason: None
    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: None
    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: None
```

```
Software Version: Cisco Nexus Operating System (NX-OS)
```

```
Software, Version 7.0(3)I6(1)
```

```
Version Source: CDP 3 entries were displayed.
```

2. Delete the replaced cluster switch CL2 if it has not been removed automatically:

```
system cluster-switch delete -device cluster-switch-name
```

3. Verify that the proper cluster switches are monitored:

```
system cluster-switch show
```

Show example

The following example shows the cluster switches are monitored because the Is Monitored state is true.

```
cluster::> system cluster-switch show
```

Switch Model	Type	Address
CL1 NX3232C	cluster-network	10.10.1.101
Serial Number: FOX000001		
Is Monitored: true		
Reason: None		
Software Version: Cisco Nexus Operating System (NX-OS)		
Software, Version 7.0(3)I6(1)		
Version Source: CDP		
C2 NX3232C	cluster-network	10.10.1.103
Serial Number: FOX000002		
Is Monitored: true		
Reason: None		
Software Version: Cisco Nexus Operating System (NX-OS)		
Software, Version 7.0(3)I6(1)		
Version Source: CDP		

4. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next?

[Configure switch health monitoring.](#)

Replace Cisco Nexus 3232C cluster switches with switchless connections

You can migrate from a cluster with a switched cluster network to one where two nodes are directly connected for ONTAP 9.3 and later.

Review requirements

Guidelines

Review the following guidelines:

- Migrating to a two-node switchless cluster configuration is a nondisruptive operation. Most systems have two dedicated cluster interconnect ports on each node, but you can also use this procedure for systems with a larger number of dedicated cluster interconnect ports on each node, such as four, six or eight.
- You cannot use the switchless cluster interconnect feature with more than two nodes.
- If you have an existing two-node cluster that uses cluster interconnect switches and is running ONTAP 9.3 or later, you can replace the switches with direct, back-to-back connections between the nodes.


What you'll need

- A healthy cluster that consists of two nodes connected by cluster switches. The nodes must be running the same ONTAP release.
- Each node with the required number of dedicated cluster ports, which provide redundant cluster interconnect connections to support your system configuration. For example, there are two redundant ports for a system with two dedicated cluster interconnect ports on each node.

Migrate the switches

About this task

The following procedure removes the cluster switches in a two-node cluster and replaces each connection to the switch with a direct connection to the partner node.



About the examples

The examples in the following procedure show nodes that are using "e0a" and "e0b" as cluster ports. Your nodes might be using different cluster ports as they vary by system.

Step 1: Prepare for migration

1. Change the privilege level to advanced, entering `y` when prompted to continue:

```
set -privilege advanced
```

The advanced prompt `*>` appears.

2. ONTAP 9.3 and later supports automatic detection of switchless clusters, which is enabled by default.

You can verify that detection of switchless clusters is enabled by running the advanced privilege command:

```
network options detect-switchless-cluster show
```

Show example

The following example output shows if the option is enabled.

```
cluster::*> network options detect-switchless-cluster show
(network options detect-switchless-cluster show)
Enable Switchless Cluster Detection: true
```

If "Enable Switchless Cluster Detection" is `false`, contact NetApp support.

3. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message
MAINT=<number_of_hours>h
```

where `h` is the duration of the maintenance window in hours. The message notifies technical support of this maintenance task so that they can suppress automatic case creation during the maintenance window.

In the following example, the command suppresses automatic case creation for two hours:

Show example

```
cluster::*> system node autosupport invoke -node * -type all
-message MAINT=2h
```

Step 2: Configure ports and cabling


1. Organize the cluster ports on each switch into groups so that the cluster ports in `group1` go to cluster switch1 and the cluster ports in `group2` go to cluster switch2. These groups are required later in the

procedure.

2. Identify the cluster ports and verify link status and health:

```
network port show -ipspace Cluster
```

In the following example for nodes with cluster ports "e0a" and "e0b", one group is identified as "node1:e0a" and "node2:e0a" and the other group as "node1:e0b" and "node2:e0b". Your nodes might be using different cluster ports because they vary by system.



Verify that the ports have a value of `up` for the "Link" column and a value of `healthy` for the "Health Status" column.

Show example

```
cluster::> network port show -ipSPACE Cluster
Node: node1

Ignore
Speed (Mbps) Health
Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status
Status
-----
-----
e0a Cluster Cluster up 9000 auto/10000 healthy
false
e0b Cluster Cluster up 9000 auto/10000 healthy
false

Node: node2

Ignore
Speed (Mbps) Health
Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status
Status
-----
-----
e0a Cluster Cluster up 9000 auto/10000 healthy
false
e0b Cluster Cluster up 9000 auto/10000 healthy
false
4 entries were displayed.
```

3. Confirm that all the cluster LIFs are on their home ports.

Verify that the “is-home” column is `true` for each of the cluster LIFs:

```
network interface show -vserver Cluster -fields is-home
```

Show example

```
cluster::*> net int show -vserver Cluster -fields is-home
(network interface show)
vserver  lif           is-home
-----  -
Cluster  node1_clus1  true
Cluster  node1_clus2  true
Cluster  node2_clus1  true
Cluster  node2_clus2  true
4 entries were displayed.
```

If there are cluster LIFs that are not on their home ports, revert those LIFs to their home ports:

```
network interface revert -vserver Cluster -lif *
```

4. Disable auto-revert for the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

5. Verify that all ports listed in the previous step are connected to a network switch:

```
network device-discovery show -port cluster_port
```

The “Discovered Device” column should be the name of the cluster switch that the port is connected to.

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to cluster switches "cs1" and "cs2".

```
cluster:::> network device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol  Port   Device (LLDP: ChassisID)  Interface  Platform
-----  -
node1/cdp
          e0a    cs1                       0/11       BES-53248
          e0b    cs2                       0/12       BES-53248
node2/cdp
          e0a    cs1                       0/9        BES-53248
          e0b    cs2                       0/9        BES-53248
4 entries were displayed.
```

6. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

7. Verify that the cluster is healthy:

```
cluster ring show
```

All units must be either master or secondary.


8. Set up the switchless configuration for the ports in group 1.



To avoid potential networking issues, you must disconnect the ports from group1 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.


a. Disconnect all the cables from the ports in group1 at the same time.

In the following example, the cables are disconnected from port "e0a" on each node, and cluster traffic continues through the switch and port "e0b" on each node:



b. Cable the ports in group1 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2:



9. The switchless cluster network option transitions from *false* to *true*. This might take up to 45 seconds. Confirm that the switchless option is set to *true*:

```
network options switchless-cluster show
```

The following example shows that the switchless cluster is enabled:

```
cluster::*> network options switchless-cluster show
Enable Switchless Cluster: true
```

10. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```



Before proceeding to the next step, you must wait at least two minutes to confirm a working back-to-back connection on group 1.


11. Set up the switchless configuration for the ports in group 2.



To avoid potential networking issues, you must disconnect the ports from group2 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.


- a. Disconnect all the cables from the ports in group2 at the same time.

In the following example, the cables are disconnected from port "e0b" on each node, and cluster traffic continues through the direct connection between the "e0a" ports:



b. Cable the ports in group2 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2 and "e0b" on node1 is connected to "e0b" on node2:



Step 3: Verify the configuration

1. Verify that the ports on both nodes are correctly connected:

```
network device-discovery show -port cluster_port
```


Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to the corresponding port on the cluster partner:

```
cluster::> net device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1/cdp
          e0a    node2                      e0a        AFF-A300
          e0b    node2                      e0b        AFF-A300
node1/lldp
          e0a    node2 (00:a0:98:da:16:44)  e0a        -
          e0b    node2 (00:a0:98:da:16:44)  e0b        -
node2/cdp
          e0a    node1                      e0a        AFF-A300
          e0b    node1                      e0b        AFF-A300
node2/lldp
          e0a    node1 (00:a0:98:da:87:49)  e0a        -
          e0b    node1 (00:a0:98:da:87:49)  e0b        -
8 entries were displayed.
```

2. Re-enable auto-revert for the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert true
```

3. Verify that all LIFs are home. This might take a few seconds.

```
network interface show -vserver Cluster -lif lif_name
```

Show example

The LIFs have been reverted if the “Is Home” column is `true`, as shown for `node1_clus2` and `node2_clus2` in the following example:

```
cluster::> network interface show -vserver Cluster -fields curr-
port,is-home
vserver  lif                curr-port  is-home
-----  -
Cluster  node1_clus1            e0a        true
Cluster  node1_clus2            e0b        true
Cluster  node2_clus1            e0a        true
Cluster  node2_clus2            e0b        true
4 entries were displayed.
```

If any cluster LIFS have not returned to their home ports, revert them manually from the local node:

```
network interface revert -vserver Cluster -lif lif_name
```

4. Check the cluster status of the nodes from the system console of either node:

```
cluster show
```

Show example

The following example shows `epsilon` on both nodes to be `false`:

```
Node  Health  Eligibility  Epsilon
-----  -
node1 true    true        false
node2 true    true        false
2 entries were displayed.
```

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

6. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

For more information, see [NetApp KB Article 1010449: How to suppress automatic case creation during scheduled maintenance windows](#).

7. Change the privilege level back to admin:

```
set -privilege admin
```

Cisco Nexus 3132Q-V

Overview

Overview of installation and configuration for Cisco Nexus 3132Q-V switches

Cisco Nexus 3132Q-V switches can be used as cluster switches in your AFF or FAS cluster. Cluster switches allow you to build ONTAP clusters with more than two nodes.

Initial configuration overview

To initially configure a Cisco Nexus 3132Q-V switch on systems running ONTAP, follow these steps:

1. [Complete Cisco Nexus 3132Q-V cabling worksheet](#). The sample cabling worksheet provides examples of recommended port assignments from the switches to the controllers. The blank worksheet provides a template that you can use in setting up your cluster.
2. [Install a Cisco Nexus 3132Q-V cluster switch in a NetApp cabinet](#). Install the Cisco Nexus 3132Q-V switch and pass-through panel in a NetApp cabinet with the standard brackets that are included with the switch.
3. [Configure the Cisco Nexus 3132Q-V switch](#). Set up and configure the Cisco Nexus 3132Q-V switch.
4. [Prepare to install NX-OS software and Reference Configuration File](#). Prepare to install the NX-OS software and the Reference Configuration File (RCF).
5. [Install the NX-OS software](#). Follow this procedure to install the NX-OS software on the Nexus 3132Q-V cluster switch.
6. [Install the Reference Configuration File \(RCF\)](#). Follow this procedure to install the RCF after setting up the Nexus 3132Q-V switch for the first time. You can also use this procedure to upgrade your RCF version.

Additional information

Before you begin installation or maintenance, be sure to review the following:

- [Configuration requirements](#)
- [Required documentation](#)
- [Smart Call Home requirements](#)

Configuration requirements for Cisco Nexus 3132Q-V switches

For Cisco Nexus 3132Q-V switch installation and maintenance, be sure to review network and configuration requirements.

Configuration requirements

To configure your cluster, you need the appropriate number and type of cables and cable connectors for your switches. Depending on the type of switch you are initially configuring, you need to connect to the switch console port with the included console cable; you also need to provide specific network information.

Network requirements

You need the following network information for all switch configurations:

- IP subnet for management network traffic.
- Host names and IP addresses for each of the storage system controllers and all applicable switches.
- Most storage system controllers are managed through the e0M interface by connecting to the Ethernet service port (wrench icon). On AFF A800 and AFF A700 systems, the e0M interface uses a dedicated Ethernet port.

Refer to the [Hardware Universe](#) for latest information.

Documentation requirements for Cisco Nexus 3132Q-V switches

For Cisco Nexus 3132Q-V switch installation and maintenance, be sure to review all the recommended documentation.

Switch documentation

To set up the Cisco Nexus 3132Q-V switches, you need the following documentation from the [Cisco Nexus 3000 Series Switches Support](#) page.

Document title	Description
<i>Nexus 3000 Series Hardware Installation Guide</i>	Provides detailed information about site requirements, switch hardware details, and installation options.
<i>Cisco Nexus 3000 Series Switch Software Configuration Guides</i> (choose the guide for the NX-OS release installed on your switches)	Provides initial switch configuration information that you need before you can configure the switch for ONTAP operation.
<i>Cisco Nexus 3000 Series NX-OS Software Upgrade and Downgrade Guide</i> (choose the guide for the NX-OS release installed on your switches)	Provides information on how to downgrade the switch to ONTAP supported switch software, if necessary.
<i>Cisco Nexus 3000 Series NX-OS Command Reference Master Index</i>	Provides links to the various command references provided by Cisco.
<i>Cisco Nexus 3000 MIBs Reference</i>	Describes the Management Information Base (MIB) files for the Nexus 3000 switches.
<i>Nexus 3000 Series NX-OS System Message Reference</i>	Describes the system messages for Cisco Nexus 3000 series switches, those that are informational, and others that might help diagnose problems with links, internal hardware, or the system software.
<i>Cisco Nexus 3000 Series NX-OS Release Notes</i> (choose the notes for the NX-OS release installed on your switches)	Describes the features, bugs, and limitations for the Cisco Nexus 3000 Series.
Regulatory, Compliance, and Safety Information for the Cisco Nexus 6000, Cisco Nexus 5000 Series, Cisco Nexus 3000 Series, and Cisco Nexus 2000 Series	Provides international agency compliance, safety, and statutory information for the Nexus 3000 series switches.

ONTAP systems documentation

To set up an ONTAP system, you need the following documents for your version of the operating system from the [ONTAP 9 Documentation Center](#).

Name	Description
Controller-specific <i>Installation and Setup Instructions</i>	Describes how to install NetApp hardware.
ONTAP documentation	Provides detailed information about all aspects of the ONTAP releases.
Hardware Universe	Provides NetApp hardware configuration and compatibility information.

Rail kit and cabinet documentation

To install a 3132Q-V Cisco switch in a NetApp cabinet, see the following hardware documentation.

Name	Description
42U System Cabinet, Deep Guide	Describes the FRUs associated with the 42U system cabinet, and provides maintenance and FRU replacement instructions.
Install Cisco Nexus 3132Q-V switch in a NetApp Cabinet	Describes how to install a Cisco Nexus 3132Q-V switch in a four-post NetApp cabinet.

Smart Call Home requirements

To use Smart Call Home feature, review the following guidelines.

Smart Call Home monitors the hardware and software components on your network. When a critical system configuration occurs, it generates an email-based notification and raises an alert to all the recipients that are configured in your destination profile. To use Smart Call Home, you must configure a cluster network switch to communicate using email with the Smart Call Home system. In addition, you can optionally set up your cluster network switch to take advantage of Cisco's embedded Smart Call Home support feature.

Before you can use Smart Call Home, be aware of the following considerations:

- An email server must be in place.
- The switch must have IP connectivity to the email server.
- The contact name (SNMP server contact), phone number, and street address information must be configured. This is required to determine the origin of messages received.
- A CCO ID must be associated with an appropriate Cisco SMARTnet Service contract for your company.
- Cisco SMARTnet Service must be in place for the device to be registered.

The [Cisco support site](#) contains information about the commands to configure Smart Call Home.

Install hardware

Configure the Cisco Nexus 3132Q-V switch

Follow this procedure to configure the Cisco Nexus 3132Q-V switch.

Before you begin

- Access to an HTTP, FTP or TFTP server at the installation site to download the applicable NX-OS and reference configuration file (RCF) releases.
- Applicable NX-OS version, downloaded from the [Cisco software download](#) page.
- Required network switch documentation, controller documentation, and ONTAP documentation. For more information, see [Required documentation](#).
- Applicable licenses, network and configuration information, and cables.
- Completed cabling worksheets. See [Complete Cisco Nexus 3132Q-V cabling worksheet](#).
- Applicable NetApp cluster network and management network RCFs, downloaded from the NetApp Support Site at mysupport.netapp.com for the switches that you receive. All Cisco cluster network and management network switches arrive with the standard Cisco factory-default configuration. These switches also have the current version of the NX-OS software, but do not have the RCFs loaded.

Steps


1. Rack the cluster network and management network switches and controllers.


If you are installing your...	Then...
Cisco Nexus 3132Q-V in a NetApp system cabinet	See the <i>Installing a Cisco Nexus 3132Q-V cluster switch and pass-through panel in a NetApp cabinet</i> guide for instructions to install the switch in a NetApp cabinet.
Equipment in a Telco rack	See the procedures provided in the switch hardware installation guides and the NetApp installation and setup instructions.

2. Cable the cluster network and management network switches to the controllers using the completed cabling worksheet, as described in [Complete Cisco Nexus 3132Q-V cabling worksheet](#).
3. Power on the cluster network and management network switches and controllers.
4. Perform an initial configuration of the cluster network switches.

Provide applicable responses to the following initial setup questions when you first boot the switch. Your site's security policy defines the responses and services to enable.

Prompt	Response
Abort Auto Provisioning and continue with normal setup? (yes/no)	Respond with yes . The default is no.
Do you want to enforce secure password standard? (yes/no)	Respond with yes . The default is yes.
Enter the password for admin:	The default password is "admin"; you must create a new, strong password. A weak password can be rejected.
Would you like to enter the basic configuration dialog? (yes/no)	Respond with yes at the initial configuration of the switch.

Prompt	Response
Create another login account? (yes/no)	Your answer depends on your site's policies on alternate administrators. The default is no .
Configure read-only SNMP community string? (yes/no)	Respond with no . The default is no.
Configure read-write SNMP community string? (yes/no)	Respond with no . The default is no.
Enter the switch name.	The switch name is limited to 63 alphanumeric characters.
Continue with Out-of-band (mgmt0) management configuration? (yes/no)	Respond with yes (the default) at that prompt. At the mgmt0 IPv4 address: prompt, enter your IP address: ip_address.
Configure the default-gateway? (yes/no)	Respond with yes . At the IPv4 address of the default-gateway: prompt, enter your default_gateway.
Configure advanced IP options? (yes/no)	Respond with no . The default is no.
Enable the telnet service? (yes/no)	Respond with no . The default is no.
Enabled SSH service? (yes/no)	Respond with yes . The default is yes. <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-left: 20px;">  SSH is recommended when using Ethernet Switch Health Monitor (CSHM) for its log collection features. SSHv2 is also recommended for enhanced security. </div>
Enter the type of SSH key you want to generate (dsa/rsa/rsa1).	The default is rsa .
Enter the number of key bits (1024-2048).	Enter the key bits from 1024-2048.
Configure the NTP server? (yes/no)	Respond with no . The default is no.
Configure default interface layer (L3/L2):	Respond with L2 . The default is L2.
Configure default switch port interface state (shut/noshut):	Respond with noshut . The default is noshut.

Prompt	Response
Configure CoPP system profile (strict/moderate/lenient/dense):	Respond with strict . The default is strict.
Would you like to edit the configuration? (yes/no)	You should see the new configuration at this point. Review and make any necessary changes to the configuration you just entered. Respond with no at the prompt if you are satisfied with the configuration. Respond with yes if you want to edit your configuration settings.
Use this configuration and save it? (yes/no)	Respond with yes to save the configuration. This automatically updates the kickstart and system images. <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  If you do not save the configuration at this stage, none of the changes will be in effect the next time you reboot the switch. </div>

5. Verify the configuration choices you made in the display that appears at the end of the setup, and make sure that you save the configuration.
6. Check the version on the cluster network switches, and if necessary, download the NetApp-supported version of the software to the switches from the [Cisco software download](#) page.

What's next?

[Prepare to install NX-OS and RCF.](#)

Install a Cisco Nexus 3132Q-V cluster switch in a NetApp cabinet

Depending on your configuration, you might need to install the Cisco Nexus 3132Q-V switch and pass-through panel in a NetApp cabinet with the standard brackets that are included with the switch.

Before you begin

- The initial preparation requirements, kit contents, and safety precautions in the [Cisco Nexus 3000 Series Hardware Installation Guide](#). Review these documents before you begin the procedure.
- The pass-through panel kit, available from NetApp (part number X8784-R6). The NetApp pass-through panel kit contains the following hardware:
 - One pass-through blanking panel
 - Four 10-32 x .75 screws
 - Four 10-32 clip nuts
- Eight 10-32 or 12-24 screws and clip nuts to mount the brackets and slider rails to the front and rear cabinet posts.
- Cisco standard rail kit to install the switch in a NetApp cabinet.




The jumper cords are not included with the pass-through kit and should be included with your switches. If they were not shipped with the switches, you can order them from NetApp (part number X1558A-R6).

Steps

1. Install the pass-through blanking panel in the NetApp cabinet.
 - a. Determine the vertical location of the switches and blanking panel in the cabinet.


In this procedure, the blanking panel will be installed in U40.

- b. Install two clip nuts on each side in the appropriate square holes for front cabinet rails.
- c. Center the panel vertically to prevent intrusion into adjacent rack space, and then tighten the screws.
- d. Insert the female connectors of both 48-inch jumper cords from the rear of the panel and through the brush assembly.




(1) Female connector of the jumper cord.

2. Install the rack-mount brackets on the Nexus 3132Q-V 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.
3. Install the clip nuts in the square hole locations for all four IEA posts.



The two 3132Q-V switches will always be mounted in the top 2U of the cabinet RU41 and 42.

4. Install the slider rails in the cabinet.

- 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 cabinet posts.

- b. Repeat step 4a for the right side rear post.


c. Repeat steps 4a and 4b at the RU41 locations on the cabinet.

5. Install the switch in the cabinet.



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) As the chassis is pushed toward the rear posts, align the two rear rack-mount guides with the slider rails.

(2) Gently slide the switch until the front rack-mount brackets are flush with the front posts.

b. Attach the switch to the cabinet.



(1) With one person holding the front of the chassis level, the other person should fully tighten the four rear screws to the cabinet 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.



By using the fully installed switch as a support, you do not need to hold the front of the second switch during the installation process.

6. When the switches are installed, connect the jumper cords to the switch power inlets.

7. Connect the male plugs of both jumper cords to the closest available PDU outlets.



To maintain redundancy, the two cords must be connected to different PDUs.

8. Connect the management port on each 3132Q-V 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.

Review cabling and configuration considerations

Before configuring your Cisco 3132Q-V switch, review the following considerations.

Support for NVIDIA CX6, CX6-DX, and CX7 Ethernet ports

If connecting a switch port to an ONTAP controller using NVIDIA ConnectX-6 (CX6), ConnectX-6 Dx (CX6-DX), or ConnectX-7 (CX7) NIC ports, you must hard-code the switch port speed.

```
(cs1)(config)# interface Ethernet1/19
(cs1)(config-if)# speed 40000
(cs1)(config-if)# no negotiate auto
(cs1)(config-if)# exit
(cs1)(config)# exit
Save the changes:
(cs1)# copy running-config startup-config
```

See the [Hardware Universe](#) for more information on switch ports.

Configure software

Prepare to install NX-OS software and Reference Configuration File

Before you install the NX-OS software and the Reference Configuration File (RCF), follow this procedure.

About the examples

The examples in this procedure use two nodes. These nodes use two 10GbE cluster interconnect ports e0a and e0b.

See the [Hardware Universe](#) to verify the correct cluster ports on your platforms.



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.

About this task

The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

Steps

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=xh
```

where *x* is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Change the privilege level to advanced, entering *y* when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (**>*) appears.

3. Display how many cluster interconnect interfaces are configured in each node for each cluster interconnect switch:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
```

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	
cluster1-02/cdp	e0a	cs1	Eth1/2	N3K-
C3132Q-V	e0b	cs2	Eth1/2	N3K-
C3132Q-V				
cluster1-01/cdp	e0a	cs1	Eth1/1	N3K-
C3132Q-V	e0b	cs2	Eth1/1	N3K-
C3132Q-V				

4. Check the administrative or operational status of each cluster interface.

a. Display the network port attributes:

```
network port show -ip-space Cluster
```


Show example

```
cluster1::*> network port show -ipSpace Cluster

Node: cluster1-02

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

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
```

b. Display information about the LIFs:

```
network interface show -vserver Cluster
```

Show example

```
cluster1::*> network interface show -vserver Cluster
```

Current Vserver Port	Logical Current Interface Home	Is	Status Admin/Oper	Network Address/Mask	Node
Cluster	cluster1-01	cluster1-01_clus1	up/up	169.254.209.69/16	
	e0a	true			
cluster1-01	cluster1-01	cluster1-01_clus2	up/up	169.254.49.125/16	
	e0b	true			
cluster1-01	cluster1-02	cluster1-02_clus1	up/up	169.254.47.194/16	
	e0a	true			
cluster1-02	cluster1-02	cluster1-02_clus2	up/up	169.254.19.183/16	
	e0b	true			

5. Verify the connectivity of the remote cluster interfaces:


```

cluster1::*> cluster ping-cluster -node local
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)

```

6. Verify that the auto-revert command is enabled on all cluster LIFs:

```
network interface show -vserver Cluster -fields auto-revert
```

Show example

```

cluster1::*> network interface show -vserver Cluster -fields auto-
revert

```

Vserver	Logical Interface	Auto-revert
Cluster	cluster1-01_clus1	true
	cluster1-01_clus2	true
	cluster1-02_clus1	true
	cluster1-02_clus2	true

What's next?

[Install NX-OS software.](#)

Install the NX-OS software

Follow this procedure to install the NX-OS software on the Nexus 3132Q-V cluster switch.

Review requirements

Before you begin

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).

Suggested documentation

- [Cisco Ethernet switch](#). Consult the switch compatibility table for the supported ONTAP and NX-OS versions.
- [Cisco Nexus 3000 Series Switches](#). Consult the appropriate software and upgrade guides available on the Cisco web site for complete documentation on the Cisco switch upgrade and downgrade procedures.

Install the software

About this task

The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

Be sure to complete the procedure in [Prepare to install NX-OS software and Reference Configuration File](#), and then follow the steps below.

Steps

1. Connect the cluster switch to the management network.
2. Use the `ping` command to verify connectivity to the server hosting the NX-OS software and the RCF.

Show example

```
cs2# ping 172.19.2.1 vrf management
Pinging 172.19.2.1 with 0 bytes of data:

Reply From 172.19.2.1: icmp_seq = 0. time= 5910 usec.
```

3. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
              e0a    cs1                      Ethernet1/7      N3K-
C3132Q-V
              e0d    cs2                      Ethernet1/7      N3K-
C3132Q-V
cluster1-02/cdp
              e0a    cs1                      Ethernet1/8      N3K-
C3132Q-V
              e0d    cs2                      Ethernet1/8      N3K-
C3132Q-V
cluster1-03/cdp
              e0a    cs1                      Ethernet1/1/1    N3K-
C3132Q-V
              e0b    cs2                      Ethernet1/1/1    N3K-
C3132Q-V
cluster1-04/cdp
              e0a    cs1                      Ethernet1/1/2    N3K-
C3132Q-V
              e0b    cs2                      Ethernet1/1/2    N3K-
C3132Q-V
cluster1::*>
```

4. 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
```

Show example

```
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
```

```
Health Health Speed (Mbps)
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
```


Show example

```
cluster1::*> network interface show -role Cluster
          Logical          Status      Network
Current   Current Is
Vserver   Interface          Admin/Oper Address/Mask      Node
Port      Home
-----
-----
Cluster
cluster1-01  cluster1-01_clus1  up/up      169.254.3.4/23
              e0a      true
cluster1-01  cluster1-01_clus2  up/up      169.254.3.5/23
              e0d      true
cluster1-02  cluster1-02_clus1  up/up      169.254.3.8/23
              e0a      true
cluster1-02  cluster1-02_clus2  up/up      169.254.3.9/23
              e0d      true
cluster1-03  cluster1-03_clus1  up/up      169.254.1.3/23
              e0a      true
cluster1-03  cluster1-03_clus2  up/up      169.254.1.1/23
              e0b      true
cluster1-04  cluster1-04_clus1  up/up      169.254.1.6/23
              e0a      true
cluster1-04  cluster1-04_clus2  up/up      169.254.1.7/23
              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
```

Show example

```
cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                                     Type                Address
Model
-----
cs1                                         cluster-network    10.233.205.90
N3K-C3132Q-V
  Serial Number: FOCXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                                     9.3(5)
  Version Source: CDP

cs2                                         cluster-network    10.233.205.91
N3K-C3132Q-V
  Serial Number: FOCXXXXXXGS
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                                     9.3(5)
  Version Source: CDP
cluster1::*>
```

5. Disable auto-revert on the cluster LIFs. The cluster LIFs fail over to the partner cluster switch and remain there as you perform the upgrade procedure on the targeted switch:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

6. Copy the NX-OS software to the Nexus 3132Q-V switch using one of the following transfer protocols: FTP, TFTP, SFTP, or SCP. For more information on Cisco commands, see the appropriate guide in [Cisco Nexus 3000 Series NX-OS Command Reference guides](#).

Show example

```
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: xxxxxxxx
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.
```

7. Verify the running version of the NX-OS software:

```
show version
```

Show example

```
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 04.25
  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 Nexus 3132QV Chassis (Nexus 9000 Series)
  Intel(R) Core(TM) i3- CPU @ 2.50GHz with 16399900 kB of memory.
  Processor Board ID FOxxxxxxxx23

  Device name: cs2
  bootflash: 15137792 kB
  usb1: 0 kB (expansion flash)

Kernel uptime is 79 day(s), 10 hour(s), 23 minute(s), 53 second(s)
```

```
Last reset at 663500 usecs after Mon Nov  2 10:50:33 2020
Reason: Reset Requested by CLI command reload
System version: 9.3(3)
Service:

plugin
  Core Plugin, Ethernet Plugin

Active Package(s) :
cs2#
```

8. Install the NX-OS image.

Installing the image file causes it to be loaded every time the switch is rebooted.

Show example

```
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      v04.25(01/28/2020):v04.25(10/18/2016)
v04.25(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#
```

9. Verify the new version of NX-OS software after the switch has rebooted:

```
show version
```

Show example

```
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 04.25
  NXOS: version 9.3(4)
  BIOS compile time: 05/22/2019
  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 Nexus 3132QV Chassis (Nexus 9000 Series)
  Intel(R) Core(TM) i3- CPU @ 2.50GHz with 16399900 kB of memory.
  Processor Board ID FOxxxxxxxx23

  Device name: cs2
  bootflash: 15137792 kB
  usb1: 0 kB (expansion flash)

Kernel uptime is 79 day(s), 10 hour(s), 23 minute(s), 53 second(s)
```



```
Last reset at 663500 usecs after Mon Nov  2 10:50:33 2020
Reason: Reset Requested by CLI command reload
System version: 9.3(4)
Service:

plugin
  Core Plugin, Ethernet Plugin

Active Package(s) :

cs2#
```

10. Verify the health of cluster ports on the cluster.

a. Verify that cluster ports are up and healthy across all nodes in the cluster:

```
network port show -role cluster
```

Show example

```
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
```

```
Health Health Speed (Mbps)
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.

```
network device-discovery show -protocol cdp
```

Show example

```
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-C3132Q-V
          e0d    cs2                      Ethernet1/7
N3K-C3132Q-V
cluster01-2/cdp
          e0a    cs1                      Ethernet1/8
N3K-C3132Q-V
          e0d    cs2                      Ethernet1/8
N3K-C3132Q-V
cluster01-3/cdp
          e0a    cs1                      Ethernet1/1/1
N3K-C3132Q-V
          e0b    cs2                      Ethernet1/1/1
N3K-C3132Q-V
cluster1-04/cdp
          e0a    cs1                      Ethernet1/1/2
N3K-C3132Q-V
          e0b    cs2                      Ethernet1/1/2
N3K-C3132Q-V

cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                Type                Address
Model
-----
-----
cs1                    cluster-network    10.233.205.90
N3K-C3132Q-V
  Serial Number: FOCXXXXXXGD
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                      9.3(5)
  Version Source: CDP

cs2                    cluster-network    10.233.205.91
```

```

N3K-C3132Q-V
  Serial Number: FOCXXXXXXGS
    Is Monitored: true
      Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                    9.3(5)
  Version Source: CDP

2 entries were displayed.

```

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.

```

11. Verify that the cluster is healthy:

```
cluster show
```

Show example

```

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::*>

```

12. Repeat steps 6 to 11 on switch cs1.

13. Enable auto-revert on the cluster LIFs.

```
network interface modify -vserver Cluster -lif * -auto-revert true
```

14. Verify that the cluster LIFs have reverted to their home port:

```
network interface show -role cluster
```

Show example

```
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::*>
```

If any cluster LIFs have not returned to their home ports, revert them manually from the local node:

```
network interface revert -vserver Cluster -lif <lif_name>
```

What's next?

[Install the Reference Configuration File \(RCF\).](#)

Install the Reference Configuration File (RCF)

Follow this procedure to install the RCF after setting up the Nexus 3132Q-V switch for the first time. You can also use this procedure to upgrade your RCF version.

Review requirements

Before you begin

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).
- The current Reference Configuration File (RCF).
- A console connection to the switch is required when installing the RCF. This requirement is optional if you have used the Knowledge Base article [How to clear the configuration on a Cisco interconnect switch while retaining remote connectivity](#) to clear the configuration, beforehand.
- [Cisco Ethernet switch](#). Consult the switch compatibility table for the supported ONTAP and RCF versions. Note that there can be command dependencies between the command syntax in the RCF and that found in versions of NX-OS.
- [Cisco Nexus 3000 Series Switches](#). Consult the appropriate software and upgrade guides available on the Cisco web site for complete documentation on the Cisco switch upgrade and downgrade procedures.

Install the file

About the examples

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.

About this task

The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure migrates all of the cluster LIFs to the operational partner switch while performing the steps on the target switch.

Be sure to complete the procedure in [Prepare to install NX-OS software and Reference Configuration File](#), and then follow the steps below.

Step 1: Check port status

1. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
              e0a    cs1                      Ethernet1/7      N3K-
C3132Q-V
              e0d    cs2                      Ethernet1/7      N3K-
C3132Q-V
cluster1-02/cdp
              e0a    cs1                      Ethernet1/8      N3K-
C3132Q-V
              e0d    cs2                      Ethernet1/8      N3K-
C3132Q-V
cluster1-03/cdp
              e0a    cs1                      Ethernet1/1/1    N3K-
C3132Q-V
              e0b    cs2                      Ethernet1/1/1    N3K-
C3132Q-V
cluster1-04/cdp
              e0a    cs1                      Ethernet1/1/2    N3K-
C3132Q-V
              e0b    cs2                      Ethernet1/1/2    N3K-
C3132Q-V
cluster1::*>
```

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 -ipSpace Cluster
```


Show example

```
cluster1::*> network port show -ipspace Cluster

Node: cluster1-01

Ignore

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

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

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
```

```
Health Health Speed (Mbps)
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 -vserver Cluster
```

Show example

```
cluster1::*> network interface show -vserver Cluster
          Logical          Status      Network
Current   Current Is
Vserver   Interface           Admin/Oper Address/Mask   Node
Port      Home
-----
-----
Cluster
cluster1-01  cluster1-01_clus1  up/up      169.254.3.4/23
            e0a      true
cluster1-01  cluster1-01_clus2  up/up      169.254.3.5/23
            e0d      true
cluster1-02  cluster1-02_clus1  up/up      169.254.3.8/23
            e0a      true
cluster1-02  cluster1-02_clus2  up/up      169.254.3.9/23
            e0d      true
cluster1-03  cluster1-03_clus1  up/up      169.254.1.3/23
            e0a      true
cluster1-03  cluster1-03_clus2  up/up      169.254.1.1/23
            e0b      true
cluster1-04  cluster1-04_clus1  up/up      169.254.1.6/23
            e0a      true
cluster1-04  cluster1-04_clus2  up/up      169.254.1.7/23
            e0b      true
cluster1::*>
```

- c. Verify that the cluster displays information for both cluster switches:

```
system cluster-switch show -is-monitoring-enabled-operational true
```

Show example

```
cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                                     Type                Address
Model
-----
cs1                                       cluster-network    10.0.0.1
NX3132QV
  Serial Number: FOXXXXXXXXGS
  Is Monitored: true
  Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                                     9.3(4)
  Version Source: CDP

cs2                                       cluster-network    10.0.0.2
NX3132QV
  Serial Number: FOXXXXXXXXGD
  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.
```



For ONTAP 9.8 and later, use the command `system switch ethernet show -is-monitoring-enabled-operational true`.

3. Disable auto-revert on the cluster LIFs.

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto
-revert false
```

Make sure that auto-revert is disabled after running this command.

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
```

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 -vserver Cluster
```

Show example

```
cluster1::*> network interface show -vserver Cluster
```

Current Is	Logical	Status	Network	Current
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			

```
cluster1::*>
```

6. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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
cluster1::*>
```

Step 2: Configure and verify the setup

1. If you have not already done so, save a copy of the current switch configuration by copying the output of the following command to a text file:

```
show running-config
```

2. Clean the configuration on switch cs2 and perform a basic setup.



When updating or applying a new RCF, you must erase the switch settings and perform basic configuration. You **must** be connected to the switch serial console port to set up the switch again. However, this requirement is optional if you have used the Knowledge Base article [How to clear the configuration on a Cisco interconnect switch while retaining remote connectivity](#) to clear the configuration, beforehand.

- a. Clean the configuration:

Show example

```
(cs2)# write erase

Warning: This command will erase the startup-configuration.

Do you wish to proceed anyway? (y/n) [n] y
```

- b. Perform a reboot of the switch:

Show example

```
(cs2)# reload

Are you sure you would like to reset the system? (y/n) y
```

3. 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.

Show example

```
cs2# copy tftp: bootflash: vrf management
Enter source filename: Nexus_3132QV_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)...
```

4. 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.

Show example

```
cs2# copy Nexus_3132QV_RCF_v1.6-Cluster-HA-Breakout.txt running-
config echo-commands
```

5. 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.

Show example

```
cs2# show banner motd

*****
*****
* NetApp Reference Configuration File (RCF)
*
* Switch   : Cisco Nexus 3132Q-V
* Filename : Nexus_3132QV_RCF_v1.6-Cluster-HA-Breakout.txt
* Date     : Nov-02-2020
* Version  : v1.6
*
* Port Usage : Breakout configuration
* Ports 1- 6: Breakout mode (4x10GbE) Intra-Cluster Ports, int
e1/1/1-4,
* e1/2/1-4, e1/3/1-4,int e1/4/1-4, e1/5/1-4, e1/6/1-4
* Ports 7-30: 40GbE Intra-Cluster/HA Ports, int e1/7-30
* Ports 31-32: Intra-Cluster ISL Ports, int e1/31-32
*
* IMPORTANT NOTES
* - Load Nexus_3132QV_RCF_v1.6-Cluster-HA.txt for non breakout
config
*
* - This RCF utilizes QoS and requires specific TCAM configuration,
requiring
*   cluster switch to be rebooted before the cluster becomes
operational.
*
* - Perform the following steps to ensure proper RCF installation:
*
*   (1) Apply RCF, 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...
*
*   (2) Save running-configuration and reboot Cluster Switch
*
*   (3) After reboot, apply same RCF second time and expect
following messages:
*       - % Invalid command at '^' marker
*
*   (4) Save running-configuration again
*
```



```

* - If running NX-OS versions 9.3(5) 9.3(6), 9.3(7), or 9.3(8)
*   - Downgrade the NX-OS firmware to version 9.3(5) or earlier if
*     NX-OS using a version later than 9.3(5).
*   - Do not upgrade NX-OS prior to applying v1.9 RCF file.
*   - After the RCF is applied and switch rebooted, then proceed to
upgrade
*     NX-OS to version 9.3(5) or later.
*
* - If running 9.3(9) 10.2(2) or later the RCF can be applied to the
switch
*     after the upgrade.
*
* - Port 1 multiplexed H/W configuration options:
*   hardware profile front portmode qsfp      (40G H/W port 1/1 is
active - default)
*   hardware profile front portmode sfp-plus  (10G H/W ports 1/1/1
- 1/1/4 are active)
*   hardware profile front portmode qsfp      (To reset to QSFP)
*
*****
*****

```

6. 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.



For steps on how to bring your 10GbE ports online after an upgrade of the RCF, see the Knowledge Base article [10GbE ports on a Cisco 3132Q cluster switch do not come online](#).

7. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
8. 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.

Show example

```
cs2# copy running-config startup-config  
[#####] 100% Copy complete
```

9. Reboot switch cs2. You can ignore both the "cluster ports down" events reported on the nodes while the switch reboots and the error % Invalid command at '^' marker output.

Show example

```
cs2# reload  
This command will reboot the system. (y/n)? [n] y
```

10. Apply the same RCF and save the running configuration for a second time. This is necessary as the RCF utilizes QoS and requires TCAM re-configuration that involves loading the RCF twice with the switch rebooted in between.

Show example

```
cs2# copy Nexus_3132QV_RCF_v1.6-Cluster-HA-Breakout.txt running-  
config echo-commands  
cs2# copy running-config startup-config  
[#####] 100% Copy complete
```

11. Verify the health of cluster ports on the cluster.
 - a. Verify that cluster ports are up and healthy across all nodes in the cluster:

```
network port show -ipSpace Cluster
```

Show example

```
cluster1::*> network port show -ipspace Cluster

Node: cluster1-01

Ignore

Health      Health      Speed (Mbps)
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

Health      Health      Speed (Mbps)
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

Health      Health      Speed (Mbps)
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
```

```
Health Health Speed (Mbps)
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
```

b. Verify the switch health from the cluster.

```
network device-discovery show -protocol cdp
```

Show example

```
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-C3132Q-V
          e0d    cs2                        Ethernet1/7
N3K-C3132Q-V
cluster01-2/cdp
          e0a    cs1                        Ethernet1/8
N3K-C3132Q-V
          e0d    cs2                        Ethernet1/8
N3K-C3132Q-V
cluster01-3/cdp
          e0a    cs1                        Ethernet1/1/1
N3K-C3132Q-V
          e0b    cs2                        Ethernet1/1/1
N3K-C3132Q-V
cluster1-04/cdp
          e0a    cs1                        Ethernet1/1/2
N3K-C3132Q-V
          e0b    cs2                        Ethernet1/1/2
N3K-C3132Q-V

cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                Type                Address
Model
-----
-----
cs1                    cluster-network    10.233.205.90
N3K-C3132Q-V
  Serial Number: FOXXXXXXXXGD
  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-C3132Q-V
  Serial Number: FOXXXXXXXXGS
    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.

```



For ONTAP 9.8 and later, use the command `system switch ethernet show -is -monitoring-enabled-operational true`.

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.

```



It can take up to 5 minutes for the cluster nodes to report as healthy.

- On cluster switch cs1, shut down the ports connected to the cluster ports of the nodes.

Show example

```

cs1(config)# interface eth1/1/1-2,eth1/7-8
cs1(config-if-range)# shutdown

```

- Verify that the cluster LIFs have migrated to the ports hosted on switch cs2. This might take a few seconds.

```

network interface show -vserver Cluster

```

Show example

```
cluster1::*> network interface show -vserver Cluster
          Logical          Status      Network          Current
Current Is
Vserver   Interface              Admin/Oper  Address/Mask     Node
Port      Home
-----
Cluster
cluster1-01 cluster1-01_clus1 up/up      169.254.3.4/23
          e0d             false
cluster1-01 cluster1-01_clus2 up/up      169.254.3.5/23
          e0d             true
cluster1-02 cluster1-02_clus1 up/up      169.254.3.8/23
          e0d             false
cluster1-02 cluster1-02_clus2 up/up      169.254.3.9/23
          e0d             true
cluster1-03 cluster1-03_clus1 up/up      169.254.1.3/23
          e0b             false
cluster1-03 cluster1-03_clus2 up/up      169.254.1.1/23
          e0b             true
cluster1-04 cluster1-04_clus1 up/up      169.254.1.6/23
          e0b             false
cluster1-04 cluster1-04_clus2 up/up      169.254.1.7/23
          e0b             true
cluster1::*>
```

14. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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::*>
```

- Repeat Steps 1 to 11 on switch cs1.
- Enable auto-revert on the cluster LIFs.

Show example

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto
-revert True
```

- 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 3: Verify the configuration

- Verify that the switch ports connected to the cluster ports are up.

```
show interface brief | grep up
```

Show example

```
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) --
.
.
```

- Verify that the ISL between cs1 and cs2 is functional:

```
show port-channel summary
```


Show example

```
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#
```

3. Verify that the cluster LIFs have reverted to their home port:

```
network interface show -vserver Cluster
```

Show example

```
cluster1::*> network interface show -vserver Cluster
          Logical          Status      Network          Current
Current Is
Vserver   Interface              Admin/Oper  Address/Mask     Node
Port      Home
-----
Cluster
cluster1-01 cluster1-01_clus1 up/up      169.254.3.4/23
          e0d             true
cluster1-01 cluster1-01_clus2 up/up      169.254.3.5/23
          e0d             true
cluster1-02 cluster1-02_clus1 up/up      169.254.3.8/23
          e0d             true
cluster1-02 cluster1-02_clus2 up/up      169.254.3.9/23
          e0d             true
cluster1-03 cluster1-03_clus1 up/up      169.254.1.3/23
          e0b             true
cluster1-03 cluster1-03_clus2 up/up      169.254.1.1/23
          e0b             true
cluster1-04 cluster1-04_clus1 up/up      169.254.1.6/23
          e0b             true
cluster1-04 cluster1-04_clus2 up/up      169.254.1.7/23
          e0b             true
cluster1::*>
```

4. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
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
cluster1::*>
```

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
-----	-----	
-----	-----	
cluster1-01		
3/5/2022 19:21:18 -06:00	cluster1-01_clus2	cluster1-02_clus1
none		
3/5/2022 19:21:20 -06:00	cluster1-01_clus2	cluster1-02_clus2
none		
cluster1-02		
3/5/2022 19:21:18 -06:00	cluster1-02_clus2	cluster1-01_clus1
none		
3/5/2022 19:21:20 -06:00	cluster1-02_clus2	cluster1-01_clus2
none		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is 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)

```

What's next?

[Verify SSH configuration.](#)

Verify your SSH configuration

If you are using the Ethernet Switch Health Monitor (CSHM) and log collection features, verify that SSH and SSH keys are enabled on the cluster switches.

Steps

1. Verify that SSH is enabled:

```

(switch) show ssh server
ssh version 2 is enabled

```

2. Verify that the SSH keys are enabled:

```
show ssh key
```

Show example

```
(switch)# show ssh key

rsa Keys generated:Fri Jun 28 02:16:00 2024

ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQGDINrD52Q586wTGJjFABjBlFaA23EpDrZ2sDCew
l7nwlIoC6HBejxluIObAH8hrW8kR+gj0ZAfPpNeLGTg3APj/yIPTBoIZZxbWRShywAM5
PqyxWwRb7kp9Zt1YHzVuHYpSO82KUDowKrL6lox/YtpKoZUDZjrZjAp8hTv3JZsPgQ==

bitcount:1024
fingerprint:
SHA256:aHwhpzo7+YCDsrp3isJv2uVGz+mjMMokqdMeXVVXfdo

could not retrieve dsa key information

ecdsa Keys generated:Fri Jun 28 02:30:56 2024

ecdsa-sha2-nistp521
AAAAE2VjZHNhLXNoYTItbmlzdHA1MjEAAAABmlzdHA1MjEAAACFBABJ+ZX5SFKhS57e
vkE273e0VoqZi4/32dt+f14fBuKv80MjMsmLfjKtCWylwgVt1Zi+C5TIBbugpzez529z
kFSF0ADb8JaGCoaAYe2HvWR/f6QLbKbqVIewCdqWgxzrIY5BPP5GBdxQJMBiOwEdnHg1
u/9Pzh/Vz9cHDcCW9qGE780QHA==

bitcount:521
fingerprint:
SHA256:TFGe2hXn6QIpcs/vyHzftHJ7Dceg0vQaULYRALZeHwQ

(switch)# show feature | include scpServer
scpServer          1          enabled
(switch)# show feature | include ssh
sshServer          1          enabled
(switch)#
```



When enabling FIPS, you must change the bitcount to 256 on the switch using the command `ssh key ecdsa 256 force`. See [Configure network security using FIPS](#) for more details.

What's next?

[Configure switch health monitoring.](#)

Migrate switches

CN1610 switch migration

Migrate from CN1610 switches to Nexus 3132Q-V switches workflow

Follow these workflow steps to migrate your CN1610 switches to Cisco Nexus 3132Q-V switches.

1

Migration requirements

Review the requirements and example switch information for the migration process.

2

Prepare for migration

Prepare your CN1610 switches for migration to Nexus 3132Q-V switches.

3

Configure your ports

Configure your ports for migration to the new Nexus 3132Q-V switches.

4

Complete your migration

Complete your migration to the new Nexus 3132Q-V switches.

Migration requirements

Cisco Nexus 3132Q-V switches can be used as cluster switches in your AFF or FAS cluster. Cluster switches allow you to build ONTAP clusters with more than two nodes.



The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

For more information, see:

- [NetApp CN1601 and CN1610 description page](#)
- [Cisco Ethernet Switch description page](#)
- [Hardware Universe](#)

CN1610 requirements

The following cluster switches are supported:

- NetApp CN1610
- Cisco Nexus 3132Q-V

The cluster switches support the following node connections:

- NetApp CN1610: ports 0/1 through 0/12 (10 GbE)
- Cisco Nexus 3132Q-V: ports e1/1-30 (40/100 GbE)

The cluster switches use the following inter-switch link (ISL) ports:

- NetApp CN1610: ports 0/13 through 0/16 (10 GbE)
- Cisco Nexus 3132Q-V: ports e1/31-32 (40/100 GbE)

The [Hardware Universe](#) contains information about supported cabling to Nexus 3132Q-V switches:

- Nodes with 10 GbE cluster connections require QSFP to SFP+ optical fiber breakout cables or QSFP to SFP+ copper breakout cables
- Nodes with 40/100 GbE cluster connections require supported QSFP/QSFP28 optical modules with optical fiber cables or QSFP/QSFP28 copper direct-attach cables

The appropriate ISL cabling is as follows:

- Beginning: For CN1610 to CN1610 (SFP+ to SFP+), four SFP+ optical fiber or copper direct-attach cables
- Interim: For CN1610 to Nexus 3132Q-V (QSFP to four SFP+ breakout), one QSFP to SFP+ optical fiber or copper breakout cable
- Final: For Nexus 3132Q-V to Nexus 3132Q-V (QSFP28 to QSFP28), two QSFP28 optical fiber or copper direct-attach cables

NetApp twinax cables are not compatible with Cisco Nexus 3132Q-V switches.

If your current CN1610 configuration uses NetApp twinax cables for cluster-node-to-switch connections or ISL connections and you want to continue using twinax in your environment, you need to procure Cisco twinax cables. Alternatively, you can use optical fiber cables for both the ISL connections and the cluster-node-to-switch connections.

About the examples used

The examples in this procedure describe replacing CN1610 switches with Cisco Nexus 3132Q-V switches. You can use these steps (with modifications) for other older Cisco switches.

The procedure also uses the following switch and node nomenclature:

- The command outputs might vary depending on different releases of ONTAP.
- The CN1610 switches to be replaced are **CL1** and **CL2**.
- The Nexus 3132Q-V switches to replace the CN1610 switches are **C1** and **C2**.
- **n1_clus1** is the first cluster logical interface (LIF) connected to cluster switch 1 (CL1 or C1) for node **n1**.
- **n1_clus2** is the first cluster LIF connected to cluster switch 2 (CL2 or C2) for node **n1**.
- **n1_clus3** is the second LIF connected to cluster switch 2 (CL2 or C2) for node **n1**.
- **n1_clus4** is the second LIF connected to cluster switch 1 (CL1 or C1) for node **n1**.
- The number of 10 GbE and 40/100 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.
- The nodes are **n1**, **n2**, **n3**, and **n4**.

The examples in this procedure use four nodes:

- Two nodes use four 10 GbE cluster interconnect ports: **e0a**, **e0b**, **e0c**, and **e0d**.
- The other two nodes use two 40 GbE cluster interconnect ports: **e4a** and **e4e**.

The [Hardware Universe](#) lists the actual cluster ports on your platforms.

What's next?

[Prepare for migration.](#)

Prepare for migration from CN1610 switches to 3132Q-V switches

Follow these steps to prepare your CN1610 switches for migration to Cisco Nexus 3132Q-V switches.

Steps

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all - message MAINT=xh
```

x is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Display information about the devices in your configuration:

```
network device-discovery show
```

Show example

The following example displays how many cluster interconnect interfaces have been configured in each node for each cluster interconnect switch:

```
cluster::> network device-discovery show
```

Node	Local Port	Discovered Device	Interface	Platform
n1	/cdp			
	e0a	CL1	0/1	CN1610
	e0b	CL2	0/1	CN1610
	e0c	CL2	0/2	CN1610
	e0d	CL1	0/2	CN1610
n2	/cdp			
	e0a	CL1	0/3	CN1610
	e0b	CL2	0/3	CN1610
	e0c	CL2	0/4	CN1610
	e0d	CL1	0/4	CN1610

8 entries were displayed.

3. Determine the administrative or operational status for each cluster interface.

a. Display the cluster network port attributes:

```
network port show
```

Show example

The following example displays the network port attributes on a system:

```
cluster::*> network port show -role Cluster
(network port show)

Node: n1

```

Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Open	Health Status	Ignore Health
e0a	cluster	cluster	up	9000	auto/10000	-	-
e0b	cluster	cluster	up	9000	auto/10000	-	-
e0c	cluster	cluster	up	9000	auto/10000	-	-
e0d	cluster	cluster	up	9000	auto/10000	-	-

```
Node: n2

```

Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Open	Health Status	Ignore Health
e0a	cluster	cluster	up	9000	auto/10000	-	-
e0b	cluster	cluster	up	9000	auto/10000	-	-
e0c	cluster	cluster	up	9000	auto/10000	-	-
e0d	cluster	cluster	up	9000	auto/10000	-	-

```
8 entries were displayed.
```

b. Display information about the logical interfaces:

```
+
network interface show
```

Show example

The following example displays the general information about all of the LIFs on your system:

```
cluster::*> network interface show -role Cluster
(network interface show)

      Logical      Status      Network      Current      Current
Is      Interface  Admin/Oper  Address/Mask  Node         Port
Vserver
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

8 entries were displayed.
```

c. Display information about the discovered cluster switches:

```
system cluster-switch show
```

Show example

The following example displays the cluster switches that are known to the cluster, along with their management IP addresses:

```
cluster::> system cluster-switch show

Switch                               Type                               Address
Model
-----                               -----
CL1                                  cluster-network                    10.10.1.101
CN1610
    Serial Number: 01234567
    Is Monitored: true
    Reason:
    Software Version: 1.2.0.7
    Version Source: ISDP

CL2                                  cluster-network                    10.10.1.102
CN1610
    Serial Number: 01234568
    Is Monitored: true
    Reason:
    Software Version: 1.2.0.7
    Version Source: ISDP

2 entries were displayed.
```

4. Set the `-auto-revert` parameter to `false` on cluster LIFs `clus1` and `clus4` on both nodes:

```
network interface modify
```

Show example

```
cluster::*> network interface modify -vserver node1 -lif clus1 -auto
-revert false
cluster::*> network interface modify -vserver node1 -lif clus4 -auto
-revert false
cluster::*> network interface modify -vserver node2 -lif clus1 -auto
-revert false
cluster::*> network interface modify -vserver node2 -lif clus4 -auto
-revert false
```

5. Verify that the appropriate RCF and image are installed on the new 3132Q-V switches as necessary for your requirements, and make any essential site customizations, such as users and passwords, network addresses, and so on.

You must prepare both switches at this time. If you need to upgrade the RCF and image, follow these steps:

- a. See the [Cisco Ethernet Switches](#) page on NetApp Support Site.
- b. Note your switch and the required software versions in the table on that page.
- c. Download the appropriate version of the RCF.
- d. Select **CONTINUE** on the **Description** page, accept the license agreement, and then follow the instructions on the **Download** page to download the RCF.
- e. Download the appropriate version of the image software.

[Cisco® Cluster and Management Network Switch Reference Configuration File Download](#)

6. Migrate the LIFs associated with the second CN1610 switch to be replaced:

```
network interface migrate
```



You must migrate the cluster LIFs from a connection to the node, either through the service processor or node management interface, which owns the cluster LIF being migrated.

Show example

The following example shows n1 and n2, but LIF migration must be done on all the nodes:

```
cluster::*> network interface migrate -vserver Cluster -lif n1_clus2
-destination-node n1 -destination-port e0a
cluster::*> network interface migrate -vserver Cluster -lif n1_clus3
-destination-node n1 -destination-port e0d
cluster::*> network interface migrate -vserver Cluster -lif n2_clus2
-destination-node n2 -destination-port e0a
cluster::*> network interface migrate -vserver Cluster -lif n2_clus3
-destination-node n2 -destination-port e0d
```

7. Verify the cluster's health:

```
network interface show
```

Show example

The following example shows the result of the previous `network interface migrate` command:

```
cluster::*> network interface show -role Cluster
(network interface show)
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
-----	-----	-----	-----	-----	-----	-----
Cluster						
true	n1_clus1	up/up	10.10.0.1/24	n1	e0a	
false	n1_clus2	up/up	10.10.0.2/24	n1	e0a	
false	n1_clus3	up/up	10.10.0.3/24	n1	e0d	
true	n1_clus4	up/up	10.10.0.4/24	n1	e0d	
true	n2_clus1	up/up	10.10.0.5/24	n2	e0a	
false	n2_clus2	up/up	10.10.0.6/24	n2	e0a	
false	n2_clus3	up/up	10.10.0.7/24	n2	e0d	
true	n2_clus4	up/up	10.10.0.8/24	n2	e0d	

8 entries were displayed.

8. Shut down the cluster interconnect ports that are physically connected to switch CL2:

```
network port modify
```

Show example

The following commands shut down the specified ports on n1 and n2, but the ports must be 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
```

9. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Node	Date		Source LIF	Destination LIF	Packet Loss
n1					
	3/5/2022 19:21:18 -06:00		n1_clus2	n1_clus1	none
	3/5/2022 19:21:20 -06:00		n1_clus2	n2_clus2	none
n2					
	3/5/2022 19:21:18 -06:00		n2_clus2	n1_clus1	none
	3/5/2022 19:21:20 -06:00		n2_clus2	n1_clus2	none

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```



```

cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1      e0a    10.10.0.1
Cluster n1_clus2 n1      e0b    10.10.0.2
Cluster n1_clus3 n1      e0c    10.10.0.3
Cluster n1_clus4 n1      e0d    10.10.0.4
Cluster n2_clus1 n2      e0a    10.10.0.5
Cluster n2_clus2 n2      e0b    10.10.0.6
Cluster n2_clus3 n2      e0c    10.10.0.7
Cluster n2_clus4 n2      e0d    10.10.0.8

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
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 16 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 10.10.0.2 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 10.10.0.3 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 10.10.0.4 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

Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
4 paths up, 0 paths down (udp check)

```

10. Shut down the ISL ports 13 through 16 on the active CN1610 switch CL1:

shutdown

Show example

The following example shows how to shut down ISL ports 13 through 16 on the CN1610 switch CL1:

```
(CL1)# configure
(CL1)(Config)# interface 0/13-0/16
(CL1)(Interface 0/13-0/16)# shutdown
(CL1)(Interface 0/13-0/16)# exit
(CL1)(Config)# exit
(CL1)#
```

11. Build a temporary ISL between CL1 and C2:

Show example

The following example builds a temporary ISL between CL1 (ports 13-16) and C2 (ports e1/24/1-4):

```
C2# configure
C2(config)# interface port-channel 2
C2(config-if)# switchport mode trunk
C2(config-if)# spanning-tree port type network
C2(config-if)# mtu 9216
C2(config-if)# interface breakout module 1 port 24 map 10g-4x
C2(config)# interface e1/24/1-4
C2(config-if-range)# switchport mode trunk
C2(config-if-range)# mtu 9216
C2(config-if-range)# channel-group 2 mode active
C2(config-if-range)# exit
C2(config-if)# exit
```

What's next?

[Configure your ports.](#)

Configure your ports for migration from CN1610 switches to 3132Q-V switches

Follow these steps to configure your ports for migration from the CN1610 switches to the new Nexus 3132Q-V switches.

Steps

1. On all nodes, remove the cables that are attached to the CN1610 switch CL2.

With supported cabling, you must reconnect the disconnected ports on all of the nodes to the Nexus 3132Q-V switch C2.

2. Remove four ISL cables from ports 13 to 16 on the CN1610 switch CL1.

You must attach appropriate Cisco QSFP to SFP+ breakout cables connecting port 1/24 on the new Cisco 3132Q-V switch C2, to ports 13 to 16 on existing CN1610 switch CL1.



When reconnecting any cables to the new Cisco 3132Q-V switch, you must use either optical fiber or Cisco twinax cables.

3. To make the ISL dynamic, configure the ISL interface 3/1 on the active CN1610 switch to disable the static mode:
`no port-channel static`

This configuration matches with the ISL configuration on the 3132Q-V switch C2 when the ISLs are brought up on both switches in step 11.

Show example

The following example shows the configuration of the ISL interface 3/1 using the `no port-channel static` command to make the ISL dynamic:

```
(CL1)# configure
(CL1)(Config)# interface 3/1
(CL1)(Interface 3/1)# no port-channel static
(CL1)(Interface 3/1)# exit
(CL1)(Config)# exit
(CL1)#
```

4. Bring up ISLs 13 through 16 on the active CN1610 switch CL1.

Show example

The following example illustrates the process of bringing up ISL ports 13 through 16 on the port-channel interface 3/1:

```
(CL1)# configure
(CL1)(Config)# interface 0/13-0/16,3/1
(CL1)(Interface 0/13-0/16,3/1)# no shutdown
(CL1)(Interface 0/13-0/16,3/1)# exit
(CL1)(Config)# exit
(CL1)#
```

5. Verify that the ISLs are up on the CN1610 switch CL1:

```
show port-channel
```

The "Link State" should be Up, "Type" should be Dynamic, and the "Port Active" column should be True for ports 0/13 to 0/16:

Show example

```
(CL1)# show port-channel 3/1
Local Interface..... 3/1
Channel Name..... ISL-LAG
Link State..... Up
Admin Mode..... Enabled
Type..... Dynamic
Load Balance Option..... 7
(Enhanced hashing mode)
```

Mbr	Device/	Port	Port
Ports	Timeout	Speed	Active
-----	-----	-----	-----
0/13	actor/long partner/long	10 Gb Full	True
0/14	actor/long partner/long	10 Gb Full	True
0/15	actor/long partner/long	10 Gb Full	True
0/16	actor/long partner/long	10 Gb Full	True

6. Verify that the ISLs are up on the 3132Q-V switch C2:

```
show port-channel summary
```

Show example

Ports Eth1/24/1 through Eth1/24/4 should indicate (P), meaning that all four ISL ports are up in the port-channel. Eth1/31 and Eth1/32 should indicate (D) as they are not connected:

```
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 (D)  Eth1/32 (D)
2      Po2 (SU)       Eth       LACP      Eth1/24/1 (P) Eth1/24/2 (P)
Eth1/24/3 (P)
                                   Eth1/24/4 (P)
```

7. Bring up all of the cluster interconnect ports that are connected to the 3132Q-V switch C2 on all of the nodes:

```
network port modify
```

Show example

The following example shows how to bring up the cluster interconnect ports connected to the 3132Q-V switch C2:

```
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
```

8. Revert all of the migrated cluster interconnect LIFs that are connected to C2 on all of the nodes:

```
network interface revert
```

Show example

```
cluster::*> network interface revert -vserver Cluster -lif n1_clus2
cluster::*> network interface revert -vserver Cluster -lif n1_clus3
cluster::*> network interface revert -vserver Cluster -lif n2_clus2
cluster::*> network interface revert -vserver Cluster -lif n2_clus3
```

9. Verify that all of the cluster interconnect ports are reverted to their home ports:

```
network interface show
```

Show example

The following example shows that the LIFs on clus2 are reverted to their home ports, and shows that the LIFs are successfully reverted if the ports in the "Current Port" column have a status of `true` in the "Is Home" column. If the Is Home value is `false`, then the LIF is not reverted.

```
cluster::*> network interface show -role cluster
(network interface show)
```

Vserver Home	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is
-----	-----	-----	-----	-----	-----	-----
Cluster						
true	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	

8 entries were displayed.

10. Verify that all of the cluster ports are connected:

```
network port show
```

Show example

The following example shows the result of the previous `network port modify` command, verifying that all of the cluster interconnects are up:

```
cluster::*> network port show -role Cluster
(network port show)

Node: n1

Port IPspace Broadcast Link MTU Speed (Mbps) Health Ignore
Status Domain Health
-----
-----
e0a cluster cluster up 9000 auto/10000 - -
e0b cluster cluster up 9000 auto/10000 - -
e0c cluster cluster up 9000 auto/10000 - -
e0d cluster cluster up 9000 auto/10000 - -

Node: n2

Port IPspace Broadcast Link MTU Speed (Mbps) Health Ignore
Status Domain Health
-----
-----
e0a cluster cluster up 9000 auto/10000 - -
e0b cluster cluster up 9000 auto/10000 - -
e0c cluster cluster up 9000 auto/10000 - -
e0d cluster cluster up 9000 auto/10000 - -

8 entries were displayed.
```

11. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Node	Date		Source LIF	Destination LIF	Packet Loss
n1					
	3/5/2022 19:21:18 -06:00		n1_clus2	n1_clus1	none
	3/5/2022 19:21:20 -06:00		n1_clus2	n2_clus2	none
n2					
	3/5/2022 19:21:18 -06:00		n2_clus2	n1_clus1	none
	3/5/2022 19:21:20 -06:00		n2_clus2	n1_clus2	none

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1      e0a    10.10.0.1
Cluster n1_clus2 n1      e0b    10.10.0.2
Cluster n1_clus3 n1      e0c    10.10.0.3
Cluster n1_clus4 n1      e0d    10.10.0.4
Cluster n2_clus1 n2      e0a    10.10.0.5
Cluster n2_clus2 n2      e0b    10.10.0.6
Cluster n2_clus3 n2      e0c    10.10.0.7
Cluster n2_clus4 n2      e0d    10.10.0.8

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
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 16 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 10.10.0.2 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 10.10.0.3 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 10.10.0.4 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

Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
4 paths up, 0 paths down (udp check)

```

12. On each node in the cluster, migrate the interfaces that are associated with the first CN1610 switch CL1,

to be replaced:

```
network interface migrate
```

Show example

The following example shows the ports or LIFs being migrated on nodes n1 and n2:

```
cluster::*> network interface migrate -vserver Cluster -lif n1_clus1
-destination-node n1 -destination-port e0b
cluster::*> network interface migrate -vserver Cluster -lif n1_clus4
-destination-node n1 -destination-port e0c
cluster::*> network interface migrate -vserver Cluster -lif n2_clus1
-destination-node n2 -destination-port e0b
cluster::*> network interface migrate -vserver Cluster -lif n2_clus4
-destination-node n2 -destination-port e0c
```

13. Verify the cluster status:

```
network interface show
```

Show example

The following example shows that the required cluster LIFs have been migrated to the appropriate cluster ports hosted on cluster switch C2:

```
cluster::*> network interface show -role Cluster
(network interface show)

      Logical      Status      Network      Current      Current      Is
Vserver Interface  Admin/Oper  Address/Mask  Node         Port         Is
Home
-----
-----
Cluster
false  n1_clus1  up/up      10.10.0.1/24  n1           e0b
true   n1_clus2  up/up      10.10.0.2/24  n1           e0b
true   n1_clus3  up/up      10.10.0.3/24  n1           e0c
false  n1_clus4  up/up      10.10.0.4/24  n1           e0c
false  n2_clus1  up/up      10.10.0.5/24  n2           e0b
true   n2_clus2  up/up      10.10.0.6/24  n2           e0b
true   n2_clus3  up/up      10.10.0.7/24  n2           e0c
false  n2_clus4  up/up      10.10.0.8/24  n2           e0c

8 entries were displayed.
```

14. Shut down the node ports that are connected to CL1 on all of the nodes:

```
network port modify
```

Show example

The following example shows how to shut down the specified ports on nodes n1 and n2:

```
cluster::*> network port modify -node n1 -port e0a -up-admin false
cluster::*> network port modify -node n1 -port e0d -up-admin false
cluster::*> network port modify -node n2 -port e0a -up-admin false
cluster::*> network port modify -node n2 -port e0d -up-admin false
```

15. Shut down the ISL ports 24, 31, and 32 on the active 3132Q-V switch C2:

```
shutdown
```

Show example

The following example shows how to shut down ISLs 24, 31, and 32 on the active 3132Q-V switch C2:

```
C2# configure
C2(config)# interface ethernet 1/24/1-4
C2(config-if-range)# shutdown
C2(config-if-range)# exit
C2(config)# interface ethernet 1/31-32
C2(config-if-range)# shutdown
C2(config-if-range)# exit
C2(config)# exit
C2#
```

16. Remove the cables that are attached to the CN1610 switch CL1 on all of the nodes.

With supported cabling, you must reconnect the disconnected ports on all of the nodes to the Nexus 3132Q-V switch C1.

17. Remove the QSFP cables from Nexus 3132Q-V C2 port e1/24.

You must connect ports e1/31 and e1/32 on C1 to ports e1/31 and e1/32 on C2 using supported Cisco QSFP optical fiber or direct-attach cables.

18. Restore the configuration on port 24 and remove the temporary port-channel 2 on C2, by copying the running-configuration file to the startup-configuration file.

Show example

The following example copies the running-configuration file to the startup-configuration file:

```
C2# configure
C2(config)# no interface breakout module 1 port 24 map 10g-4x
C2(config)# no interface port-channel 2
C2(config-if)# interface e1/24
C2(config-if)# description 40GbE Node Port
C2(config-if)# spanning-tree port type edge
C2(config-if)# spanning-tree bpduguard enable
C2(config-if)# mtu 9216
C2(config-if-range)# exit
C2(config)# exit
C2# copy running-config startup-config
[#####] 100%
Copy Complete.
```

19. Bring up ISL ports 31 and 32 on C2, the active 3132Q-V switch:

no shutdown

Show example

The following example shows how to bring up ISLs 31 and 32 on the 3132Q-V switch C2:

```
C2# configure
C2(config)# interface ethernet 1/31-32
C2(config-if-range)# no shutdown
C2(config-if-range)# exit
C2(config)# exit
C2# copy running-config startup-config
[#####] 100%
Copy Complete.
```

What's next?

[Complete your migration.](#)

Complete your migration from CN1610 switches to Nexus 3132Q-V switches

Complete the following steps to finalize the CN1610 switches migration to Nexus 3132Q-V switches.

Steps

1. Verify that the ISL connections are up on the 3132Q-V switch C2:

```
show port-channel summary
```

Ports Eth1/31 and Eth1/32 should indicate (P), meaning that both the ISL ports are up in the port-channel.

Show example

```
C1# show port-channel summary
Flags:  D - Down          P - Up in port-channel (members)
        I - Individual   H - Hot-standby (LACP only)
        s - Suspended    r - Module-removed
        S - Switched     R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met

-----
-----
Group Port-          Type      Protocol  Member Ports
Channel
-----
-----
1      Po1(SU)       Eth       LACP      Eth1/31(P)  Eth1/32(P)
```

2. Bring up all of the cluster interconnect ports connected to the new 3132Q-V switch C1 on all of the nodes:

```
network port modify
```

Show example

The following example shows how to bring up all of the cluster interconnect ports connected to the new 3132Q-V switch C1:

```
cluster::*> network port modify -node n1 -port e0a -up-admin true
cluster::*> network port modify -node n1 -port e0d -up-admin true
cluster::*> network port modify -node n2 -port e0a -up-admin true
cluster::*> network port modify -node n2 -port e0d -up-admin true
```

3. Verify the status of the cluster node port:

```
network port show
```

Show example

The following example verifies that all of the cluster interconnect ports on n1 and n2 on the new 3132Q-V switch C1 are up:

```
cluster::*> network port show -role Cluster
(network port show)

Node: n1

```

Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Open	Health Status	Ignore Health
e0a	cluster	cluster	up	9000	auto/10000	-	-
e0b	cluster	cluster	up	9000	auto/10000	-	-
e0c	cluster	cluster	up	9000	auto/10000	-	-
e0d	cluster	cluster	up	9000	auto/10000	-	-

```
Node: n2

```

Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Open	Health Status	Ignore Health
e0a	cluster	cluster	up	9000	auto/10000	-	-
e0b	cluster	cluster	up	9000	auto/10000	-	-
e0c	cluster	cluster	up	9000	auto/10000	-	-
e0d	cluster	cluster	up	9000	auto/10000	-	-

```
8 entries were displayed.
```

4. Revert all of the migrated cluster interconnect LIFs that were originally connected to C1 on all of the nodes:

```
network interface revert
```


Show example

The following example shows how to revert the migrated cluster LIFs to their home ports:

```
cluster::*> network interface revert -vserver Cluster -lif n1_clus1
cluster::*> network interface revert -vserver Cluster -lif n1_clus4
cluster::*> network interface revert -vserver Cluster -lif n2_clus1
cluster::*> network interface revert -vserver Cluster -lif n2_clus4
```

5. Verify that the interface is now home:

```
network interface show
```

Show example

The following example shows the status of cluster interconnect interfaces is up and Is home for n1 and n2:

```
cluster::*> network interface show -role Cluster
(network interface show)
```

Vserver Home	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is
-----	-----	-----	-----	-----	-----	-----
Cluster						
true	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	

8 entries were displayed.

6. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Node	Date		Source LIF	Destination LIF	Packet Loss
n1					
	3/5/2022 19:21:18 -06:00		n1_clus2	n1_clus1	none
	3/5/2022 19:21:20 -06:00		n1_clus2	n2_clus2	none
n2					
	3/5/2022 19:21:18 -06:00		n2_clus2	n1_clus1	none
	3/5/2022 19:21:20 -06:00		n2_clus2	n1_clus2	none

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1      e0a    10.10.0.1
Cluster n1_clus2 n1      e0b    10.10.0.2
Cluster n1_clus3 n1      e0c    10.10.0.3
Cluster n1_clus4 n1      e0d    10.10.0.4
Cluster n2_clus1 n2      e0a    10.10.0.5
Cluster n2_clus2 n2      e0b    10.10.0.6
Cluster n2_clus3 n2      e0c    10.10.0.7
Cluster n2_clus4 n2      e0d    10.10.0.8

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
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 16 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 10.10.0.2 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 10.10.0.3 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 10.10.0.4 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

Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
4 paths up, 0 paths down (udp check)

```

7. Expand the cluster by adding nodes to the Nexus 3132Q-V cluster switches.

8. Display the information about the devices in your configuration:

- `network device-discovery show`
- `network port show -role cluster`
- `network interface show -role cluster`
- `system cluster-switch show`

Show example

The following examples show nodes n3 and n4 with 40 GbE cluster ports connected to ports e1/7 and e1/8, respectively on both the Nexus 3132Q-V cluster switches, and both nodes have joined the cluster. The 40 GbE cluster interconnect ports used are e4a and e4e.

```
cluster::*> network device-discovery show
```

Node	Local Port	Discovered Device	Interface	Platform
n1	/cdp			
	e0a	C1	Ethernet1/1/1	N3K-C3132Q-V
	e0b	C2	Ethernet1/1/1	N3K-C3132Q-V
	e0c	C2	Ethernet1/1/2	N3K-C3132Q-V
	e0d	C1	Ethernet1/1/2	N3K-C3132Q-V
n2	/cdp			
	e0a	C1	Ethernet1/1/3	N3K-C3132Q-V
	e0b	C2	Ethernet1/1/3	N3K-C3132Q-V
	e0c	C2	Ethernet1/1/4	N3K-C3132Q-V
	e0d	C1	Ethernet1/1/4	N3K-C3132Q-V
n3	/cdp			
	e4a	C1	Ethernet1/7	N3K-C3132Q-V
	e4e	C2	Ethernet1/7	N3K-C3132Q-V
n4	/cdp			
	e4a	C1	Ethernet1/8	N3K-C3132Q-V
	e4e	C2	Ethernet1/8	N3K-C3132Q-V

```
12 entries were displayed.
```

```
cluster::*> network port show -role cluster
(network port show)
```

```
Node: n1
```

Ignore	Broadcast	Speed (Mbps)	Health				
Port	IPspace	Domain	Link	MTU	Admin/Open	Status	
Health	Status						
e0a	cluster	cluster	up	9000	auto/10000	-	-
e0b	cluster	cluster	up	9000	auto/10000	-	-
e0c	cluster	cluster	up	9000	auto/10000	-	-
e0d	cluster	cluster	up	9000	auto/10000	-	-

Node: n2

		Broadcast			Speed (Mbps)	Health	
Ignore							
Port	IPspace	Domain	Link	MTU	Admin/Open	Status	
Health	Status						
-----	-----	-----	-----	-----	-----	-----	-----
-----	-----	-----	-----	-----	-----	-----	-----
e0a	cluster	cluster	up	9000	auto/10000	-	-
e0b	cluster	cluster	up	9000	auto/10000	-	-
e0c	cluster	cluster	up	9000	auto/10000	-	-
e0d	cluster	cluster	up	9000	auto/10000	-	-

Node: n3

		Broadcast			Speed (Mbps)	Health	
Ignore							
Port	IPspace	Domain	Link	MTU	Admin/Open	Status	
Health	Status						
-----	-----	-----	-----	-----	-----	-----	-----
-----	-----	-----	-----	-----	-----	-----	-----
e4a	cluster	cluster	up	9000	auto/40000	-	-
e4e	cluster	cluster	up	9000	auto/40000	-	-

Node: n4

		Broadcast			Speed (Mbps)	Health	
Ignore							
Port	IPspace	Domain	Link	MTU	Admin/Open	Status	
Health	Status						
-----	-----	-----	-----	-----	-----	-----	-----
-----	-----	-----	-----	-----	-----	-----	-----
e4a	cluster	cluster	up	9000	auto/40000	-	-
e4e	cluster	cluster	up	9000	auto/40000	-	-

12 entries were displayed.

```
cluster::*> network interface show -role Cluster
(network interface show)
```

Is	Logical	Status	Network	Current	Current
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.
```



```
cluster::> system cluster-switch show
```

```
Switch                               Type                Address             Model
-----
C1                                   cluster-network    10.10.1.103
NX3132V
  Serial Number: FOX000001
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS)
  Software, Version
                        7.0(3)I4(1)
  Version Source: CDP

C2                                   cluster-network    10.10.1.104
NX3132V
  Serial Number: FOX000002
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS)
  Software, Version
                        7.0(3)I4(1)
  Version Source: CDP

CL1                                   cluster-network    10.10.1.101        CN1610
  Serial Number: 01234567
  Is Monitored: true
  Reason:
  Software Version: 1.2.0.7
  Version Source: ISDP

CL2                                   cluster-network    10.10.1.102
CN1610
  Serial Number: 01234568
  Is Monitored: true
  Reason:
  Software Version: 1.2.0.7
  Version Source: ISDP

4 entries were displayed.
```

9. Remove the replaced CN1610 switches if they are not automatically removed:

```
system cluster-switch delete
```

Show example

The following example shows how to remove the CN1610 switches:

```
cluster::> system cluster-switch delete -device CL1
cluster::> system cluster-switch delete -device CL2
```

10. Configure clusters clus1 and clus4 to `-auto-revert` on each node and confirm:

Show example

```
cluster::*> network interface modify -vserver nodel -lif clus1 -auto
-revert true
cluster::*> network interface modify -vserver nodel -lif clus4 -auto
-revert true
cluster::*> network interface modify -vserver node2 -lif clus1 -auto
-revert true
cluster::*> network interface modify -vserver node2 -lif clus4 -auto
-revert true
```

11. Verify that the proper cluster switches are monitored:

```
system cluster-switch show
```

Show example

```
cluster::> system cluster-switch show

Switch                               Type                               Address
Model
-----
C1                                   cluster-network                    10.10.1.103
NX3132V
  Serial Number: FOX000001
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        7.0(3)I4(1)
  Version Source: CDP

C2                                   cluster-network                    10.10.1.104
NX3132V
  Serial Number: FOX000002
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        7.0(3)I4(1)
  Version Source: CDP

2 entries were displayed.
```

12. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next?

[Configure switch health monitoring.](#)

Cisco Nexus 5596 switch migration

Migrate from Nexus 5596 switches to Nexus 3132Q-V switches workflow

Follow these workflow steps to migrate your Cisco Nexus 5596 switches to Cisco Nexus 3132Q-V switches.

1

Migration requirements

Review the requirements and example switch information for the migration process.

2

Prepare for migration

Prepare your Nexus 5596 switches for migration to Nexus 3132Q-V switches.

3

Configure your ports

Configure your ports for migration to the new Nexus 3132Q-V switches.

4

Complete your migration

Complete your migration to the new Nexus 3132Q-V switches.

Migration requirements

Cisco Nexus 3132Q-V switches can be used as cluster switches in your AFF or FAS cluster. Cluster switches allow you to build ONTAP clusters with more than two nodes.



The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

For more information, see:

- [Cisco Ethernet Switch description page](#)
- [Hardware Universe](#)

Cisco Nexus 5596 requirements

The cluster switches use the following ports for connections to nodes:

- Nexus 5596: ports e1/1-40 (10 GbE)
- Nexus 3132Q-V: ports e1/1-30 (10/40/100 GbE)

The cluster switches use the following Inter-Switch Link (ISL) ports:

- Nexus 5596: ports e1/41-48 (10 GbE)
- Nexus 3132Q-V: ports e1/31-32 (40/100 GbE)

The [Hardware Universe](#) contains information about supported cabling to Nexus 3132Q-V switches:

- Nodes with 10 GbE cluster connections require QSFP to SFP+ optical fiber breakout cables or QSFP to SFP+ copper breakout cables.
- Nodes with 40/100 GbE cluster connections require supported QSFP/QSFP28 optical modules with fiber cables or QSFP/QSFP28 copper direct-attach cables.

The cluster switches use the appropriate ISL cabling:

- Beginning: Nexus 5596 (SFP+ to SFP+)
 - 8x SFP+ fiber or copper direct-attach cables
- Interim: Nexus 5596 to Nexus 3132Q-V (QSFP to 4xSFP+ break-out)
 - 1x QSFP to SFP+ fiber break-out or copper break-out cables
- Final: Nexus 3132Q-V to Nexus 3132Q-V (QSFP28 to QSFP28)
 - 2x QSFP28 fiber or copper direct-attach cables
- On Nexus 3132Q-V switches, you can operate QSFP/QSFP28 ports in either 40/100 Gigabit Ethernet or 4 x10 Gigabit Ethernet modes.

By default, there are 32 ports in the 40/100 Gigabit Ethernet mode. These 40 Gigabit Ethernet ports are numbered in a 2-tuple naming convention. For example, the second 40 Gigabit Ethernet port is numbered as 1/2.

The process of changing the configuration from 40 Gigabit Ethernet to 10 Gigabit Ethernet is called *breakout* and the process of changing the configuration from 10 Gigabit Ethernet to 40 Gigabit Ethernet is called *breakin*.

When you break out a 40/100 Gigabit Ethernet port into 10 Gigabit Ethernet ports, the resulting ports are numbered using a 3-tuple naming convention. For example, the break-out ports of the second 40/100 Gigabit Ethernet port are numbered as 1/2/1, 1/2/2, 1/2/3, and 1/2/4.

- On the left side of Nexus 3132Q-V switches are 2 SFP+ ports, called 1/33 and 1/34.
- You have configured some of the ports on Nexus 3132Q-V switches to run at 10 GbE or 40/100 GbE.



You can break out the first six ports into 4x10 GbE mode by using the `interface breakout module 1 port 1-6 map 10g-4x` command. Similarly, you can regroup the first six QSFP+ ports from breakout configuration by using the `no interface breakout module 1 port 1-6 map 10g-4x` command.

- You have done the planning, migration, and read the required documentation on 10 GbE and 40/100 GbE connectivity from nodes to Nexus 3132Q-V cluster switches.
- The ONTAP and NX-OS versions supported in this procedure are on the [Cisco Ethernet Switches page](#).

About the examples used

The examples in this procedure describe replacing Cisco Nexus 5596 switches with Cisco Nexus 3132Q-V switches. You can use these steps (with modifications) for other older Cisco switches.

The procedure also uses the following switch and node nomenclature:

- The command outputs might vary depending on different releases of ONTAP.
- The Nexus 5596 switches to be replaced are **CL1** and **CL2**.
- The Nexus 3132Q-V switches to replace the Nexus 5596 switches are **C1** and **C2**.
- **n1_clus1** is the first cluster logical interface (LIF) connected to cluster switch 1 (CL1 or C1) for node **n1**.
- **n1_clus2** is the first cluster LIF connected to cluster switch 2 (CL2 or C2) for node **n1**.
- **n1_clus3** is the second LIF connected to cluster switch 2 (CL2 or C2) for node **n1**.
- **n1_clus4** is the second LIF connected to cluster switch 1 (CL1 or C1) for node **n1**.
- The number of 10 GbE and 40/100 GbE ports are defined in the reference configuration files (RCFs)

available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.

- The nodes are **n1**, **n2**, **n3**, and **n4**.

The examples in this procedure use four nodes:

- Two nodes use four 10 GbE cluster interconnect ports: **e0a**, **e0b**, **e0c**, and **e0d**.
- The other two nodes use two 40 GbE cluster interconnect ports: **e4a** and **e4e**.

The [Hardware Universe](#) lists the actual cluster ports on your platforms.

Scenarios covered

This procedure covers the following scenarios:

- The cluster starts with two nodes connected and functioning in a two Nexus 5596 cluster switches.
- The cluster switch CL2 to be replaced by C2 (steps 1 to 19):
 - Traffic on all cluster ports and LIFs on all nodes connected to CL2 are migrated onto the first cluster ports and LIFs connected to CL1.
 - Disconnect cabling from all cluster ports on all nodes connected to CL2, and then use supported break-out cabling to reconnect the ports to new cluster switch C2.
 - Disconnect cabling between ISL ports between CL1 and CL2, and then use supported break-out cabling to reconnect the ports from CL1 to C2.
 - Traffic on all cluster ports and LIFs connected to C2 on all nodes is reverted.
- The cluster switch CL2 to be replaced by C2.
 - Traffic on all cluster ports or LIFs on all nodes connected to CL1 are migrated onto the second cluster ports or LIFs connected to C2.
 - Disconnect cabling from all cluster port on all nodes connected to CL1 and reconnect, using supported break-out cabling, to new cluster switch C1.
 - Disconnect cabling between ISL ports between CL1 and C2, and reconnect using supported cabling, from C1 to C2.
 - Traffic on all cluster ports or LIFs connected to C1 on all nodes is reverted.
- Two FAS9000 nodes have been added to cluster with examples showing cluster details.

What's next?

[Prepare for migration.](#)

Prepare for migration from Nexus 5596 switches to Nexus 3132Q-V switches

Follow these steps to prepare your Cisco Nexus 5596 switches for migration to Cisco Nexus 3132Q-V switches.

Steps

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=xh
```


x is the duration of the maintenance window in hours.



The message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Display information about the devices in your configuration:

```
network device-discovery show
```

Show example

The following example shows how many cluster interconnect interfaces have been configured in each node for each cluster interconnect switch:

```
cluster::> network device-discovery show
      Local  Discovered
Node   Port    Device           Interface      Platform
-----
n1     /cdp
      e0a    CL1              Ethernet1/1    N5K-C5596UP
      e0b    CL2              Ethernet1/1    N5K-C5596UP
      e0c    CL2              Ethernet1/2    N5K-C5596UP
      e0d    CL1              Ethernet1/2    N5K-C5596UP
n2     /cdp
      e0a    CL1              Ethernet1/3    N5K-C5596UP
      e0b    CL2              Ethernet1/3    N5K-C5596UP
      e0c    CL2              Ethernet1/4    N5K-C5596UP
      e0d    CL1              Ethernet1/4    N5K-C5596UP
8 entries were displayed.
```

3. Determine the administrative or operational status for each cluster interface:

a. Display the network port attributes:

```
network port show
```

Show example

The following example displays the network port attributes on a system:

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health      Health      Speed (Mbps)
Port        IPspace     Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
-----
e0a         Cluster     Cluster     up    9000 auto/10000 -
-
e0b         Cluster     Cluster     up    9000 auto/10000 -
-
e0c         Cluster     Cluster     up    9000 auto/10000 -
-
e0d         Cluster     Cluster     up    9000 auto/10000 -
-

Node: n2

Ignore

Health      Health      Speed (Mbps)
Port        IPspace     Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
-----
e0a         Cluster     Cluster     up    9000 auto/10000 -
-
e0b         Cluster     Cluster     up    9000 auto/10000 -
-
e0c         Cluster     Cluster     up    9000 auto/10000 -
-
e0d         Cluster     Cluster     up    9000 auto/10000 -
-

8 entries were displayed.
```

b. Display information about the logical interfaces:


```
network interface show
```

Show example

The following example displays the general information about all of the LIFs on your system:

```
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
e0a      n1_clus1   up/up      10.10.0.1/24  n1
true
e0b      n1_clus2   up/up      10.10.0.2/24  n1
true
e0c      n1_clus3   up/up      10.10.0.3/24  n1
true
e0d      n1_clus4   up/up      10.10.0.4/24  n1
true
e0a      n2_clus1   up/up      10.10.0.5/24  n2
true
e0b      n2_clus2   up/up      10.10.0.6/24  n2
true
e0c      n2_clus3   up/up      10.10.0.7/24  n2
true
e0d      n2_clus4   up/up      10.10.0.8/24  n2
true
8 entries were displayed.
```

c. Display information about the discovered cluster switches:

```
system cluster-switch show
```

Show example

The following example displays the cluster switches that are known to the cluster, along with their management IP addresses:

```
cluster::*> system cluster-switch show

Switch                               Type                               Address
Model
-----
CL1                                  cluster-network                    10.10.1.101
NX5596
    Serial Number: 01234567
    Is Monitored: true
    Reason:
    Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                               7.1(1)N1(1)
    Version Source: CDP
CL2                                  cluster-network                    10.10.1.102
NX5596
    Serial Number: 01234568
    Is Monitored: true
    Reason:
    Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                               7.1(1)N1(1)
    Version Source: CDP

2 entries were displayed.
```

4. Set the `-auto-revert` parameter to `false` on cluster LIFs `clus1` and `clus2` on both nodes:

```
network interface modify
```

Show example

```
cluster::*> network interface modify -vserver node1 -lif clus1 -auto
-revert false
cluster::*> network interface modify -vserver node1 -lif clus2 -auto
-revert false
cluster::*> network interface modify -vserver node2 -lif clus1 -auto
-revert false
cluster::*> network interface modify -vserver node2 -lif clus2 -auto
-revert false
```

5. Verify that the appropriate RCF and image are installed on the new 3132Q-V switches as necessary for your requirements, and make the essential site customizations, such as users and passwords, network addresses, and so on.

You must prepare both switches at this time. If you need to upgrade the RCF and image, follow these steps:

- a. Go to the [Cisco Ethernet Switches](#) page on the NetApp Support Site.
- b. Note your switch and the required software versions in the table on that page.
- c. Download the appropriate version of the RCF.
- d. Select **CONTINUE** on the **Description** page, accept the license agreement, and then follow the instructions on the **Download** page to download the RCF.
- e. Download the appropriate version of the image software.

See the *ONTAP 8.x or later Cluster and Management Network Switch Reference Configuration FilesDownload* page, and then select the appropriate version.

To find the correct version, see the *ONTAP 8.x or later Cluster Network Switch Download page*.

6. Migrate the LIFs associated with the second Nexus 5596 switch to be replaced:

```
network interface migrate
```

Show example

The following example shows n1 and n2, but LIF migration must be done on all of the nodes:

```
cluster::*> network interface migrate -vserver Cluster -lif n1_clus2
-source-node n1 -
destination-node n1 -destination-port e0a
cluster::*> network interface migrate -vserver Cluster -lif n1_clus3
-source-node n1 -
destination-node n1 -destination-port e0d
cluster::*> network interface migrate -vserver Cluster -lif n2_clus2
-source-node n2 -
destination-node n2 -destination-port e0a
cluster::*> network interface migrate -vserver Cluster -lif n2_clus3
-source-node n2 -
destination-node n2 -destination-port e0d
```

7. Verify the cluster's health:

```
network interface show
```

Show example

The following example shows the result of the previous `network interface migrate` command:

```
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
e0a      n1_clus1    up/up      10.10.0.1/24  n1
true
e0a      n1_clus2    up/up      10.10.0.2/24  n1
false
e0d      n1_clus3    up/up      10.10.0.3/24  n1
false
e0d      n1_clus4    up/up      10.10.0.4/24  n1
true
e0a      n2_clus1    up/up      10.10.0.5/24  n2
true
e0a      n2_clus2    up/up      10.10.0.6/24  n2
false
e0d      n2_clus3    up/up      10.10.0.7/24  n2
false
e0d      n2_clus4    up/up      10.10.0.8/24  n2
true
8 entries were displayed.
```

8. Shut down the cluster interconnect ports that are physically connected to switch CL2:

```
network port modify
```

Show example

The following commands shut down the specified ports on n1 and n2, but the ports must be 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
```

9. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Node	Date	Loss	Source LIF	Destination LIF	
	n1	3/5/2022 19:21:18	-06:00	n1_clus2	n2_clus1	none
		3/5/2022 19:21:20	-06:00	n1_clus2	n2_clus2	none
	n2	3/5/2022 19:21:18	-06:00	n2_clus2	n1_clus1	none
		3/5/2022 19:21:20	-06:00	n2_clus2	n1_clus2	none

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1      e0a 10.10.0.1
Cluster n1_clus2 n1      e0b 10.10.0.2
Cluster n1_clus3 n1      e0c 10.10.0.3
Cluster n1_clus4 n1      e0d 10.10.0.4
Cluster n2_clus1 n2      e0a 10.10.0.5
Cluster n2_clus2 n2      e0b 10.10.0.6
Cluster n2_clus3 n2      e0c 10.10.0.7
Cluster n2_clus4 n2      e0d 10.10.0.8

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
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 16 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 10.10.0.2 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 10.10.0.3 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 10.10.0.4 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
Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
4 paths up, 0 paths down (udp check)

```

10. Shut down the ISL ports 41 through 48 on the active Nexus 5596 switch CL1:

Show example

The following example shows how to shut down ISL ports 41 through 48 on the Nexus 5596 switch CL1:

```
(CL1)# configure
(CL1)(Config)# interface e1/41-48
(CL1)(config-if-range)# shutdown
(CL1)(config-if-range)# exit
(CL1)(Config)# exit
(CL1)#
```

If you are replacing a Nexus 5010 or 5020, specify the appropriate port numbers for ISL.

11. Build a temporary ISL between CL1 and C2.

Show example

The following example shows a temporary ISL being set up between CL1 and C2:

```
C2# configure
C2(config)# interface port-channel 2
C2(config-if)# switchport mode trunk
C2(config-if)# spanning-tree port type network
C2(config-if)# mtu 9216
C2(config-if)# interface breakout module 1 port 24 map 10g-4x
C2(config)# interface e1/24/1-4
C2(config-if-range)# switchport mode trunk
C2(config-if-range)# mtu 9216
C2(config-if-range)# channel-group 2 mode active
C2(config-if-range)# exit
C2(config-if)# exit
```

What's next?

[Configure your ports.](#)

Configure your ports for migration from 5596 switches to 3132Q-V switches

Follow these steps to configure your ports for migration from the Nexus 5596 switches to the new Nexus 3132Q-V switches.

Steps

1. On all nodes, remove all cables attached to the Nexus 5596 switch CL2.

With supported cabling, reconnect disconnected ports on all nodes to the Nexus 3132Q-V switch C2.

2. Remove all the cables from the Nexus 5596 switch CL2.

Attach the appropriate Cisco QSFP to SFP+ break-out cables connecting port 1/24 on the new Cisco 3132Q-V switch, C2, to ports 45 to 48 on existing Nexus 5596, CL1.

3. Verify that interfaces eth1/45-48 already have channel-group 1 mode active in their running configuration.
4. Bring up ISLs ports 45 through 48 on the active Nexus 5596 switch CL1.

Show example

The following example shows ISLs ports 45 through 48 being brought up:

```
(CL1)# configure
(CL1) (Config)# interface e1/45-48
(CL1) (config-if-range)# no shutdown
(CL1) (config-if-range)# exit
(CL1) (Config)# exit
(CL1)#
```

5. Verify that the ISLs are up on the Nexus 5596 switch CL1:

```
show port-channel summary
```

Show example

Ports eth1/45 through eth1/48 should indicate (P) meaning that the ISL ports are up in the port-channel:

Example

```
CL1# show port-channel summary
```

```
Flags: D - Down          P - Up in port-channel (members)
       I - Individual    H - Hot-standby (LACP only)
       s - Suspended     r - Module-removed
       S - Switched      R - Routed
       U - Up (port-channel)
       M - Not in use. Min-links not met
```

```
-----
-----
Group Port-          Type   Protocol  Member Ports
Channel
-----
-----
1      Po1(SU)       Eth    LACP     Eth1/41(D)  Eth1/42(D)
Eth1/43(D)
                                     Eth1/44(D)  Eth1/45(P)
Eth1/46(P)
                                     Eth1/47(P)  Eth1/48(P)
```

6. Verify that the ISLs are up on the 3132Q-V switch C2:

```
show port-channel summary
```

Show example

Ports eth1/24/1, eth1/24/2, eth1/24/3, and eth1/24/4 should indicate (P) meaning that the ISL ports are up in the port-channel:

```
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 (D)  Eth1/32 (D)
2      Po2 (SU)       Eth     LACP      Eth1/24/1 (P) Eth1/24/2 (P)
Eth1/24/3 (P)
                                   Eth1/24/4 (P)
```

7. On all nodes, bring up all the cluster interconnect ports connected to the 3132Q-V switch C2:

```
network port modify
```

Show example

The following example shows the specified ports being brought up on nodes n1 and n2:

```
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
```

8. On all nodes, revert all of the migrated cluster interconnect LIFs connected to C2:

```
network interface revert
```

Show example

The following example shows the migrated cluster LIFs being reverted to their home ports on nodes n1 and n2:

```
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
```

9. Verify all the cluster interconnect ports are now reverted to their home:

```
network interface show
```

Show example

The following example shows that the LIFs on clus2 reverted to their home ports and shows that the LIFs are successfully reverted if the ports in the Current Port column have a status of `true` in the `Is Home` column. If the `Is Home` value is `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
e0a      n1_clus1    up/up      10.10.0.1/24  n1
      true
e0b      n1_clus2    up/up      10.10.0.2/24  n1
      true
e0c      n1_clus3    up/up      10.10.0.3/24  n1
      true
e0d      n1_clus4    up/up      10.10.0.4/24  n1
      true
e0a      n2_clus1    up/up      10.10.0.5/24  n2
      true
e0b      n2_clus2    up/up      10.10.0.6/24  n2
      true
e0c      n2_clus3    up/up      10.10.0.7/24  n2
      true
e0d      n2_clus4    up/up      10.10.0.8/24  n2
      true
8 entries were displayed.
```

10. Verify that the clustered ports are connected:

```
network port show
```

Show example

The following example shows the result of the previous `network port modify` command, verifying that all the cluster interconnects are up:

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Speed(Mbps) Health
Status
-----
-----
e0a      Cluster      Cluster      up   9000  auto/10000  -
-
e0b      Cluster      Cluster      up   9000  auto/10000  -
-
e0c      Cluster      Cluster      up   9000  auto/10000  -
-
e0d      Cluster      Cluster      up   9000  auto/10000  -
-

Node: n2

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Speed(Mbps) Health
Status
-----
-----
e0a      Cluster      Cluster      up   9000  auto/10000  -
-
e0b      Cluster      Cluster      up   9000  auto/10000  -
-
e0c      Cluster      Cluster      up   9000  auto/10000  -
-
e0d      Cluster      Cluster      up   9000  auto/10000  -
-

8 entries were displayed.
```

11. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		

n1		
3/5/2022 19:21:18 -06:00	n1_clus2	n2_clus1
3/5/2022 19:21:20 -06:00	n1_clus2	n2_clus2
		none
		none
n2		
3/5/2022 19:21:18 -06:00	n2_clus2	n1_clus1
3/5/2022 19:21:20 -06:00	n2_clus2	n1_clus2
		none
		none

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1      e0a 10.10.0.1
Cluster n1_clus2 n1      e0b 10.10.0.2
Cluster n1_clus3 n1      e0c 10.10.0.3
Cluster n1_clus4 n1      e0d 10.10.0.4
Cluster n2_clus1 n2      e0a 10.10.0.5
Cluster n2_clus2 n2      e0b 10.10.0.6
Cluster n2_clus3 n2      e0c 10.10.0.7
Cluster n2_clus4 n2      e0d 10.10.0.8

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
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 16 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 10.10.0.2 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 10.10.0.3 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 10.10.0.4 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
Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
4 paths up, 0 paths down (udp check)

```

12. On each node in the cluster, migrate the interfaces associated with the first Nexus 5596 switch, CL1, to be replaced:

```
network interface migrate
```

Show example

The following example shows the ports or LIFs being migrated on nodes n1 and n2:

```
cluster::*> network interface migrate -vserver Cluster -lif n1_clus1
-source-node n1 -
destination-node n1 -destination-port e0b
cluster::*> network interface migrate -vserver Cluster -lif n1_clus4
-source-node n1 -
destination-node n1 -destination-port e0c
cluster::*> network interface migrate -vserver Cluster -lif n2_clus1
-source-node n2 -
destination-node n2 -destination-port e0b
cluster::*> network interface migrate -vserver Cluster -lif n2_clus4
-source-node n2 -
destination-node n2 -destination-port e0c
```

13. Verify the cluster status:

```
network interface show
```

Show example

The following example shows that the required cluster LIFs have been migrated to appropriate cluster ports hosted on cluster switch C2:

```
(network interface show)
      Logical      Status      Network      Current
Current Is
Vserver      Interface  Admin/Oper  Address/Mask  Node
Port      Home
-----
Cluster
e0b          n1_clus1   up/up       10.10.0.1/24  n1
false
e0b          n1_clus2   up/up       10.10.0.2/24  n1
true
e0c          n1_clus3   up/up       10.10.0.3/24  n1
true
e0c          n1_clus4   up/up       10.10.0.4/24  n1
false
e0b          n2_clus1   up/up       10.10.0.5/24  n2
false
e0b          n2_clus2   up/up       10.10.0.6/24  n2
true
e0c          n2_clus3   up/up       10.10.0.7/24  n2
true
e0c          n2_clus4   up/up       10.10.0.8/24  n2
false
8 entries were displayed.
-----
```

14. On all the nodes, shut down the node ports that are connected to CL1:

```
network port modify
```

Show example

The following example shows the specified ports being shut down on nodes n1 and n2:

```
cluster::*> network port modify -node n1 -port e0a -up-admin false
cluster::*> network port modify -node n1 -port e0d -up-admin false
cluster::*> network port modify -node n2 -port e0a -up-admin false
cluster::*> network port modify -node n2 -port e0d -up-admin false
```

15. Shut down the ISL ports 24, 31, and 32 on the active 3132Q-V switch C2:

```
shutdown
```

Show example

The following example shows how to shut down ISLs 24, 31, and 32:

```
C2# configure
C2(Config)# interface e1/24/1-4
C2(config-if-range)# shutdown
C2(config-if-range)# exit
C2(config)# interface 1/31-32
C2(config-if-range)# shutdown
C2(config-if-range)# exit
C2(config-if)# exit
C2#
```

16. On all nodes, remove all cables attached to the Nexus 5596 switch CL1.

With supported cabling, reconnect disconnected ports on all nodes to the Nexus 3132Q-V switch C1.

17. Remove the QSFP breakout cable from Nexus 3132Q-V C2 ports e1/24.

Connect ports e1/31 and e1/32 on C1 to ports e1/31 and e1/32 on C2 using supported Cisco QSFP optical fiber or direct-attach cables.

18. Restore the configuration on port 24 and remove the temporary Port Channel 2 on C2:

```

C2# configure
C2(config)# no interface breakout module 1 port 24 map 10g-4x
C2(config)# no interface port-channel 2
C2(config-if)# int e1/24
C2(config-if)# description 40GbE Node Port
C2(config-if)# spanning-tree port type edge
C2(config-if)# spanning-tree bpduguard enable
C2(config-if)# mtu 9216
C2(config-if-range)# exit
C2(config)# exit
C2# copy running-config startup-config
[#####] 100%
Copy Complete.

```

19. Bring up ISL ports 31 and 32 on C2, the active 3132Q-V switch: no shutdown

Show example

The following example shows how to bring up ISLs 31 and 32 on the 3132Q-V switch C2:

```

C2# configure
C2(config)# interface ethernet 1/31-32
C2(config-if-range)# no shutdown
C2(config-if-range)# exit
C2(config)# exit
C2# copy running-config startup-config
[#####] 100%
Copy Complete.

```

What's next?

[Complete your migration.](#)

Complete the migration from Nexus 5596 switches to Nexus 3132Q-V switches

Complete the following steps to finalize the Nexus 5596 switches migration to Nexus 3132Q-V switches.

Steps

1. Verify that the ISL connections are up on the 3132Q-V switch C2:

```
show port-channel summary
```

Show example

Ports Eth1/31 and Eth1/32 should indicate (P), meaning that both the ISL ports are up in the port-channel:

```
C1# show port-channel summary
Flags: D - Down          P - Up in port-channel (members)
       I - Individual    H - Hot-standby (LACP only)
       s - Suspended     r - Module-removed
       S - Switched      R - Routed
       U - Up (port-channel)
       M - Not in use. Min-links not met

-----
-----
Group Port-          Type   Protocol  Member Ports
Channel
-----
-----
1      Po1 (SU)      Eth     LACP      Eth1/31 (P)  Eth1/32 (P)
```

2. On all nodes, bring up all the cluster interconnect ports connected to the new 3132Q-V switch C1:

```
network port modify
```

Show example

The following example shows all the cluster interconnect ports being brought up for n1 and n2 on the 3132Q-V switch C1:

```
cluster::*> network port modify -node n1 -port e0a -up-admin true
cluster::*> network port modify -node n1 -port e0d -up-admin true
cluster::*> network port modify -node n2 -port e0a -up-admin true
cluster::*> network port modify -node n2 -port e0d -up-admin true
```

3. Verify the status of the cluster node port:

```
network port show
```

Show example

The following example verifies that all cluster interconnect ports on all nodes on the new 3132Q-V switch C1 are up:

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e0a      Cluster      Cluster      up   9000  auto/10000  -
-
e0b      Cluster      Cluster      up   9000  auto/10000  -
-
e0c      Cluster      Cluster      up   9000  auto/10000  -
-
e0d      Cluster      Cluster      up   9000  auto/10000  -
-

Node: n2

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e0a      Cluster      Cluster      up   9000  auto/10000  -
-
e0b      Cluster      Cluster      up   9000  auto/10000  -
-
e0c      Cluster      Cluster      up   9000  auto/10000  -
-
e0d      Cluster      Cluster      up   9000  auto/10000  -
-

8 entries were displayed.
```

4. On all nodes, revert the specific cluster LIFs to their home ports:


```
network interface revert
```

Show example

The following example shows the specific cluster LIFs being reverted to their home ports on nodes n1 and n2:

```
cluster::*> network interface revert -vserver Cluster -lif n1_clus1
cluster::*> network interface revert -vserver Cluster -lif n1_clus4
cluster::*> network interface revert -vserver Cluster -lif n2_clus1
cluster::*> network interface revert -vserver Cluster -lif n2_clus4
```

5. Verify that the interface is home:

```
network interface show
```

Show example

The following example shows the status of cluster interconnect interfaces is up and Is home for n1 and n2:

```
cluster::*> network interface show -role cluster
(network interface show)

Current Is Logical Status Network Current
Vserver Is Interface Admin/Oper Address/Mask Node
Port Home
-----
-----
Cluster
e0a true n1_clus1 up/up 10.10.0.1/24 n1
e0b true n1_clus2 up/up 10.10.0.2/24 n1
e0c true n1_clus3 up/up 10.10.0.3/24 n1
e0d true n1_clus4 up/up 10.10.0.4/24 n1
e0a true n2_clus1 up/up 10.10.0.5/24 n2
e0b true n2_clus2 up/up 10.10.0.6/24 n2
e0c true n2_clus3 up/up 10.10.0.7/24 n2
e0d true n2_clus4 up/up 10.10.0.8/24 n2
8 entries were displayed.
```

6. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		

n1		
3/5/2022 19:21:18 -06:00	n1_clus2	n2_clus1
3/5/2022 19:21:20 -06:00	n1_clus2	n2_clus2
		none
		none
n2		
3/5/2022 19:21:18 -06:00	n2_clus2	n1_clus1
3/5/2022 19:21:20 -06:00	n2_clus2	n1_clus2
		none
		none

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1      e0a 10.10.0.1
Cluster n1_clus2 n1      e0b 10.10.0.2
Cluster n1_clus3 n1      e0c 10.10.0.3
Cluster n1_clus4 n1      e0d 10.10.0.4
Cluster n2_clus1 n2      e0a 10.10.0.5
Cluster n2_clus2 n2      e0b 10.10.0.6
Cluster n2_clus3 n2      e0c 10.10.0.7
Cluster n2_clus4 n2      e0d 10.10.0.8

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
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 16 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 16 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 10.10.0.2 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 10.10.0.3 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 10.10.0.4 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
Larger than PMTU communication succeeds on 16 path(s)
RPC status:
4 paths up, 0 paths down (tcp check)
4 paths up, 0 paths down (udp check)

```

7. Expand the cluster by adding nodes to the Nexus 3132Q-V cluster switches.
8. Display the information about the devices in your configuration:

- network device-discovery show
- network port show -role cluster
- network interface show -role cluster
- system cluster-switch show

Show example

The following examples show nodes n3 and n4 with 40 GbE cluster ports connected to ports e1/7 and e1/8, respectively on both the Nexus 3132Q-V cluster switches, and both nodes have joined the cluster. The 40 GbE cluster interconnect ports used are e4a and e4e.

```
cluster::> network device-discovery show
      Local  Discovered
Node   Port    Device           Interface           Platform
-----
n1     /cdp
      e0a    C1               Ethernet1/1/1       N3K-
C3132Q-V
      e0b    C2               Ethernet1/1/1       N3K-
C3132Q-V
      e0c    C2               Ethernet1/1/2       N3K-
C3132Q-V
      e0d    C1               Ethernet1/1/2       N3K-
C3132Q-V
n2     /cdp
      e0a    C1               Ethernet1/1/3       N3K-
C3132Q-V
      e0b    C2               Ethernet1/1/3       N3K-
C3132Q-V
      e0c    C2               Ethernet1/1/4       N3K-
C3132Q-V
      e0d    C1               Ethernet1/1/4       N3K-
C3132Q-V
n3     /cdp
      e4a    C1               Ethernet1/7         N3K-
C3132Q-V
      e4e    C2               Ethernet1/7         N3K-
C3132Q-V
n4     /cdp
      e4a    C1               Ethernet1/8         N3K-
C3132Q-V
      e4e    C2               Ethernet1/8         N3K-
C3132Q-V
12 entries were displayed.
```

```
cluster::*> network port show -role cluster
(network port show)
Node: n1
```

```

Ignore
                                                    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

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper  Speed (Mbps)
Status      Status

-----
e4a         Cluster    Cluster    up    9000 auto/40000  -
-
e4e         Cluster    Cluster    up    9000 auto/40000  -
-
12 entries were displayed.

```



```

cluster::*> network interface show -role cluster
(network interface show)

```

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
e0a	n1_clus1	up/up	10.10.0.1/24	n1
e0b	n1_clus2	up/up	10.10.0.2/24	n1
e0c	n1_clus3	up/up	10.10.0.3/24	n1
e0d	n1_clus4	up/up	10.10.0.4/24	n1
e0a	n2_clus1	up/up	10.10.0.5/24	n2
e0b	n2_clus2	up/up	10.10.0.6/24	n2
e0c	n2_clus3	up/up	10.10.0.7/24	n2
e0d	n2_clus4	up/up	10.10.0.8/24	n2
e4a	n3_clus1	up/up	10.10.0.9/24	n3
e4e	n3_clus2	up/up	10.10.0.10/24	n3
e4a	n4_clus1	up/up	10.10.0.11/24	n4
e4e	n4_clus2	up/up	10.10.0.12/24	n4

12 entries were displayed.

```
cluster::*> system cluster-switch show
```

Switch Model	Type	Address

C1 NX3132V	cluster-network	10.10.1.103
Serial Number: FOX000001		
Is Monitored: true		
Reason:		
Software Version: Cisco Nexus Operating System (NX-OS)		
Software, Version		
7.0(3)I4(1)		
Version Source: CDP		
C2 NX3132V	cluster-network	10.10.1.104
Serial Number: FOX000002		
Is Monitored: true		
Reason:		
Software Version: Cisco Nexus Operating System (NX-OS)		
Software, Version		
7.0(3)I4(1)		
Version Source: CDP		
CL1 NX5596	cluster-network	10.10.1.101
Serial Number: 01234567		
Is Monitored: true		
Reason:		
Software Version: Cisco Nexus Operating System (NX-OS)		
Software, Version		
7.1(1)N1(1)		
Version Source: CDP		
CL2 NX5596	cluster-network	10.10.1.102
Serial Number: 01234568		
Is Monitored: true		
Reason:		
Software Version: Cisco Nexus Operating System (NX-OS)		
Software, Version		
7.1(1)N1(1)		
Version Source: CDP		

4 entries were displayed.

9. Remove the replaced Nexus 5596 if they are not automatically removed:

```
system cluster-switch delete
```

Show example

The following example shows how to remove the Nexus 5596:

```
cluster::> system cluster-switch delete -device CL1
cluster::> system cluster-switch delete -device CL2
```

10. Configure clusters clus1 and clus2 to auto revert on each node and confirm.

Show example

```
cluster::*> network interface modify -vserver node1 -lif clus1 -auto
-revert true
cluster::*> network interface modify -vserver node1 -lif clus2 -auto
-revert true
cluster::*> network interface modify -vserver node2 -lif clus1 -auto
-revert true
cluster::*> network interface modify -vserver node2 -lif clus2 -auto
-revert true
```

11. Verify that the proper cluster switches are monitored:

```
system cluster-switch show
```

Show example

```
cluster::> system cluster-switch show

Switch                               Type                               Address
Model
-----
C1                                   cluster-network                   10.10.1.103
NX3132V
  Serial Number: FOX000001
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                    7.0(3)I4(1)
  Version Source: CDP

C2                                   cluster-network                   10.10.1.104
NX3132V
  Serial Number: FOX000002
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                    7.0(3)I4(1)
  Version Source: CDP

2 entries were displayed.
```

12. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next?

[Configure switch health monitoring.](#)

Migrate from switchless clusters to two-node switched clusters

Migrate from switchless clusters to two-node switched clusters workflow

Follow these workflow steps to migrate from a two-node switchless cluster to a two-node switched cluster that includes Cisco Nexus 3132Q-V cluster network switches.

1

Migration requirements

Review the requirements and example switch information for the migration process.

2

Prepare for migration

Prepare your switchless clusters for migration to two-node switched clusters.

3

Configure your ports

Configure your ports for migration from two-node switchless clusters to two-node switched clusters.

4

Complete your migration

Complete your migration from switchless clusters to two-node switched clusters.

Migration requirements

If you have a two-node switchless cluster, review this procedure for the applicable requirements to migrate to a two-node switched cluster.



The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

For more information, see:

- [NetApp CN1601 and CN1610 description page](#)
- [Cisco Ethernet Switch description page](#)
- [Hardware Universe](#)

Port and node connections

Make sure you understand the port and node connections and cabling requirements when you migrate to a two-node switched cluster with Cisco Nexus 3132Q-V cluster switches.

- The cluster switches use the Inter-Switch Link (ISL) ports e1/31-32.
- The [Hardware Universe](#) contains information about supported cabling to Nexus 3132Q-V switches:
 - The nodes with 10 GbE cluster connections require QSFP optical modules with breakout fiber cables or QSFP to SFP+ copper break-out cables.
 - The nodes with 40 GbE cluster connections require supported QSFP/QSFP28 optical modules with fiber cables or QSFP/QSFP28 copper direct-attach cables.
 - The cluster switches use the appropriate ISL cabling: 2x QSFP28 fiber or copper direct-attach cables.
- On Nexus 3132Q-V, you can operate QSFP ports as either 40 Gb Ethernet or 4x10 Gb Ethernet modes.

By default, there are 32 ports in the 40 Gb Ethernet mode. These 40 Gb Ethernet ports are numbered in a 2-tuple naming convention. For example, the second 40 Gb Ethernet port is numbered as 1/2. The process of changing the configuration from 40 Gb Ethernet to 10 Gb Ethernet is called *breakout* and the process of changing the configuration from 10 Gb Ethernet to 40 Gb Ethernet is called *breakin*. When you break out a

40 Gb Ethernet port into 10 Gb Ethernet ports, the resulting ports are numbered using a 3-tuple naming convention. For example, the breakout ports of the second 40 Gb Ethernet port are numbered as 1/2/1, 1/2/2, 1/2/3, and 1/2/4.

- On the left side of Nexus 3132Q-V is a set of four SFP+ ports multiplexed to the first QSFP port.

By default, the RCF is structured to use the first QSFP port.

You can make four SFP+ ports active instead of a QSFP port for Nexus 3132Q-V by using the `hardware profile front portmode sfp-plus` command. Similarly, you can reset Nexus 3132Q-V to use a QSFP port instead of four SFP+ ports by using the `hardware profile front portmode qsfp` command.

- Make sure you configured some of the ports on Nexus 3132Q-V to run at 10 GbE or 40 GbE.

You can break-out the first six ports into 4x10 GbE mode by using the `interface breakout module 1 port 1-6 map 10g-4x` command. Similarly, you can regroup the first six QSFP+ ports from breakout configuration by using the `no interface breakout module 1 port 1-6 map 10g-4x` command.

- The number of 10 GbE and 40 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.

Before you begin

- Configurations properly set up and functioning.
- Nodes running ONTAP 9.4 or later.
- All cluster ports in the `up` state.
- The Cisco Nexus 3132Q-V cluster switch is supported.
- The existing cluster network configuration has:
 - The Nexus 3132 cluster infrastructure that is redundant and fully functional on both switches.
 - The latest RCF and NX-OS versions on your switches.

The [Cisco Ethernet Switches](#) page has information about the ONTAP and NX-OS versions supported in this procedure.

- Management connectivity on both switches.
- Console access to both switches.
- All cluster logical interfaces (LIFs) in the `up` state without being migrated.
- Initial customization of the switch.
- All the ISL ports enabled and cabled.

In addition, you must plan, migrate, and read the required documentation on 10 GbE and 40 GbE connectivity from nodes to Nexus 3132Q-V cluster switches.

About the examples used

The examples in this procedure use the following switch and node nomenclature:

- Nexus 3132Q-V cluster switches, C1 and C2.
- The nodes are n1 and n2.



The examples in this procedure use two nodes, each using two 40 GbE cluster interconnect ports **e4a** and **e4e**. The [Hardware Universe](#) has details about the cluster ports on your platforms.

This procedure covers the following scenarios:

- **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 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.



The procedure requires the use of both ONTAP commands and Cisco Nexus 3000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

- The cluster starts with two nodes connected and functioning in a two-node switchless cluster setting.
- The first cluster port is moved to C1.
- The second cluster port is moved to C2.
- The two-node switchless cluster option is disabled.

What's next?

[Prepare for migration.](#)

Prepare for migration from switchless clusters to switched clusters

Follow these steps to prepare your switchless cluster for migration to a two-node switched cluster.

Steps

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all - message MAINT=xh
```

x is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Determine the administrative or operational status for each cluster interface:
 - a. Display the network port attributes:

```
network port show
```

Show example

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health      Health      Speed (Mbps)
Port        IPspace     Broadcast Domain Link MTU Admin/Oper
Status      Status
-----
e4a         Cluster     Cluster     up   9000 auto/40000 -
-
e4e         Cluster     Cluster     up   9000 auto/40000 -
-

Node: n2

Ignore

Health      Health      Speed (Mbps)
Port        IPspace     Broadcast Domain Link MTU Admin/Oper
Status      Status
-----
e4a         Cluster     Cluster     up   9000 auto/40000 -
-
e4e         Cluster     Cluster     up   9000 auto/40000 -
-

4 entries were displayed.
```

b. Display information about the logical interfaces:

```
network interface show
```


Show example

```
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
e4a      true      n1_clus1   up/up         10.10.0.1/24  n1
e4e      true      n1_clus2   up/up         10.10.0.2/24  n1
e4a      true      n2_clus1   up/up         10.10.0.3/24  n2
e4e      true      n2_clus2   up/up         10.10.0.4/24  n2
4 entries were displayed.
```

3. Verify that the appropriate RCFs and image are installed on the new 3132Q-V switches as necessary for your requirements, and make any essential site customizations, such as users and passwords, network addresses, and so on.

You must prepare both switches at this time. If you need to upgrade the RCF and image software, you must follow these steps:

- a. Go to the [Cisco Ethernet Switches](#) page on the NetApp Support Site.
 - b. Note your switch and the required software versions in the table on that page.
 - c. Download the appropriate version of RCF.
 - d. Select **CONTINUE** on the **Description** page, accept the license agreement, and then follow the instructions on the **Download** page to download the RCF.
 - e. Download the appropriate version of the image software.
4. Select **CONTINUE** on the **Description** page, accept the license agreement, and then follow the instructions on the **Download** page to download the RCF.

What's next?

[Configure your ports.](#)

Configure your ports for migration from switchless clusters to switched clusters

Follow these steps to configure your ports for migration from two-node switchless clusters to two-node switched clusters.

Steps

1. On Nexus 3132Q-V switches C1 and C2, disable all node-facing ports C1 and C2, but do not disable the ISL ports.

Show example

The following example shows ports 1 through 30 being disabled on Nexus 3132Q-V cluster switches C1 and C2 using a configuration supported in RCF NX3132_RCF_v1.1_24p10g_26p40g.txt:

```
C1# copy running-config startup-config
[#####] 100%
Copy complete.
C1# configure
C1(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C1(config-if-range)# shutdown
C1(config-if-range)# exit
C1(config)# exit

C2# copy running-config startup-config
[#####] 100%
Copy complete.
C2# configure
C2(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C2(config-if-range)# shutdown
C2(config-if-range)# exit
C2(config)# exit
```

2. Connect ports 1/31 and 1/32 on C1 to the same ports on C2 using supported cabling.
3. Verify that the ISL ports are operational on C1 and C2:

```
show port-channel summary
```

Show example

```
C1# show port-channel summary
Flags: D - Down          P - Up in port-channel (members)
      I - Individual     H - Hot-standby (LACP only)
      s - Suspended      r - Module-removed
      S - Switched       R - Routed
      U - Up (port-channel)
      M - Not in use. Min-links not met

-----
-----
Group Port-          Type   Protocol  Member Ports
Channel
-----
-----
1      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)
```

4. Display the list of neighboring devices on the switch:

```
show cdp neighbors
```

Show example

```
C1# show cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-
Bridge
                S - Switch, H - Host, I - IGMP, r - Repeater,
                V - VoIP-Phone, D - Remotely-Managed-Device,
                s - Supports-STP-Dispute

Device-ID          Local Intrfce  Hldtme Capability  Platform
Port ID
C2                  Eth1/31        174    R S I s       N3K-C3132Q-V
Eth1/31
C2                  Eth1/32        174    R S I s       N3K-C3132Q-V
Eth1/32

Total entries displayed: 2

C2# show cdp neighbors
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-
Bridge
                S - Switch, H - Host, I - IGMP, r - Repeater,
                V - VoIP-Phone, D - Remotely-Managed-Device,
                s - Supports-STP-Dispute

Device-ID          Local Intrfce  Hldtme Capability  Platform
Port ID
C1                  Eth1/31        178    R S I s       N3K-C3132Q-V
Eth1/31
C1                  Eth1/32        178    R S I s       N3K-C3132Q-V
Eth1/32

Total entries displayed: 2
```

5. Display the cluster port connectivity on each node:

```
network device-discovery show
```

Show example

The following example shows a two-node switchless cluster configuration.

```
cluster::*> network device-discovery show
```

Node	Local Port	Discovered Device	Interface	Platform
n1	/cdp			
	e4a	n2	e4a	FAS9000
	e4e	n2	e4e	FAS9000
n2	/cdp			
	e4a	n1	e4a	FAS9000
	e4e	n1	e4e	FAS9000

6. Migrate the clus1 interface to the physical port hosting clus2:

```
network interface migrate
```

Execute this command from each local node.

Show example

```
cluster::*> network interface migrate -vserver Cluster -lif n1_clus1  
-source-node n1  
-destination-node n1 -destination-port e4e  
cluster::*> network interface migrate -vserver Cluster -lif n2_clus1  
-source-node n2  
-destination-node n2 -destination-port e4e
```

7. Verify the cluster interfaces migration:

```
network interface show
```

Show example

```
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
e4e      false
      n1_clus2    up/up      10.10.0.2/24    n1
e4e      true
      n2_clus1    up/up      10.10.0.3/24    n2
e4e      false
      n2_clus2    up/up      10.10.0.4/24    n2
e4e      true
4 entries were displayed.
```

8. Shut down cluster ports clus1 LIF on both nodes:

```
network port modify
```

```
cluster::*> network port modify -node n1 -port e4a -up-admin false
cluster::*> network port modify -node n2 -port e4a -up-admin false
```

9. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Node	Date	Loss	Source LIF	Destination LIF	
	n1	3/5/2022 19:21:18	-06:00	n1_clus2	n2_clus1	none
		3/5/2022 19:21:20	-06:00	n1_clus2	n2_clus2	none
	n2	3/5/2022 19:21:18	-06:00	n2_clus2	n1_clus1	none
		3/5/2022 19:21:20	-06:00	n2_clus2	n1_clus2	none

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1      e4a 10.10.0.1
Cluster n1_clus2 n1      e4e 10.10.0.2
Cluster n2_clus1 n2      e4a 10.10.0.3
Cluster n2_clus2 n2      e4e 10.10.0.4

Local = 10.10.0.1 10.10.0.2
Remote = 10.10.0.3 10.10.0.4
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 32 path(s):
    Local 10.10.0.1 to Remote 10.10.0.3
    Local 10.10.0.1 to Remote 10.10.0.4
    Local 10.10.0.2 to Remote 10.10.0.3
    Local 10.10.0.2 to Remote 10.10.0.4
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
1 paths up, 0 paths down (tcp check)
1 paths up, 0 paths down (ucp check)

```

10. Disconnect the cable from e4a on node n1.

You can refer to the running configuration and connect the first 40 GbE port on the switch C1 (port 1/7 in this example) to e4a on n1 using supported cabling on Nexus 3132Q-V.



When reconnecting any cables to a new Cisco cluster switch, the cables used must be either fiber or cabling supported by Cisco.

11. Disconnect the cable from e4a on node n2.

You can refer to the running configuration and connect e4a to the next available 40 GbE port on C1, port 1/8, using supported cabling.

12. Enable all node-facing ports on C1.

Show example

The following example shows ports 1 through 30 being enabled on Nexus 3132Q-V cluster switches C1 and C2 using the configuration supported in RCF NX3132_RCF_v1.1_24p10g_26p40g.txt:

```
C1# configure
C1(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C1(config-if-range)# no shutdown
C1(config-if-range)# exit
C1(config)# exit
```

13. Enable the first cluster port, e4a, on each node:

```
network port modify
```

Show example

```
cluster::*> network port modify -node n1 -port e4a -up-admin true
cluster::*> network port modify -node n2 -port e4a -up-admin true
```

14. Verify that the clusters are up on both nodes:

```
network port show
```

Show example

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e4a      Cluster      Cluster      up    9000 auto/40000  -
-
e4e      Cluster      Cluster      up    9000 auto/40000  -
-

Node: n2

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e4a      Cluster      Cluster      up    9000 auto/40000  -
-
e4e      Cluster      Cluster      up    9000 auto/40000  -
-

4 entries were displayed.
```

15. For each node, revert all of the migrated cluster interconnect LIFs:

```
network interface revert
```

Show example

The following example shows the migrated LIFs being reverted to their home ports.

```
cluster::*> network interface revert -vserver Cluster -lif n1_clus1
cluster::*> network interface revert -vserver Cluster -lif n2_clus1
```

16. Verify that all of the cluster interconnect ports are now reverted to their home ports:

```
network interface show
```

The `Is Home` column should display a value of `true` for all of the ports listed in the `Current Port` column. If the displayed value is `false`, the port has not been reverted.

Show example

```
cluster::*> network interface show -role cluster
(network interface show)
Current Is Logical Status Network Current
Vserver Interface Admin/Oper Address/Mask Node
Port Home
-----
Cluster
e4a true n1_clus1 up/up 10.10.0.1/24 n1
e4e true n1_clus2 up/up 10.10.0.2/24 n1
e4a true n2_clus1 up/up 10.10.0.3/24 n2
e4e true n2_clus2 up/up 10.10.0.4/24 n2
4 entries were displayed.
```

17. Display the cluster port connectivity on each node:

```
network device-discovery show
```

Show example

```
cluster::*> network device-discovery show
          Local  Discovered
Node      Port    Device           Interface      Platform
-----
n1        /cdp
          e4a    C1               Ethernet1/7    N3K-C3132Q-V
          e4e    n2               e4e            FAS9000
n2        /cdp
          e4a    C1               Ethernet1/8    N3K-C3132Q-V
          e4e    n1               e4e            FAS9000
```

18. On the console of each node, migrate clus2 to port e4a:

```
network interface migrate
```

Show example

```
cluster::*> network interface migrate -vserver Cluster -lif n1_clus2
-source-node n1
-destination-node n1 -destination-port e4a
cluster::*> network interface migrate -vserver Cluster -lif n2_clus2
-source-node n2
-destination-node n2 -destination-port e4a
```

19. Shut down cluster ports clus2 LIF on both nodes:

```
network port modify
```

The following example shows the specified ports being shut down on both nodes:

```
cluster::*> network port modify -node n1 -port e4e -up-admin false
cluster::*> network port modify -node n2 -port e4e -up-admin false
```

20. Verify the cluster LIF status:

```
network interface show
```

Show example

```
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
e4a       true
          n1_clus2    up/up      10.10.0.2/24  n1
e4a       false
          n2_clus1    up/up      10.10.0.3/24  n2
e4a       true
          n2_clus2    up/up      10.10.0.4/24  n2
e4a       false
4 entries were displayed.
```

21. Disconnect the cable from e4e on node n1.

You can refer to the running configuration and connect the first 40 GbE port on the switch C2 (port 1/7 in this example) to e4e on n1 using supported cabling on Nexus 3132Q-V.

22. Disconnect the cable from e4e on node n2.

You can refer to the running configuration and connect e4e to the next available 40 GbE port on C2, port 1/8, using supported cabling.

23. Enable all node-facing ports on C2.

Show example

The following example shows ports 1 through 30 being enabled on Nexus 3132Q-V cluster switches C1 and C2 using a configuration supported in RCF NX3132_RCF_v1.1_24p10g_26p40g.txt:

```
C2# configure
C2(config)# int e1/1/1-4,e1/2/1-4,e1/3/1-4,e1/4/1-4,e1/5/1-4,e1/6/1-4,e1/7-30
C2(config-if-range)# no shutdown
C2(config-if-range)# exit
C2(config)# exit
```

24. Enable the second cluster port, e4e, on each node:

```
network port modify
```

The following example shows the specified ports being brought up:

```
cluster::*> network port modify -node n1 -port e4e -up-admin true
cluster::*> network port modify -node n2 -port e4e -up-admin true
```

25. 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
```

26. Verify that all of the cluster interconnect ports are now reverted to their home ports:

```
network interface show
```

The `Is Home` column should display a value of `true` for all of the ports listed in the `Current Port` column. If the displayed value is `false`, the port has not been reverted.

Show example

```
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
e4a          n1_clus1   up/up       10.10.0.1/24  n1
true
e4e          n1_clus2   up/up       10.10.0.2/24  n1
true
e4a          n2_clus1   up/up       10.10.0.3/24  n2
true
e4e          n2_clus2   up/up       10.10.0.4/24  n2
true
4 entries were displayed.
```

27. Verify that all of the cluster interconnect ports are in the up state.

```
network port show -role cluster
```

Show example

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e4a      Cluster      Cluster      up    9000 auto/40000  -
-
e4e      Cluster      Cluster      up    9000 auto/40000  -
-

Node: n2

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
-----
e4a      Cluster      Cluster      up    9000 auto/40000  -
-
e4e      Cluster      Cluster      up    9000 auto/40000  -
-

4 entries were displayed.
```

What's next?

[Complete your migration.](#)

Complete the migration from two-node switchless clusters to two-node switched clusters

Follow these steps to complete the migration from switchless clusters to two-node switched clusters.

Steps

1. Display the cluster switch port numbers each cluster port is connected to on each node:

```
network device-discovery show
```


Show example

```
cluster::*> network device-discovery show
```

Node	Local Port	Discovered Device	Interface	Platform
n1	/cdp			
	e4a	C1	Ethernet1/7	N3K-C3132Q-V
	e4e	C2	Ethernet1/7	N3K-C3132Q-V
n2	/cdp			
	e4a	C1	Ethernet1/8	N3K-C3132Q-V
	e4e	C2	Ethernet1/8	N3K-C3132Q-V

2. Display discovered and monitored cluster switches:

```
system cluster-switch show
```

Show example

```
cluster::*> system cluster-switch show

Switch                               Type                               Address
Model
-----
C1                                   cluster-network                    10.10.1.101
NX3132V
  Serial Number: FOX000001
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        7.0(3)I4(1)
  Version Source: CDP

C2                                   cluster-network                    10.10.1.102
NX3132V
  Serial Number: FOX000002
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        7.0(3)I4(1)
  Version Source: CDP

2 entries were displayed.
```

3. Disable the two-node switchless configuration settings on any node:

```
network options switchless-cluster
```

```
network options switchless-cluster modify -enabled false
```

4. Verify that the switchless-cluster option has been disabled.

```
network options switchless-cluster show
```

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

				Source	Destination	
Packet				LIF	LIF	
Node	Date					
Loss						

n1						
	3/5/2022	19:21:18	-06:00	n1_clus2	n2_clus1	none
	3/5/2022	19:21:20	-06:00	n1_clus2	n2_clus2	none
n2						
	3/5/2022	19:21:18	-06:00	n2_clus2	n1_clus1	none
	3/5/2022	19:21:20	-06:00	n2_clus2	n1_clus2	none

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1      e4a 10.10.0.1
Cluster n1_clus2 n1      e4e 10.10.0.2
Cluster n2_clus1 n2      e4a 10.10.0.3
Cluster n2_clus2 n2      e4e 10.10.0.4

Local = 10.10.0.1 10.10.0.2
Remote = 10.10.0.3 10.10.0.4
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 32 path(s):
    Local 10.10.0.1 to Remote 10.10.0.3
    Local 10.10.0.1 to Remote 10.10.0.4
    Local 10.10.0.2 to Remote 10.10.0.3
    Local 10.10.0.2 to Remote 10.10.0.4
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
1 paths up, 0 paths down (tcp check)
1 paths up, 0 paths down (ucp check)

```

6. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next?

[Configure switch health monitoring.](#)

Replace switches

Requirements for replacing Cisco Nexus 3132Q-V cluster switches

Make sure you understand the configuration requirements, port connections, and cabling requirements when you replace cluster switches.

Cisco Nexus 3132Q-V requirements

- The Cisco Nexus 3132Q-V cluster switch is supported.
- The number of 10 GbE and 40 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.

- The cluster switches use the Inter-Switch Link (ISL) ports e1/31-32.
- The [Hardware Universe](#) contains information about supported cabling to Nexus 3132Q-V switches:
 - The nodes with 10 GbE cluster connections require QSFP optical modules with breakout fiber cables or QSFP to SFP+ copper break-out cables.
 - The nodes with 40 GbE cluster connections require supported QSFP/QSFP28 optical modules with fiber cables or QSFP/QSFP28 copper direct-attach cables.
 - The cluster switches use the appropriate ISL cabling: 2x QSFP28 fiber or copper direct-attach cables.
- On Nexus 3132Q-V, you can operate QSFP ports as either 40 Gb Ethernet or 4x10 Gb Ethernet modes.

By default, there are 32 ports in the 40 Gb Ethernet mode. These 40 Gb Ethernet ports are numbered in a 2-tuple naming convention. For example, the second 40 Gb Ethernet port is numbered as 1/2. The process of changing the configuration from 40 Gb Ethernet to 10 Gb Ethernet is called *breakout* and the process of changing the configuration from 10 Gb Ethernet to 40 Gb Ethernet is called *breakin*. When you break out a 40 Gb Ethernet port into 10 Gb Ethernet ports, the resulting ports are numbered using a 3-tuple naming convention. For example, the breakout ports of the second 40 Gb Ethernet port are numbered as 1/2/1, 1/2/2, 1/2/3, and 1/2/4.

- On the left side of Nexus 3132Q-V is a set of four SFP+ ports multiplexed to the first QSFP port.

By default, the RCF is structured to use the first QSFP port.

You can make four SFP+ ports active instead of a QSFP port for Nexus 3132Q-V by using the `hardware profile front portmode sfp-plus` command. Similarly, you can reset Nexus 3132Q-V to use a QSFP port instead of four SFP+ ports by using the `hardware profile front portmode qsfp` command.

- You must have configured some of the ports on Nexus 3132Q-V to run at 10 GbE or 40 GbE.

You can break-out the first six ports into 4x10 GbE mode by using the `interface breakout module 1 port 1-6 map 10g-4x` command. Similarly, you can regroup the first six QSFP+ ports from breakout configuration by using the `no interface breakout module 1 port 1-6 map 10g-4x` command.

- You must have done the planning, migration, and read the required documentation on 10 GbE and 40 GbE connectivity from nodes to Nexus 3132Q-V cluster switches.

The [Cisco Ethernet Switches](#) page has information about the ONTAP and NX-OS versions supported in this procedure.

Cisco Nexus 5596 requirements

- The following cluster switches are supported:
 - Nexus 5596
 - Nexus 3132Q-V
- The number of 10 GbE and 40 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.
- The cluster switches use the following ports for connections to nodes:
 - Ports e1/1-40 (10 GbE): Nexus 5596
 - Ports e1/1-30 (40 GbE): Nexus 3132Q-V

- The cluster switches use the following Inter-Switch Link (ISL) ports:
 - Ports e1/41-48 (10 GbE): Nexus 5596
 - Ports e1/31-32 (40 GbE): Nexus 3132Q-V
- The [Hardware Universe](#) contains information about supported cabling to Nexus 3132Q-V switches:
 - Nodes with 10 GbE cluster connections require QSFP to SFP+ optical fiber breakout cables or QSFP to SFP+ copper breakout cables.
 - Nodes with 40 GbE cluster connections require supported QSFP/QSFP28 optical modules with fiber cables or QSFP/QSFP28 copper direct-attach cables.
- The cluster switches use the appropriate ISL cabling:
 - Beginning: Nexus 5596 to Nexus 5596 (SFP+ to SFP+)
 - 8x SFP+ fiber or copper direct-attach cables
 - Interim: Nexus 5596 to Nexus 3132Q-V (QSFP to 4xSFP+ break-out)
 - 1x QSFP to SFP+ fiber break-out or copper break-out cables
 - Final: Nexus 3132Q-V to Nexus 3132Q-V (QSFP28 to QSFP28)
 - 2x QSFP28 fiber or copper direct-attach cables
- On Nexus 3132Q-V switches, you can operate QSFP/QSFP28 ports as either 40 Gigabit Ethernet or 4 x10 Gigabit Ethernet modes.

By default, there are 32 ports in the 40 Gigabit Ethernet mode. These 40 Gigabit Ethernet ports are numbered in a 2-tuple naming convention. For example, the second 40 Gigabit Ethernet port is numbered as 1/2. The process of changing the configuration from 40 Gigabit Ethernet to 10 Gigabit Ethernet is called *breakout* and the process of changing the configuration from 10 Gigabit Ethernet to 40 Gigabit Ethernet is called *breakin*. When you break out a 40 Gigabit Ethernet port into 10 Gigabit Ethernet ports, the resulting ports are numbered using a 3-tuple naming convention. For example, the break-out ports of the second 40 Gigabit Ethernet port are numbered as 1/2/1, 1/2/2, 1/2/3, and 1/2/4.

- On the left side of Nexus 3132Q-V switches is a set of 4 SFP+ ports multiplexed to that QSFP28 port.

By default, the RCF is structured to use the QSFP28 port.



You can make 4x SFP+ ports active instead of a QSFP port for Nexus 3132Q-V switches by using the `hardware profile front portmode sfp-plus` command. Similarly, you can reset Nexus 3132Q-V switches to use a QSFP port instead of 4x SFP+ ports by using the `hardware profile front portmode qsfp` command.

- You have configured some of the ports on Nexus 3132Q-V switches to run at 10 GbE or 40 GbE.



You can break out the first six ports into 4x10 GbE mode by using the `interface breakout module 1 port 1-6 map 10g-4x` command. Similarly, you can regroup the first six QSFP+ ports from breakout configuration by using the `no interface breakout module 1 port 1-6 map 10g-4x` command.

- You have done the planning, migration, and read the required documentation on 10 GbE and 40 GbE connectivity from nodes to Nexus 3132Q-V cluster switches.
- The ONTAP and NX-OS versions supported in this procedure are on the [Cisco Ethernet Switches](#) page.

NetApp CN1610 requirements

- The following cluster switches are supported:
 - NetApp CN1610
 - Cisco Nexus 3132Q-V
- The cluster switches support the following node connections:
 - NetApp CN1610: ports 0/1 through 0/12 (10 GbE)
 - Cisco Nexus 3132Q-V: ports e1/1-30 (40 GbE)
- The cluster switches use the following inter-switch link (ISL) ports:
 - NetApp CN1610: ports 0/13 through 0/16 (10 GbE)
 - Cisco Nexus 3132Q-V: ports e1/31-32 (40 GbE)
- The [Hardware Universe](#) contains information about supported cabling to Nexus 3132Q-V switches:
 - Nodes with 10 GbE cluster connections require QSFP to SFP+ optical fiber breakout cables or QSFP to SFP+ copper breakout cables
 - Nodes with 40 GbE cluster connections require supported QSFP/QSFP28 optical modules with optical fiber cables or QSFP/QSFP28 copper direct-attach cables
- The appropriate ISL cabling is as follows:
 - Beginning: For CN1610 to CN1610 (SFP+ to SFP+), four SFP+ optical fiber or copper direct-attach cables
 - Interim: For CN1610 to Nexus 3132Q-V (QSFP to four SFP+ breakout), one QSFP to SFP+ optical fiber or copper breakout cable
 - Final: For Nexus 3132Q-V to Nexus 3132Q-V (QSFP28 to QSFP28), two QSFP28 optical fiber or copper direct-attach cables
- NetApp twinax cables are not compatible with Cisco Nexus 3132Q-V switches.

If your current CN1610 configuration uses NetApp twinax cables for cluster-node-to-switch connections or ISL connections and you want to continue using twinax in your environment, you need to procure Cisco twinax cables. Alternatively, you can use optical fiber cables for both the ISL connections and the cluster-node-to-switch connections.

- On Nexus 3132Q-V switches, you can operate QSFP/QSFP28 ports as either 40 Gb Ethernet or 4x 10 Gb Ethernet modes.

By default, there are 32 ports in the 40 Gb Ethernet mode. These 40 Gb Ethernet ports are numbered in a 2-tuple naming convention. For example, the second 40 Gb Ethernet port is numbered as 1/2. The process of changing the configuration from 40 Gb Ethernet to 10 Gb Ethernet is called *breakout* and the process of changing the configuration from 10 Gb Ethernet to 40 Gb Ethernet is called *breakin*. When you break out a 40 Gb Ethernet port into 10 Gb Ethernet ports, the resulting ports are numbered using a 3-tuple naming convention. For example, the breakout ports of the second 40 Gb Ethernet port are numbered as 1/2/1, 1/2/2, 1/2/3, and 1/2/4.

- On the left side of Nexus 3132Q-V switches is a set of four SFP+ ports multiplexed to the first QSFP port.

By default, the reference configuration file (RCF) is structured to use the first QSFP port.

You can make four SFP+ ports active instead of a QSFP port for Nexus 3132Q-V switches by using the hardware profile `front portmode sfp-plus` command. Similarly, you can reset Nexus 3132Q-V switches to use a QSFP port instead of four SFP+ ports by using the hardware profile `front`

portmode qsfp command.



When you use the first four SFP+ ports, it will disable the first 40GbE QSFP port.

- You must have configured some of the ports on Nexus 3132Q-V switches to run at 10 GbE or 40 GbE.

You can break out the first six ports into 4x10 GbE mode by using the `interface breakout module 1 port 1-6 map 10g-4x` command. Similarly, you can regroup the first six QSFP+ ports from *breakout* configuration by using the `no interface breakout module 1 port 1-6 map 10g-4x` command.

- You must have done the planning, migration, and read the required documentation on 10 GbE and 40 GbE connectivity from nodes to Nexus 3132Q-V cluster switches.
- The ONTAP and NX-OS versions that are supported in this procedure are listed on the [Cisco Ethernet Switches](#) page.
- The ONTAP and FASTPATH versions that are supported in this procedure are listed on the [NetApp CN1601 and CN1610 Switches](#) page.

Replace Cisco Nexus 3132Q-V cluster switches

Follow this procedure to replace a defective Cisco Nexus 3132Q-V switch in a cluster network. The replacement procedure is a nondisruptive procedure (NDO).

Review requirements

Switch requirements

Review the [Requirements for replacing Cisco Nexus 3132Q-V cluster switches](#).

Before you begin

- The existing cluster and network configuration has:
 - The Nexus 3132Q-V cluster infrastructure is redundant and fully functional on both switches.
The [Cisco Ethernet Switch](#) page has the latest RCF and NX-OS versions on your switches.
 - All cluster ports are in the `up` state.
 - Management connectivity exists on both switches.
 - All cluster logical interfaces (LIFs) are in the `up` state and have been migrated.
- For the Nexus 3132Q-V replacement switch, make sure that:
 - Management network connectivity on the replacement switch is functional.
 - Console access to the replacement switch is in place.
 - The desired RCF and NX-OS operating system image switch is loaded onto the switch.
 - Initial customization of the switch is complete.
- [Hardware Universe](#)

Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when replacing your switch:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport before and after maintenance to disable case creation for the duration of the maintenance. See this Knowledge Base article [SU92: How to suppress automatic case creation during scheduled maintenance windows](#) for further details.
- Enable session logging for any CLI sessions. For instructions on how to enable session logging, review the "Logging Session Output" section in this Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

Replace the switch

This procedure replaces the second Nexus 3132Q-V cluster switch CL2 with new 3132Q-V switch C2.

About the examples

The examples in this procedure use the following switch and node nomenclature:

- n1_clus1 is the first cluster logical interface (LIF) connected to cluster switch C1 for node n1.
- n1_clus2 is the first cluster LIF connected to cluster switch CL2 or C2, for node n1.
- n1_clus3 is the second LIF connected to cluster switch C2, for node n1.
- n1_clus4 is the second LIF connected to cluster switch CL1, for node n1.
- The number of 10 GbE and 40 GbE ports are defined in the reference configuration files (RCFs) available on the [Cisco® Cluster Network Switch Reference Configuration File Download](#) page.
- The nodes are n1, n2, n3, and n4.

-

The examples in this procedure use four nodes: Two nodes use four 10 GB cluster interconnect ports: e0a, e0b, e0c, and e0d. The other two nodes use two 40 GB cluster interconnect ports: e4a and e4e. See the [Hardware Universe](#) for the actual cluster ports on your platforms.

About this task

This procedure covers the following scenario:

- The cluster starts with four nodes connected to two Nexus 3132Q-V cluster switches, CL1 and CL2.
- Cluster switch CL2 is to be replaced by C2
 - On each node, cluster LIFs connected to CL2 are migrated onto cluster ports connected to CL1.
 - Disconnect cabling from all ports on CL2 and reconnect cabling to the same ports on the replacement switch C2.
 - On each node, its migrated cluster LIFs are reverted.

Step 1: Prepare for replacement

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all - message MAINT=xh
```

x is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Display information about the devices in your configuration:

```
network device-discovery show
```

Show example

```
cluster::> network device-discovery show
```

Node	Local Port	Discovered Device	Interface	Platform
n1	/cdp			
	e0a	CL1	Ethernet1/1/1	N3K-C3132Q-V
	e0b	CL2	Ethernet1/1/1	N3K-C3132Q-V
	e0c	CL2	Ethernet1/1/2	N3K-C3132Q-V
	e0d	CL1	Ethernet1/1/2	N3K-C3132Q-V
n2	/cdp			
	e0a	CL1	Ethernet1/1/3	N3K-C3132Q-V
	e0b	CL2	Ethernet1/1/3	N3K-C3132Q-V
	e0c	CL2	Ethernet1/1/4	N3K-C3132Q-V
	e0d	CL1	Ethernet1/1/4	N3K-C3132Q-V
n3	/cdp			
	e4a	CL1	Ethernet1/7	N3K-C3132Q-V
	e4e	CL2	Ethernet1/7	N3K-C3132Q-V
n4	/cdp			
	e4a	CL1	Ethernet1/8	N3K-C3132Q-V
	e4e	CL2	Ethernet1/8	N3K-C3132Q-V

12 entries were displayed

3. Determine the administrative or operational status for each cluster interface:

a. Display the network port attributes:

```
network port show
```

Show example

```
cluster::*> network port show -role cluster
(network port show)

Node: n1

Ignore

Health      Health      Speed (Mbps)
Port        IPspace     Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
e0a         Cluster     Cluster     up   9000 auto/10000 -
-
e0b         Cluster     Cluster     up   9000 auto/10000 -
-
e0c         Cluster     Cluster     up   9000 auto/10000 -
-
e0d         Cluster     Cluster     up   9000 auto/10000 -
-

Node: n2

Ignore

Health      Health      Speed (Mbps)
Port        IPspace     Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
e0a         Cluster     Cluster     up   9000 auto/10000 -
-
e0b         Cluster     Cluster     up   9000 auto/10000 -
-
e0c         Cluster     Cluster     up   9000 auto/10000 -
-
e0d         Cluster     Cluster     up   9000 auto/10000 -
-

Node: n3

Ignore

Health      Health      Speed (Mbps)
```

```
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper
Status    Status
-----
-----
e4a      Cluster      Cluster      up   9000 auto/40000 -
-
e4e      Cluster      Cluster      up   9000 auto/40000 -
-

Node: n4

Ignore

Health    Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper
Status    Status
-----
-----
e4a      Cluster      Cluster      up   9000 auto/40000 -
-
e4e      Cluster      Cluster      up   9000 auto/40000 -
-

Speed (Mbps)

12 entries were displayed.
```

b. Display information about the logical interfaces:

```
network interface show
```

Show example

```
cluster::*> network interface show -role cluster
(network interface show)
```

Current Is	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node

Cluster				
e0a	n1_clus1	up/up	10.10.0.1/24	n1
e0b	n1_clus2	up/up	10.10.0.2/24	n1
e0c	n1_clus3	up/up	10.10.0.3/24	n1
e0d	n1_clus4	up/up	10.10.0.4/24	n1
e0a	n2_clus1	up/up	10.10.0.5/24	n2
e0b	n2_clus2	up/up	10.10.0.6/24	n2
e0c	n2_clus3	up/up	10.10.0.7/24	n2
e0d	n2_clus4	up/up	10.10.0.8/24	n2
e0a	n3_clus1	up/up	10.10.0.9/24	n3
e0e	n3_clus2	up/up	10.10.0.10/24	n3
e0a	n4_clus1	up/up	10.10.0.11/24	n4
e0e	n4_clus2	up/up	10.10.0.12/24	n4

12 entries were displayed.

c. Display the information on the discovered cluster switches:

```
system cluster-switch show
```

Show example

```
cluster::> system cluster-switch show

Switch                               Type                               Address
Model
-----
CL1                                   cluster-network                    10.10.1.101
NX3132V
  Serial Number: FOX000001
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                               7.0(3)I4(1)
  Version Source: CDP

CL2                                   cluster-network                    10.10.1.102
NX3132V
  Serial Number: FOX000002
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                               7.0(3)I4(1)
  Version Source: CDP

2 entries were displayed.
```

4. Verify that the appropriate RCF and image are installed on the new Nexus 3132Q-V switch as necessary for your requirements, and make any essential site customizations.

You must prepare the replacement switch at this time. If you need to upgrade the RCF and image, you must follow these steps:

- a. On the NetApp Support Site, go to the [Cisco Ethernet Switch](#) page.
 - b. Note your switch and the required software versions in the table on that page.
 - c. Download the appropriate version of the RCF.
 - d. Click **CONTINUE** on the **Description** page, accept the license agreement, and then follow the instructions on the **Download** page to download the RCF.
 - e. Download the appropriate version of the image software.
5. Migrate the LIFs associated to the cluster ports connected to switch C2:

```
network interface migrate
```

Show example

This example shows that the LIF migration is done on all the nodes:

```
cluster::*> network interface migrate -vserver Cluster -lif n1_clus2
-source-node n1 -destination-node n1 -destination-port e0a
cluster::*> network interface migrate -vserver Cluster -lif n1_clus3
-source-node n1 -destination-node n1 -destination-port e0d
cluster::*> network interface migrate -vserver Cluster -lif n2_clus2
-source-node n2 -destination-node n2 -destination-port e0a
cluster::*> network interface migrate -vserver Cluster -lif n2_clus3
-source-node n2 -destination-node n2 -destination-port e0d
cluster::*> network interface migrate -vserver Cluster -lif n3_clus2
-source-node n3 -destination-node n3 -destination-port e4a
cluster::*> network interface migrate -vserver Cluster -lif n4_clus2
-source-node n4 -destination-node n4 -destination-port e4a
```

6. Verify cluster's health:

```
network interface show
```

Show example

```
cluster::*> network interface show -role cluster
(network interface show)
```

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
e0a	n1_clus1	up/up	10.10.0.1/24	n1
e0a	true			
e0a	n1_clus2	up/up	10.10.0.2/24	n1
e0a	false			
e0d	n1_clus3	up/up	10.10.0.3/24	n1
e0d	false			
e0d	n1_clus4	up/up	10.10.0.4/24	n1
e0d	true			
e0a	n2_clus1	up/up	10.10.0.5/24	n2
e0a	true			
e0a	n2_clus2	up/up	10.10.0.6/24	n2
e0a	false			
e0d	n2_clus3	up/up	10.10.0.7/24	n2
e0d	false			
e0d	n2_clus4	up/up	10.10.0.8/24	n2
e0d	true			
e4a	n3_clus1	up/up	10.10.0.9/24	n3
e4a	true			
e4a	n3_clus2	up/up	10.10.0.10/24	n3
e4a	false			
e4a	n4_clus1	up/up	10.10.0.11/24	n4
e4a	true			
e4a	n4_clus2	up/up	10.10.0.12/24	n4
e4a	false			

12 entries were displayed.

7. Shut down the cluster interconnect ports that are physically connected to switch CL2:

```
network port modify
```


Show example

This example shows the specified ports being shut down on all nodes:

```
cluster::*> network port modify -node n1 -port e0b -up-admin false
cluster::*> network port modify -node n1 -port e0c -up-admin false
cluster::*> network port modify -node n2 -port e0b -up-admin false
cluster::*> network port modify -node n2 -port e0c -up-admin false
cluster::*> network port modify -node n3 -port e4e -up-admin false
cluster::*> network port modify -node n4 -port e4e -up-admin false
```

8. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Node	Date		Source LIF	Destination LIF	Packet Loss
n1					
	3/5/2022 19:21:18 -06:00		n1_clus2	n2_clus1	none
	3/5/2022 19:21:20 -06:00		n1_clus2	n2_clus2	none
n2					
	3/5/2022 19:21:18 -06:00		n2_clus2	n1_clus1	none
	3/5/2022 19:21:20 -06:00		n2_clus2	n1_clus2	none
n3					
...					
...					
n4					
...					
...					

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```
cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1 e0a 10.10.0.1
Cluster n1_clus2 n1 e0b 10.10.0.2
Cluster n1_clus3 n1 e0c 10.10.0.3
Cluster n1_clus4 n1 e0d 10.10.0.4
Cluster n2_clus1 n2 e0a 10.10.0.5
Cluster n2_clus2 n2 e0b 10.10.0.6
```

```
Cluster n2_clus3 n2      e0c 10.10.0.7
Cluster n2_clus4 n2      e0d 10.10.0.8
Cluster n3_clus1 n4      e0a 10.10.0.9
Cluster n3_clus2 n3      e0e 10.10.0.10
Cluster n4_clus1 n4      e0a 10.10.0.11
Cluster n4_clus2 n4      e0e 10.10.0.12

Local = 10.10.0.1 10.10.0.2 10.10.0.3 10.10.0.4
Remote = 10.10.0.5 10.10.0.6 10.10.0.7 10.10.0.8 10.10.0.9 10.10.0.10
10.10.0.11 10.10.0.12
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 32 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 32 path(s):
  Local 10.10.0.1 to Remote 10.10.0.5
  Local 10.10.0.1 to Remote 10.10.0.6
  Local 10.10.0.1 to Remote 10.10.0.7
  Local 10.10.0.1 to Remote 10.10.0.8
  Local 10.10.0.1 to Remote 10.10.0.9
  Local 10.10.0.1 to Remote 10.10.0.10
  Local 10.10.0.1 to Remote 10.10.0.11
  Local 10.10.0.1 to Remote 10.10.0.12
  Local 10.10.0.2 to Remote 10.10.0.5
  Local 10.10.0.2 to Remote 10.10.0.6
  Local 10.10.0.2 to Remote 10.10.0.7
  Local 10.10.0.2 to Remote 10.10.0.8
  Local 10.10.0.2 to Remote 10.10.0.9
  Local 10.10.0.2 to Remote 10.10.0.10
  Local 10.10.0.2 to Remote 10.10.0.11
  Local 10.10.0.2 to Remote 10.10.0.12
  Local 10.10.0.3 to Remote 10.10.0.5
  Local 10.10.0.3 to Remote 10.10.0.6
  Local 10.10.0.3 to Remote 10.10.0.7
  Local 10.10.0.3 to Remote 10.10.0.8
  Local 10.10.0.3 to Remote 10.10.0.9
  Local 10.10.0.3 to Remote 10.10.0.10
  Local 10.10.0.3 to Remote 10.10.0.11
  Local 10.10.0.3 to Remote 10.10.0.12
  Local 10.10.0.4 to Remote 10.10.0.5
  Local 10.10.0.4 to Remote 10.10.0.6
  Local 10.10.0.4 to Remote 10.10.0.7
  Local 10.10.0.4 to Remote 10.10.0.8
  Local 10.10.0.4 to Remote 10.10.0.9
```

```
Local 10.10.0.4 to Remote 10.10.0.10
Local 10.10.0.4 to Remote 10.10.0.11
Local 10.10.0.4 to Remote 10.10.0.12
```

Larger than PMTU communication succeeds on 32 path(s)

RPC status:

8 paths up, 0 paths down (tcp check)

8 paths up, 0 paths down (udp check)

9. Shut down the ports 1/31 and 1/32 on CL1, and the active Nexus 3132Q-V switch:

```
shutdown
```

Show example

This example shows the ISL ports 1/31 and 1/32 being shut down on switch CL1:

```
(CL1)# configure
(CL1) (Config)# interface e1/31-32
(CL1) (config-if-range)# shutdown
(CL1) (config-if-range)# exit
(CL1) (Config)# exit
(CL1)#
```

Step 2: Configure ports

1. Remove all the cables attached to the Nexus 3132Q-V switch CL2 and reconnect them to the replacement switch C2 on all nodes.
2. Remove the ISL cables from ports e1/31 and e1/32 on CL2 and reconnect them to the same ports on the replacement switch C2.
3. Bring up ISLs ports 1/31 and 1/32 on the Nexus 3132Q-V switch CL1:

```
(CL1)# configure
(CL1) (Config)# interface e1/31-32
(CL1) (config-if-range)# no shutdown
(CL1) (config-if-range)# exit
(CL1) (Config)# exit
(CL1)#
```

4. Verify that the ISLs are up on CL1:

```
show port-channel
```

Ports Eth1/31 and Eth1/32 should indicate (P), which means that the ISL ports are up in the port-channel.

Show example

```
CL1# show port-channel summary
Flags: D - Down          P - Up in port-channel (members)
      I - Individual    H - Hot-standby (LACP only)
      s - Suspended     r - Module-removed
      S - Switched      R - Routed
      U - Up (port-channel)
      M - Not in use. Min-links not met

-----
-----
Group Port-          Type  Protocol  Member
Ports
  Channel
-----
-----
1      Po1 (SU)      Eth    LACP      Eth1/31 (P)  Eth1/32 (P)
```

5. Verify that the ISLs are up on C2:

```
show port-channel summary
```

Ports Eth1/31 and Eth1/32 should indicate (P), which means that both ISL ports are up in the port-channel.

Show example

```
C2# show port-channel summary
Flags: D - Down          P - Up in port-channel (members)
      I - Individual    H - Hot-standby (LACP only)
      s - Suspended     r - Module-removed
      S - Switched      R - Routed
      U - Up (port-channel)
      M - Not in use. Min-links not met

-----
-----
Group Port-          Type  Protocol  Member Ports
  Channel
-----
-----
1      Po1 (SU)      Eth    LACP      Eth1/31 (P)  Eth1/32 (P)
```

6. On all nodes, bring up all the cluster interconnect ports connected to the Nexus 3132Q-V switch C2:

```
network port modify
```

Show example

```
cluster::*> network port modify -node n1 -port e0b -up-admin true
cluster::*> network port modify -node n1 -port e0c -up-admin true
cluster::*> network port modify -node n2 -port e0b -up-admin true
cluster::*> network port modify -node n2 -port e0c -up-admin true
cluster::*> network port modify -node n3 -port e4e -up-admin true
cluster::*> network port modify -node n4 -port e4e -up-admin true
```

7. For all nodes, revert all of the migrated cluster interconnect LIFs:

```
network interface revert
```

Show example

```
cluster::*> network interface revert -vserver Cluster -lif n1_clus2
cluster::*> network interface revert -vserver Cluster -lif n1_clus3
cluster::*> network interface revert -vserver Cluster -lif n2_clus2
cluster::*> network interface revert -vserver Cluster -lif n2_clus3
Cluster::*> network interface revert -vserver Cluster -lif n3_clus2
Cluster::*> network interface revert -vserver Cluster -lif n4_clus2
```

8. Verify that the cluster interconnect ports are now reverted to their home:

```
network interface show
```

Show example

This example shows that all the LIFs are successfully reverted because the ports listed under the Current Port column have a status of true in the Is Home column. If the Is Home column value is false, the LIF has not been reverted.

```
cluster::*> network interface show -role cluster
(network interface show)

```

Current Port	Is Home	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node

Cluster					
e0a	true	n1_clus1	up/up	10.10.0.1/24	n1
e0b	true	n1_clus2	up/up	10.10.0.2/24	n1
e0c	true	n1_clus3	up/up	10.10.0.3/24	n1
e0d	true	n1_clus4	up/up	10.10.0.4/24	n1
e0a	true	n2_clus1	up/up	10.10.0.5/24	n2
e0b	true	n2_clus2	up/up	10.10.0.6/24	n2
e0c	true	n2_clus3	up/up	10.10.0.7/24	n2
e0d	true	n2_clus4	up/up	10.10.0.8/24	n2
e4a	true	n3_clus1	up/up	10.10.0.9/24	n3
e4e	true	n3_clus2	up/up	10.10.0.10/24	n3
e4a	true	n4_clus1	up/up	10.10.0.11/24	n4
e4e	true	n4_clus2	up/up	10.10.0.12/24	n4

12 entries were displayed.

9. Verify that the cluster ports are connected:

```
network port show
```

Show example

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
e0a      Cluster      Cluster      up    9000  auto/10000  -
-
e0b      Cluster      Cluster      up    9000  auto/10000  -
-
e0c      Cluster      Cluster      up    9000  auto/10000  -
-
e0d      Cluster      Cluster      up    9000  auto/10000  -
-

Node: n2

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
e0a      Cluster      Cluster      up    9000  auto/10000  -
-
e0b      Cluster      Cluster      up    9000  auto/10000  -
-
e0c      Cluster      Cluster      up    9000  auto/10000  -
-
e0d      Cluster      Cluster      up    9000  auto/10000  -
-

Node: n3

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
```



```

Status
-----
e4a      Cluster      Cluster      up    9000 auto/40000 -
-
e4e      Cluster      Cluster      up    9000 auto/40000 -
-

Node: n4

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
e4a      Cluster      Cluster      up    9000 auto/40000 -
-
e4e      Cluster      Cluster      up    9000 auto/40000 -
-

12 entries were displayed.

```

10. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Node	Date		Source LIF	Destination LIF	Packet Loss
n1					
	3/5/2022 19:21:18 -06:00		n1_clus2	n2_clus1	none
	3/5/2022 19:21:20 -06:00		n1_clus2	n2_clus2	none
n2					
	3/5/2022 19:21:18 -06:00		n2_clus2	n1_clus1	none
	3/5/2022 19:21:20 -06:00		n2_clus2	n1_clus2	none
n3					
...					
...					
n4					
...					
...					

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```
cluster::*> cluster ping-cluster -node n1
Host is n1
Getting addresses from network interface table...
Cluster n1_clus1 n1 e0a 10.10.0.1
Cluster n1_clus2 n1 e0b 10.10.0.2
Cluster n2_clus1 n2 e0a 10.10.0.5
Cluster n2_clus2 n2 e0b 10.10.0.6
Cluster n2_clus3 n2 e0c 10.10.0.7
Cluster n2_clus4 n2 e0d 10.10.0.8
```

```
Cluster n3_clus1 n3      e0a 10.10.0.9
Cluster n3_clus2 n3      e0e 10.10.0.10
Cluster n4_clus1 n4      e0a 10.10.0.11
Cluster n4_clus2 n4      e0e 10.10.0.12

Local = 10.10.0.1 10.10.0.2 10.10.0.3 10.10.0.4
Remote = 10.10.0.5 10.10.0.6 10.10.0.7 10.10.0.8 10.10.0.9 10.10.0.10
10.10.0.11 10.10.0.12
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 32 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 1500 byte MTU on 32 path(s):
  Local 10.10.0.1 to Remote 10.10.0.5
  Local 10.10.0.1 to Remote 10.10.0.6
  Local 10.10.0.1 to Remote 10.10.0.7
  Local 10.10.0.1 to Remote 10.10.0.8
  Local 10.10.0.1 to Remote 10.10.0.9
  Local 10.10.0.1 to Remote 10.10.0.10
  Local 10.10.0.1 to Remote 10.10.0.11
  Local 10.10.0.1 to Remote 10.10.0.12
  Local 10.10.0.2 to Remote 10.10.0.5
  Local 10.10.0.2 to Remote 10.10.0.6
  Local 10.10.0.2 to Remote 10.10.0.7
  Local 10.10.0.2 to Remote 10.10.0.8
  Local 10.10.0.2 to Remote 10.10.0.9
  Local 10.10.0.2 to Remote 10.10.0.10
  Local 10.10.0.2 to Remote 10.10.0.11
  Local 10.10.0.2 to Remote 10.10.0.12
  Local 10.10.0.3 to Remote 10.10.0.5
  Local 10.10.0.3 to Remote 10.10.0.6
  Local 10.10.0.3 to Remote 10.10.0.7
  Local 10.10.0.3 to Remote 10.10.0.8
  Local 10.10.0.3 to Remote 10.10.0.9
  Local 10.10.0.3 to Remote 10.10.0.10
  Local 10.10.0.3 to Remote 10.10.0.11
  Local 10.10.0.3 to Remote 10.10.0.12
  Local 10.10.0.4 to Remote 10.10.0.5
  Local 10.10.0.4 to Remote 10.10.0.6
  Local 10.10.0.4 to Remote 10.10.0.7
  Local 10.10.0.4 to Remote 10.10.0.8
  Local 10.10.0.4 to Remote 10.10.0.9
  Local 10.10.0.4 to Remote 10.10.0.10
  Local 10.10.0.4 to Remote 10.10.0.11
```

```
Local 10.10.0.4 to Remote 10.10.0.12
```

```
Larger than PMTU communication succeeds on 32 path(s)
```

```
RPC status:
```

```
8 paths up, 0 paths down (tcp check)
```

```
8 paths up, 0 paths down (udp check)
```

Step 3: Verify the configuration

1. Display the information about the devices in your configuration:

- network device-discovery show
- network port show -role cluster
- network interface show -role cluster
- system cluster-switch show

Show example

```
cluster::> network device-discovery show
      Local   Discovered
Node   Port     Device           Interface         Platform
-----
n1     /cdp
      e0a    C1               Ethernet1/1/1     N3K-C3132Q-V
      e0b    C2               Ethernet1/1/1     N3K-C3132Q-V
      e0c    C2               Ethernet1/1/2     N3K-C3132Q-V
      e0d    C1               Ethernet1/1/2     N3K-C3132Q-V
n2     /cdp
      e0a    C1               Ethernet1/1/3     N3K-C3132Q-V
      e0b    C2               Ethernet1/1/3     N3K-C3132Q-V
      e0c    C2               Ethernet1/1/4     N3K-C3132Q-V
      e0d    C1               Ethernet1/1/4     N3K-C3132Q-V
n3     /cdp
      e4a    C1               Ethernet1/7       N3K-C3132Q-V
      e4e    C2               Ethernet1/7       N3K-C3132Q-V
n4     /cdp
      e4a    C1               Ethernet1/8       N3K-C3132Q-V
      e4e    C2               Ethernet1/8       N3K-C3132Q-V
12 entries were displayed.
```

```
cluster::*> network port show -role cluster
(network port show)
Node: n1

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
e0a      Cluster      Cluster      up    9000  auto/10000  -
-
e0b      Cluster      Cluster      up    9000  auto/10000  -
-
e0c      Cluster      Cluster      up    9000  auto/10000  -
-
e0d      Cluster      Cluster      up    9000  auto/10000  -
-
```

Node: n2

Ignore

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
e0a	Cluster	Cluster		up	9000	auto/10000	-
-							
e0b	Cluster	Cluster		up	9000	auto/10000	-
-							
e0c	Cluster	Cluster		up	9000	auto/10000	-
-							
e0d	Cluster	Cluster		up	9000	auto/10000	-
-							

Node: n3

Ignore

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
-----	-----						
e4a	Cluster	Cluster		up	9000	auto/40000	-
-							
e4e	Cluster	Cluster		up	9000	auto/40000	-
-							

Node: n4

Ignore

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
-----	-----						
e4a	Cluster	Cluster		up	9000	auto/40000	-
-							
e4e	Cluster	Cluster		up	9000	auto/40000	-
-							

12 entries were displayed.

```
cluster::*> network interface show -role cluster
(network interface show)
```

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
	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.

```
cluster::*> system cluster-switch show
```

```
Switch                               Type                               Address
Model
-----
CL1                                   cluster-network                   10.10.1.101
NX3132V
  Serial Number: FOX000001
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        7.0(3)I4(1)
  Version Source: CDP

CL2                                   cluster-network                   10.10.1.102
NX3132V
  Serial Number: FOX000002
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        7.0(3)I4(1)
  Version Source: CDP

C2                                   cluster-network                   10.10.1.103
NX3132V
  Serial Number: FOX000003
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        7.0(3)I4(1)
  Version Source: CDP

3 entries were displayed.
```

2. Remove the replaced Nexus 3132Q-V switch, if it is not already removed automatically:

```
system cluster-switch delete
```

```
cluster::*> system cluster-switch delete -device CL2
```


3. Verify that the proper cluster switches are monitored:

```
system cluster-switch show
```

Show example

```
cluster::> system cluster-switch show

Switch                                Type                                Address
Model
-----
CL1                                    cluster-network                    10.10.1.101
NX3132V
  Serial Number: FOX000001
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        7.0(3)I4(1)
  Version Source: CDP

C2                                    cluster-network                    10.10.1.103
NX3132V
  Serial Number: FOX000002
  Is Monitored: true
  Reason:
  Software Version: Cisco Nexus Operating System (NX-OS) Software,
Version
                        7.0(3)I4(1)
  Version Source: CDP

2 entries were displayed.
```

4. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next?

[Configure switch health monitoring](#)

Replace Cisco Nexus 3132Q-V cluster switches with switchless connections

You can migrate from a cluster with a switched cluster network to one where two nodes are directly connected for ONTAP 9.3 and later.

Review requirements

Guidelines

Review the following guidelines:

- Migrating to a two-node switchless cluster configuration is a nondisruptive operation. Most systems have two dedicated cluster interconnect ports on each node, but you can also use this procedure for systems with a larger number of dedicated cluster interconnect ports on each node, such as four, six or eight.
- You cannot use the switchless cluster interconnect feature with more than two nodes.
- If you have an existing two-node cluster that uses cluster interconnect switches and is running ONTAP 9.3 or later, you can replace the switches with direct, back-to-back connections between the nodes.


What you'll need

- A healthy cluster that consists of two nodes connected by cluster switches. The nodes must be running the same ONTAP release.
- Each node with the required number of dedicated cluster ports, which provide redundant cluster interconnect connections to support your system configuration. For example, there are two redundant ports for a system with two dedicated cluster interconnect ports on each node.

Migrate the switches

About this task

The following procedure removes the cluster switches in a two-node cluster and replaces each connection to the switch with a direct connection to the partner node.



About the examples

The examples in the following procedure show nodes that are using "e0a" and "e0b" as cluster ports. Your nodes might be using different cluster ports as they vary by system.

Step 1: Prepare for migration

1. Change the privilege level to advanced, entering `y` when prompted to continue:

```
set -privilege advanced
```

The advanced prompt `*>` appears.

2. ONTAP 9.3 and later supports automatic detection of switchless clusters, which is enabled by default.

You can verify that detection of switchless clusters is enabled by running the advanced privilege command:

```
network options detect-switchless-cluster show
```

Show example

The following example output shows if the option is enabled.

```
cluster::*> network options detect-switchless-cluster show
(network options detect-switchless-cluster show)
Enable Switchless Cluster Detection: true
```

If "Enable Switchless Cluster Detection" is `false`, contact NetApp support.

3. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message
MAINT=<number_of_hours>h
```

where `h` is the duration of the maintenance window in hours. The message notifies technical support of this maintenance task so that they can suppress automatic case creation during the maintenance window.

In the following example, the command suppresses automatic case creation for two hours:

Show example


```
cluster::*> system node autosupport invoke -node * -type all
-message MAINT=2h
```

Step 2: Configure ports and cabling

1. Organize the cluster ports on each switch into groups so that the cluster ports in group1 go to cluster switch1 and the cluster ports in group2 go to cluster switch2. These groups are required later in the procedure.
2. Identify the cluster ports and verify link status and health:

```
network port show -ipspace Cluster
```

In the following example for nodes with cluster ports "e0a" and "e0b", one group is identified as "node1:e0a" and "node2:e0a" and the other group as "node1:e0b" and "node2:e0b". Your nodes might be using different cluster ports because they vary by system.



Verify that the ports have a value of `up` for the “Link” column and a value of `healthy` for the “Health Status” column.

Show example

```

cluster::> network port show -ipspace Cluster
Node: node1

Ignore

Health
Speed (Mbps) Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status
Status
-----
-----
e0a Cluster Cluster up 9000 auto/10000 healthy
false
e0b Cluster Cluster up 9000 auto/10000 healthy
false

Node: node2

Ignore

Health
Speed (Mbps) Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status
Status
-----
-----
e0a Cluster Cluster up 9000 auto/10000 healthy
false
e0b Cluster Cluster up 9000 auto/10000 healthy
false
4 entries were displayed.
  
```

3. Confirm that all the cluster LIFs are on their home ports.

Verify that the “is-home” column is true for each of the cluster LIFs:

```
network interface show -vserver Cluster -fields is-home
```

Show example

```
cluster::*> net int show -vserver Cluster -fields is-home
(network interface show)
vserver  lif          is-home
-----  -
Cluster  node1_clus1  true
Cluster  node1_clus2  true
Cluster  node2_clus1  true
Cluster  node2_clus2  true
4 entries were displayed.
```

If there are cluster LIFs that are not on their home ports, revert those LIFs to their home ports:

```
network interface revert -vserver Cluster -lif *
```

4. Disable auto-revert for the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

5. Verify that all ports listed in the previous step are connected to a network switch:

```
network device-discovery show -port cluster_port
```

The “Discovered Device” column should be the name of the cluster switch that the port is connected to.

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to cluster switches "cs1" and "cs2".

```
cluster::> network device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol  Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1/cdp
          e0a    cs1                        0/11       BES-53248
          e0b    cs2                        0/12       BES-53248
node2/cdp
          e0a    cs1                        0/9        BES-53248
          e0b    cs2                        0/9        BES-53248
4 entries were displayed.
```

6. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node2		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node1		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node2		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

7. Verify that the cluster is healthy:

```
cluster ring show
```

All units must be either master or secondary.


8. Set up the switchless configuration for the ports in group 1.



To avoid potential networking issues, you must disconnect the ports from group1 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.


a. Disconnect all the cables from the ports in group1 at the same time.

In the following example, the cables are disconnected from port "e0a" on each node, and cluster traffic continues through the switch and port "e0b" on each node:



b. Cable the ports in group1 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2:



9. The switchless cluster network option transitions from *false* to *true*. This might take up to 45 seconds. Confirm that the switchless option is set to *true*:

```
network options switchless-cluster show
```

The following example shows that the switchless cluster is enabled:

```
cluster::*> network options switchless-cluster show
Enable Switchless Cluster: true
```

10. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```



Before proceeding to the next step, you must wait at least two minutes to confirm a working back-to-back connection on group 1.


11. Set up the switchless configuration for the ports in group 2.



To avoid potential networking issues, you must disconnect the ports from group2 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.


- a. Disconnect all the cables from the ports in group2 at the same time.

In the following example, the cables are disconnected from port "e0b" on each node, and cluster traffic continues through the direct connection between the "e0a" ports:



b. Cable the ports in group2 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2 and "e0b" on node1 is connected to "e0b" on node2:



Step 3: Verify the configuration

1. Verify that the ports on both nodes are correctly connected:

```
network device-discovery show -port cluster_port
```

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to the corresponding port on the cluster partner:

```
cluster::> net device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1/cdp
          e0a    node2                      e0a        AFF-A300
          e0b    node2                      e0b        AFF-A300
node1/lldp
          e0a    node2 (00:a0:98:da:16:44)  e0a        -
          e0b    node2 (00:a0:98:da:16:44)  e0b        -
node2/cdp
          e0a    node1                      e0a        AFF-A300
          e0b    node1                      e0b        AFF-A300
node2/lldp
          e0a    node1 (00:a0:98:da:87:49)  e0a        -
          e0b    node1 (00:a0:98:da:87:49)  e0b        -
8 entries were displayed.
```

2. Re-enable auto-revert for the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert true
```

3. Verify that all LIFs are home. This might take a few seconds.

```
network interface show -vserver Cluster -lif lif_name
```

Show example

The LIFs have been reverted if the “Is Home” column is `true`, as shown for `node1_clus2` and `node2_clus2` in the following example:

```
cluster::> network interface show -vserver Cluster -fields curr-
port,is-home
vserver  lif                curr-port  is-home
-----  -
Cluster  node1_clus1            e0a        true
Cluster  node1_clus2            e0b        true
Cluster  node2_clus1            e0a        true
Cluster  node2_clus2            e0b        true
4 entries were displayed.
```

If any cluster LIFS have not returned to their home ports, revert them manually from the local node:

```
network interface revert -vserver Cluster -lif lif_name
```

4. Check the cluster status of the nodes from the system console of either node:

```
cluster show
```

Show example

The following example shows `epsilon` on both nodes to be `false`:

```
Node  Health  Eligibility  Epsilon
-----  -
node1 true    true        false
node2 true    true        false
2 entries were displayed.
```

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

6. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

For more information, see [NetApp KB Article 1010449: How to suppress automatic case creation during scheduled maintenance windows](#).

7. Change the privilege level back to admin:

```
set -privilege admin
```

Cisco Nexus 92300YC

Overview

Overview of installation and configuration for Cisco Nexus 92300YC switches

Before configuring Cisco Nexus 92300YC switches, review the procedure overview.

To initially configure a Cisco Nexus 92300YC switch on systems running ONTAP, follow these steps:

1. [Complete Cisco Nexus 92300YC cabling worksheet](#). The sample cabling worksheet provides examples of recommended port assignments from the switches to the controllers. The blank worksheet provides a template that you can use in setting up your cluster.
2. [Configure the Cisco Nexus 92300YC switch](#). Set up and configure the Cisco Nexus 92300YC switch.
3. [Prepare to install NX-OS software and Reference Configuration File \(RCF\)](#). Prepare for installing the NX-OS software and the Reference Configuration File (RCF).
4. [Install the NX-OS software](#). Install the NX-OS software on the Nexus 92300YC switch. NX-OS is a network operating system for the Nexus series of Ethernet switches and MDS series of Fibre Channel (FC) storage area network switches provided by Cisco Systems.
5. [Install the Reference Configuration File \(RCF\)](#). Install the RCF after setting up the Nexus 92300YC switch for the first time. You can also use this procedure to upgrade your RCF version.

Additional information

Before you begin installation or maintenance, be sure to review the following:

- [Configuration requirements](#)
- [Components and part numbers](#)
- [Required documentation](#)
- [Smart Call Home requirements](#)

Configuration requirements for Cisco Nexus 92300YC switches

For Cisco Nexus 92300YC switch installation and maintenance, be sure to review all configuration and network requirements.

If you want to build ONTAP clusters with more than two nodes, you need two supported cluster network switches. You can use additional management switches, which are optional.

Configuration requirements

To configure your cluster, you need the appropriate number and type of cables and cable connectors for your switches. Depending on the type of switch you are initially configuring, you need to connect to the switch console port with the included console cable; you also need to provide specific network information.

Network requirements

You need the following network information for all switch configurations:

- IP subnet for management network traffic
- Host names and IP addresses for each of the storage system controllers and all applicable switches
- Most storage system controllers are managed through the e0M interface by connecting to the Ethernet service port (wrench icon). On AFF A800 and AFF A700 systems, the e0M interface uses a dedicated Ethernet port.

Refer to the [Hardware Universe](#) for latest information.

Components for Cisco Nexus 92300YC switches

For Cisco Nexus 92300YC switch installation and maintenance, be sure to review all

switch components and part numbers. See the [Hardware Universe](#) for details.

The following table lists the part number and description for the 92300YC switch, fans, and power supplies:

Part number	Description
190003	Cisco 92300YC, CLSW, 48Pt10/25GB, 18Pt100G, PTSX (PTSX = Port Side Exhaust)
190003R	Cisco 92300YC, CLSW, 48Pt10/25GB, 18Pt100G, PSIN (PSIN = Port Side Intake)
X-NXA-FAN-35CFM-B	Fan, Cisco N9K port side intake airflow
X-NXA-FAN-35CFM-F	Fan, Cisco N9K port side exhaust airflow
X-NXA-PAC-650W-B	Power supply, Cisco 650W - port side intake
X-NXA-PAC-650W-F	Power supply, Cisco 650W - port side exhaust

Cisco Nexus 92300YC switch airflow details:

- Port-side exhaust airflow (standard air) — Cool air enters the chassis through the fan and power supply modules in the cold aisle and exhausts through the port end of the chassis in the hot aisle. Port-side exhaust airflow with blue coloring.
- Port-side intake airflow (reverse air) — Cool air enters the chassis through the port end in the cold aisle and exhausts through the fan and power supply modules in the hot aisle. Port-side intake airflow with burgundy coloring.

Documentation requirements for Cisco Nexus 92300YC switches

For Cisco Nexus 92300YC switch installation and maintenance, be sure to review all the recommended documentation.

Switch documentation

To set up the Cisco Nexus 92300YC switches, you need the following documentation from the [Cisco Nexus 9000 Series Switches Support](#) page:

Document title	Description
<i>Nexus 9000 Series Hardware Installation Guide</i>	Provides detailed information about site requirements, switch hardware details, and installation options.
<i>Cisco Nexus 9000 Series Switch Software Configuration Guides</i> (choose the guide for the NX-OS release installed on your switches)	Provides initial switch configuration information that you need before you can configure the switch for ONTAP operation.

Document title	Description
<i>Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide</i> (choose the guide for the NX-OS release installed on your switches)	Provides information on how to downgrade the switch to ONTAP supported switch software, if necessary.
<i>Cisco Nexus 9000 Series NX-OS Command Reference Master Index</i>	Provides links to the various command references provided by Cisco.
<i>Cisco Nexus 9000 MIBs Reference</i>	Describes the Management Information Base (MIB) files for the Nexus 9000 switches.
<i>Nexus 9000 Series NX-OS System Message Reference</i>	Describes the system messages for Cisco Nexus 9000 series switches, those that are informational, and others that might help diagnose problems with links, internal hardware, or the system software.
<i>Cisco Nexus 9000 Series NX-OS Release Notes</i> (choose the notes for the NX-OS release installed on your switches)	Describes the features, bugs, and limitations for the Cisco Nexus 9000 Series.
Regulatory Compliance and Safety Information for Cisco Nexus 9000 Series	Provides international agency compliance, safety, and statutory information for the Nexus 9000 series switches.

ONTAP systems documentation

To set up an ONTAP system, you need the following documents for your version of the operating system from the [ONTAP 9 Documentation Center](#).

Name	Description
Controller-specific <i>Installation and Setup Instructions</i>	Describes how to install NetApp hardware.
ONTAP documentation	Provides detailed information about all aspects of the ONTAP releases.
Hardware Universe	Provides NetApp hardware configuration and compatibility information.

Rail kit and cabinet documentation

To install a Cisco Nexus 92300YC switch in a NetApp cabinet, see the following hardware documentation.

Name	Description
42U System Cabinet, Deep Guide	Describes the FRUs associated with the 42U system cabinet, and provides maintenance and FRU replacement instructions.

Name	Description
[Install a Cisco Nexus 92300YC switch in a NetApp Cabinet]	Describes how to install a Cisco Nexus 92300YC switch in a four-post NetApp cabinet.

Smart Call Home requirements

To use Smart Call Home feature, review the following guidelines.

Smart Call Home monitors the hardware and software components on your network. When a critical system configuration occurs, it generates an email-based notification and raises an alert to all the recipients that are configured in your destination profile. To use Smart Call Home, you must configure a cluster network switch to communicate using email with the Smart Call Home system. In addition, you can optionally set up your cluster network switch to take advantage of Cisco's embedded Smart Call Home support feature.

Before you can use Smart Call Home, be aware of the following considerations:

- An email server must be in place.
- The switch must have IP connectivity to the email server.
- The contact name (SNMP server contact), phone number, and street address information must be configured. This is required to determine the origin of messages received.
- A CCO ID must be associated with an appropriate Cisco SMARTnet Service contract for your company.
- Cisco SMARTnet Service must be in place for the device to be registered.

The [Cisco support site](#) contains information about the commands to configure Smart Call Home.

Install hardware

Complete Cisco Nexus 92300YC cabling worksheet

If you want to document the supported platforms, download a PDF of this page and complete the cabling worksheet.

The sample cabling worksheet provides examples of recommended port assignments from the switches to the controllers. The blank worksheet provides a template that you can use in setting up your cluster.

Sample cabling worksheet

The sample port definition on each pair of switches is as follows:

Cluster switch A		Cluster switch B	
Switch port	Node and port usage	Switch port	Node and port usage
1	10/25 GbE node	1	10/25 GbE node
2	10/25 GbE node	2	10/25 GbE node
3	10/25 GbE node	3	10/25 GbE node

Cluster switch A		Cluster switch B	
4	10/25 GbE node	4	10/25 GbE node
5	10/25 GbE node	5	10/25 GbE node
6	10/25 GbE node	6	10/25 GbE node
7	10/25 GbE node	7	10/25 GbE node
8	10/25 GbE node	8	10/25 GbE node
9	10/25 GbE node	9	10/25 GbE node
10	10/25 GbE node	10	10/25 GbE node
11	10/25 GbE node	11	10/25 GbE node
12	10/25 GbE node	12	10/25 GbE node
13	10/25 GbE node	13	10/25 GbE node
14	10/25 GbE node	14	10/25 GbE node
15	10/25 GbE node	15	10/25 GbE node
16	10/25 GbE node	16	10/25 GbE node
17	10/25 GbE node	17	10/25 GbE node
18	10/25 GbE node	18	10/25 GbE node
19	10/25 GbE node	19	10/25 GbE node
20	10/25 GbE node	20	10/25 GbE node
21	10/25 GbE node	21	10/25 GbE node
22	10/25 GbE node	22	10/25 GbE node
23	10/25 GbE node	23	10/25 GbE node
24	10/25 GbE node	24	10/25 GbE node
25	10/25 GbE node	25	10/25 GbE node

Cluster switch A		Cluster switch B	
26	10/25 GbE node	26	10/25 GbE node
27	10/25 GbE node	27	10/25 GbE node
28	10/25 GbE node	28	10/25 GbE node
29	10/25 GbE node	29	10/25 GbE node
30	10/25 GbE node	30	10/25 GbE node
31	10/25 GbE node	31	10/25 GbE node
32	10/25 GbE node	32	10/25 GbE node
33	10/25 GbE node	33	10/25 GbE node
34	10/25 GbE node	34	10/25 GbE node
35	10/25 GbE node	35	10/25 GbE node
36	10/25 GbE node	36	10/25 GbE node
37	10/25 GbE node	37	10/25 GbE node
38	10/25 GbE node	38	10/25 GbE node
39	10/25 GbE node	39	10/25 GbE node
40	10/25 GbE node	40	10/25 GbE node
41	10/25 GbE node	41	10/25 GbE node
42	10/25 GbE node	42	10/25 GbE node
43	10/25 GbE node	43	10/25 GbE node
44	10/25 GbE node	44	10/25 GbE node
45	10/25 GbE node	45	10/25 GbE node
46	10/25 GbE node	46	10/25 GbE node
47	10/25 GbE node	47	10/25 GbE node

Cluster switch A		Cluster switch B	
48	10/25 GbE node	48	10/25 GbE node
49	40/100 GbE node	49	40/100 GbE node
50	40/100 GbE node	50	40/100 GbE node
51	40/100 GbE node	51	40/100 GbE node
52	40/100 GbE node	52	40/100 GbE node
53	40/100 GbE node	53	40/100 GbE node
54	40/100 GbE node	54	40/100 GbE node
55	40/100 GbE node	55	40/100 GbE node
56	40/100 GbE node	56	40/100 GbE node
57	40/100 GbE node	57	40/100 GbE node
58	40/100 GbE node	58	40/100 GbE node
59	40/100 GbE node	59	40/100 GbE node
60	40/100 GbE node	60	40/100 GbE node
61	40/100 GbE node	61	40/100 GbE node
62	40/100 GbE node	62	40/100 GbE node
63	40/100 GbE node	63	40/100 GbE node
64	40/100 GbE node	64	40/100 GbE node
65	100 GbE ISL to switch B port 65	65	100 GbE ISL to switch A port 65
66	100 GbE ISL to switch B port 66	66	100 GbE ISL to switch A port 65

Blank cabling worksheet

You can use the blank cabling worksheet to document the platforms that are supported as nodes in a cluster. The *Supported Cluster Connections* section of the [Hardware Universe](#) defines the cluster ports used by the

platform.

Cluster switch A		Cluster switch B	
Switch port	Node/port usage	Switch port	Node/port usage
1		1	
2		2	
3		3	
4		4	
5		5	
6		6	
7		7	
8		8	
9		9	
10		10	
11		11	
12		12	
13		13	
14		14	
15		15	
16		16	
17		17	
18		18	
19		19	
20		20	

Cluster switch A		Cluster switch B	
21		21	
22		22	
23		23	
24		24	
25		25	
26		26	
27		27	
28		28	
29		29	
30		30	
31		31	
32		32	
33		33	
34		34	
35		35	
36		36	
37		37	
38		38	
39		39	
40		40	
41		41	
42		42	

Cluster switch A		Cluster switch B	
43		43	
44		44	
45		45	
46		46	
47		47	
48		48	
49		49	
50		50	
51		51	
52		52	
53		53	
54		54	
55		55	
56		56	
57		57	
58		58	
59		59	
60		60	
61		61	
62		62	
63		63	
64		64	

Cluster switch A		Cluster switch B	
65	ISL to switch B port 65	65	ISL to switch A port 65
66	ISL to switch B port 66	66	ISL to switch A port 66

Configure the Cisco Nexus 92300YC switch

Follow this procedure to set up and configure the Cisco Nexus 92300YC switch.

Steps

1. Connect the serial port to a host or serial port.
2. Connect the management port (on the non-port side of the switch) to the same network where your SFTP server is located.
3. At the console, set the host side serial settings:
 - 9600 baud
 - 8 data bits
 - 1 stop bit
 - parity: none
 - flow control: none
4. When booting for the first time or rebooting after erasing the running configuration, the Nexus 92300YC switch loops in a boot cycle. Interrupt this cycle by typing **yes** to abort Power on Auto Provisioning.

The System Admin Account setup is displayed.

Show example

```
$ VDC-1 %$ %POAP-2-POAP_INFO:   - Abort Power On Auto Provisioning
[yes - continue with normal setup, skip - bypass password and basic
configuration, no - continue with Power On Auto Provisioning]
(yes/skip/no) [no]: y
Disabling POAP.....Disabling POAP
2019 Apr 10 00:36:17 switch %$ VDC-1 %$ poap: Rolling back, please
wait... (This may take 5-15 minutes)

      ---- System Admin Account Setup ----

Do you want to enforce secure password standard (yes/no) [y]:
```

5. Type **y** to enforce secure password standard:

```
Do you want to enforce secure password standard (yes/no) [y]: y
```

6. Enter and confirm the password for user admin:

```
Enter the password for "admin":  
Confirm the password for "admin":
```

7. Type **yes** to enter the Basic System Configuration dialog.

Show example

```
This setup utility will guide you through the basic configuration of  
the system. Setup configures only enough connectivity for management  
of the system.
```

```
Please register Cisco Nexus9000 Family devices promptly with your  
supplier. Failure to register may affect response times for initial  
service calls. Nexus9000 devices must be registered to receive  
entitled support services.
```

```
Press Enter at anytime to skip a dialog. Use ctrl-c at anytime  
to skip the remaining dialogs.
```

```
Would you like to enter the basic configuration dialog (yes/no):
```

8. Create another login account:

```
Create another login account (yes/no) [n]:
```

9. Configure read-only and read-write SNMP community strings:

```
Configure read-only SNMP community string (yes/no) [n]:
```

```
Configure read-write SNMP community string (yes/no) [n]:
```

10. Configure the cluster switch name:

```
Enter the switch name : cs2
```

11. Configure the out-of-band management interface:

```
Continue with Out-of-band (mgmt0) management configuration? (yes/no)
[y]: y

Mgmt0 IPv4 address : 172.22.133.216

Mgmt0 IPv4 netmask : 255.255.224.0

Configure the default gateway? (yes/no) [y]: y

IPv4 address of the default gateway : 172.22.128.1
```

12. Configure advanced IP options:

```
Configure advanced IP options? (yes/no) [n]: n
```

13. Configure Telnet services:

```
Enable the telnet service? (yes/no) [n]: n
```

14. Configure SSH services and SSH keys:

```
Enable the ssh service? (yes/no) [y]: y

Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa

Number of rsa key bits <1024-2048> [1024]: 2048
```

15. Configure other settings:

```
Configure the ntp server? (yes/no) [n]: n

Configure default interface layer (L3/L2) [L2]: L2

Configure default switchport interface state (shut/noshut) [noshut]:
noshut

Configure CoPP system profile (strict/moderate/lenient/dense)
[strict]: strict
```

16. Confirm switch information and save the configuration:

```
Would you like to edit the configuration? (yes/no) [n]: n

Use this configuration and save it? (yes/no) [y]: y

[] 100%
Copy complete, now saving to disk (please wait)...
Copy complete.
```

What's next?

[Prepare to install NX-OS software and RCF.](#)

Review cabling and configuration considerations

Before configuring your Cisco 92300YC switch, review the following considerations.

Support for NVIDIA CX6, CX6-DX, and CX7 Ethernet ports

If connecting a switch port to an ONTAP controller using NVIDIA ConnectX-6 (CX6), ConnectX-6 Dx (CX6-DX), or ConnectX-7 (CX7) NIC ports, you must hard-code the switch port speed.

```
(cs1)(config)# interface Ethernet1/19
For 100GbE speed:
(cs1)(config-if)# speed 100000
For 40GbE speed:
(cs1)(config-if)# speed 40000
(cs1)(config-if)# no negotiate auto
(cs1)(config-if)# exit
(cs1)(config)# exit
Save the changes:
(cs1)# copy running-config startup-config
```

See the [Hardware Universe](#) for more information on switch ports.

Configure software

Install the NX-OS software

Follow this procedure to install the NX-OS software on the Nexus 92300YC switch.

NX-OS is a network operating system for the Nexus series of Ethernet switches and MDS series of Fibre Channel (FC) storage area network switches provided by Cisco Systems.

Review requirements

Supported ports and node connections

- The Inter-Switch Links (ISLs) supported for the Nexus 92300YC switches are ports 1/65 and 1/66.

- The node connections supported for the Nexus 92300YC switches are ports 1/1 through 1/66.

What you'll need

- Applicable NetApp Cisco NX-OS software for your switches from the NetApp Support Site, available from mysupport.netapp.com
- A fully functioning cluster (no errors in the logs or similar issues).
- [Cisco Ethernet switch page](#). Consult the switch compatibility table for the supported ONTAP and NX-OS versions.

Install the software

The examples in this procedure use two nodes, but you can have up to 24 nodes in a cluster.

About the examples

The examples in this procedure use the following switch and node nomenclature:

- The Nexus 92300YC switch names are `cs1` and `cs2`.
- The example used in this procedure starts the upgrade on the second switch, `*cs2*`.
- The cluster LIF names are `node1_clus1` and `node1_clus2` for node1, and `node2_clus1` and `node2_clus2` for node2.
- The IPspace name is `Cluster`.
- The `cluster1::*>` prompt indicates the name of the cluster.
- The cluster ports on each node are named `e0a` and `e0b`.

See the [Hardware Universe^](#) for the actual cluster ports supported on your platform.

Steps

1. Connect the cluster switch to the management network.
2. Use the `ping` command to verify connectivity to the server hosting the NX-OS software and the RCF.

Show example

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.
```

3. Copy the NX-OS software and EPLD images to the Nexus 92300YC switch.

Show example

```
cs2# copy sftp: bootflash: vrf management
Enter source filename: /code/nxos.9.2.2.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.9.2.2.bin /bootflash/nxos.9.2.2.bin
/code/nxos.9.2.2.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.9.2.2.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.9.2.2.img /bootflash/n9000-
epld.9.2.2.img
/code/n9000-epld.9.2.2.img 100% 161MB 9.5MB/s 00:16
sftp> exit
Copy complete, now saving to disk (please wait)...
Copy complete.
```

4. Verify the running version of the NX-OS software:

```
show version
```


Show example

```
cs2# show version
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2018, 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 05.31
  NXOS: version 9.2(1)
  BIOS compile time: 05/17/2018
  NXOS image file is: bootflash:///nxos.9.2.1.bin
  NXOS compile time: 7/17/2018 16:00:00 [07/18/2018 00:21:19]

Hardware
  cisco Nexus9000 C92300YC Chassis
  Intel(R) Xeon(R) CPU D-1526 @ 1.80GHz with 16337884 kB of memory.
  Processor Board ID FDO220329V5

  Device name: cs2
  bootflash: 115805356 kB
  Kernel uptime is 0 day(s), 4 hour(s), 23 minute(s), 11 second(s)

  Last reset at 271444 usecs after Wed Apr 10 00:25:32 2019
  Reason: Reset Requested by CLI command reload
```

```
System version: 9.2(1)
```

```
Service:
```

```
plugin
```

```
Core Plugin, Ethernet Plugin
```

```
Active Package(s):
```

```
cs2#
```

5. Install the NX-OS image.

Installing the image file causes it to be loaded every time the switch is rebooted.

Show example

```
cs2# install all nxos bootflash:nxos.9.2.2.bin
```

```
Installer will perform compatibility check first. Please wait.  
Installer is forced disruptive
```

```
Verifying image bootflash:/nxos.9.2.2.bin for boot variable "nxos".  
[] 100% -- SUCCESS
```

```
Verifying image type.  
[] 100% -- SUCCESS
```

```
Preparing "nxos" version info using image bootflash:/nxos.9.2.2.bin.  
[] 100% -- SUCCESS
```

```
Preparing "bios" version info using image bootflash:/nxos.9.2.2.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 Version	Image	Running-Version(pri:alt)	New-
		Upg-Required	
1	nxos		9.2(1)
9.2(2)		yes	
1	bios	v05.31(05/17/2018):v05.28(01/18/2018)	
v05.33(09/08/2018)		yes	

```
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
```

```
2019 Apr 10 04:59:35 cs2 %$ VDC-1 %$ %VMAN-2-ACTIVATION_STATE:  
Successfully deactivated virtual service 'guestshell+'
```

```
Finishing the upgrade, switch will reboot in 10 seconds.
```

6. Verify the new version of NX-OS software after the switch has rebooted:

```
show version
```

Show example

```
cs2# show version
```

```
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2018, 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 05.33
NXOS: version 9.2(2)
BIOS compile time: 09/08/2018
NXOS image file is: bootflash:///nxos.9.2.2.bin
NXOS compile time: 11/4/2018 21:00:00 [11/05/2018 06:11:06]
```

Hardware

```
cisco Nexus9000 C92300YC Chassis
Intel(R) Xeon(R) CPU D-1526 @ 1.80GHz with 16337884 kB of memory.
Processor Board ID FDO220329V5

Device name: cs2
bootflash: 115805356 kB
Kernel uptime is 0 day(s), 0 hour(s), 3 minute(s), 52 second(s)
```

```
Last reset at 182004 usecs after Wed Apr 10 04:59:48 2019
```

Reason: Reset due to upgrade

System version: 9.2(1)

Service:

plugin

Core Plugin, Ethernet Plugin

Active Package(s):

7. Upgrade the EPLD image and reboot the switch.

Show example

```
cs2# show version module 1 epld
```

EPLD Device	Version
MI FPGA	0x7
IO FPGA	0x17
MI FPGA2	0x2
GEM FPGA	0x2
GEM FPGA	0x2
GEM FPGA	0x2
GEM FPGA	0x2

```
cs2# install epld bootflash:n9000-epld.9.2.2.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	0x07	0x07	No
1	SUP	IO FPGA	0x17	0x19	Yes
1	SUP	MI FPGA2	0x02	0x02	No

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
```

```
EPLDs upgraded.
```

```
Module 1 EPLD upgrade is successful.
```

8. After the switch reboot, log in again and verify that the new version of EPLD loaded successfully.

Show example

```
cs2# *show version module 1 epld*
```

EPLD Device	Version
MI FPGA	0x7
IO FPGA	0x19
MI FPGA2	0x2
GEM FPGA	0x2
GEM FPGA	0x2
GEM FPGA	0x2
GEM FPGA	0x2

What's next?

[Install the Reference Configuration File](#)

Verify your SSH configuration

If you are using the Ethernet Switch Health Monitor (CSHM) and log collection features, verify that SSH and SSH keys are enabled on the cluster switches.

Steps

1. Verify that SSH is enabled:

```
(switch) show ssh server  
ssh version 2 is enabled
```

2. Verify that the SSH keys are enabled:

```
show ssh key
```


Show example

```
(switch)# show ssh key

rsa Keys generated:Fri Jun 28 02:16:00 2024

ssh-rsa
AAAAB3NzaC1yc2EAAAADAQABAAQGDINrD52Q586wTGJjFABjBlFaA23EpDrZ2sDCew
l7nwlIoC6HBejxluIObAH8hrW8kR+gj0ZAfPpNeLGTg3APj/yIPTBoIZZxbWRShywAM5
PqyxWwRb7kp9Zt1YHzVuHYpSO82KUDowKrL6lox/YtpKoZUDZjrZjAp8hTv3JZsPgQ==

bitcount:1024
fingerprint:
SHA256:aHwhpzo7+YCDSrp3isJv2uVGz+mjMMokqdMeXVVXfdo

could not retrieve dsa key information

ecdsa Keys generated:Fri Jun 28 02:30:56 2024

ecdsa-sha2-nistp521
AAAAE2VjZHNhLXNoYTItbmlzdHA1MjEAAAABmlzdHA1MjEAAACFBABJ+ZX5SFKhS57e
vkE273e0VoqZi4/32dt+f14fBuKv80MjMsmLfjKtCWylwgVt1Zi+C5TIBbugpzez529z
kFSF0ADb8JaGCoaAYe2HvWR/f6QLbKbqVIewCdqWgxzrIY5BPP5GBdxQJMBiOwEdnHg1
u/9Pzh/Vz9cHDcCW9qGE780QHA==

bitcount:521
fingerprint:
SHA256:TFGe2hXn6QIpcs/vyHzftHJ7Dceg0vQaULYRALZeHwQ

(switch)# show feature | include scpServer
scpServer          1          enabled
(switch)# show feature | include ssh
sshServer          1          enabled
(switch)#
```



When enabling FIPS, you must change the bitcount to 256 on the switch using the command `ssh key ecdsa 256 force`. See [Configure network security using FIPS](#) for more details.

What's next?

[Configure switch health monitoring.](#)

Migrate switches

Migrate to a two-node switched cluster with a Cisco Nexus 92300YC switch

If you have an existing two-node *switchless* cluster environment, you can migrate to a two-node *switched* cluster environment using Cisco Nexus 92300YC switches to enable you to scale beyond two nodes in the cluster.

The procedure you use depends on whether you have two dedicated cluster-network ports on each controller or a single cluster port on each controller. The process documented works for all nodes using optical or twinax ports, but is not supported on this switch if nodes are using onboard 10Gb BASE-T RJ45 ports for the cluster-network ports.

Most systems require two dedicated cluster-network ports on each controller.



After your migration completes, you might need to install the required configuration file to support the Cluster Switch Health Monitor (CSHM) for 92300YC cluster switches. See [Install the Cluster Switch Health Monitor \(CSHM\)](#).

Review requirements

What you'll need

For a two-node switchless configuration, ensure that:

- The two-node switchless configuration is properly set up and functioning.
- The nodes are running ONTAP 9.6 and later.
- All cluster ports are in the **up** state.
- All cluster logical interfaces (LIFs) are in the **up** state and on their home ports.

For the Cisco Nexus 92300YC switch configuration:

- Both switches have management network connectivity.
- There is console access to the cluster switches.
- Nexus 92300YC node-to-node switch and switch-to-switch connections use twinax or fiber cables.

[Hardware Universe - Switches](#) contains more information about cabling.

- Inter-Switch Link (ISL) cables are connected to ports 1/65 and 1/66 on both 92300YC switches.
- Initial customization of both the 92300YC switches are completed. So that the:
 - 92300YC switches are running the latest version of software
 - Reference Configuration Files (RCFs) are applied to the switchesAny site customization, such as SMTP, SNMP, and SSH is configured on the new switches.

Migrate the switch

About the examples

The examples in this procedure use the following cluster switch and node nomenclature:

- The names of the 92300YC switches are cs1 and cs2.
- The names of the cluster SVMs are node1 and node2.
- The names of the LIFs are node1_clus1 and node1_clus2 on node 1, and node2_clus1 and node2_clus2

on node 2 respectively.

- The `cluster1::*>` prompt indicates the name of the cluster.
- The cluster ports used in this procedure are e0a and e0b.

[Hardware Universe](#) contains the latest information about the actual cluster ports for your platforms.

Step 1: Prepare for migration

1. Change the privilege level to advanced, entering `y` when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (`*>`) appears.

2. 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
```

where `x` is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

Show example

The following command suppresses automatic case creation for two hours:

```
cluster1::*> system node autosupport invoke -node * -type all  
-message MAINT=2h
```

Step 2: Configure cables and ports

1. Disable all node-facing ports (not ISL ports) on both the new cluster switches `cs1` and `cs2`.

You must not disable the ISL ports.

Show example

The following example shows that node-facing ports 1 through 64 are disabled on switch `cs1`:

```
cs1# config  
Enter configuration commands, one per line. End with CNTL/Z.  
cs1(config)# interface e/1-64  
cs1(config-if-range)# shutdown
```

2. Verify that the ISL and the physical ports on the ISL between the two 92300YC switches cs1 and cs2 are up on ports 1/65 and 1/66:

```
show port-channel summary
```

Show example

The following example shows that the ISL ports are up on switch cs1:

```
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      Eth1/65 (P)  Eth1/66 (P)
```

+

The following example shows that the ISL ports are up on switch cs2 :

+

```
(cs2)# 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      Eth1/65 (P)  Eth1/66 (P)
```

3. Display the list of neighboring devices:

```
show cdp neighbors
```

This command provides information about the devices that are connected to the system.

Show example

The following example lists the neighboring devices on switch cs1:

```
cs1# 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
cs2 (FDO220329V5)  Eth1/65       175     R S I s     N9K-C92300YC
Eth1/65
cs2 (FDO220329V5)  Eth1/66       175     R S I s     N9K-C92300YC
Eth1/66

Total entries displayed: 2
```

+

The following example lists the neighboring devices on switch cs2:

+

```
cs2# 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
cs1 (FDO220329KU)  Eth1/65       177     R S I s     N9K-C92300YC
Eth1/65
cs1 (FDO220329KU)  Eth1/66       177     R S I s     N9K-C92300YC
Eth1/66

Total entries displayed: 2
```

4. Verify that all cluster ports are up:

```
network port show -ipspace Cluster
```

Each port should display up for Link and healthy for Health Status.

Show example

```
cluster1::*> network port show -ipspace Cluster
```

```
Node: node1
```

Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Oper	Health Status
e0a	Cluster	Cluster	up	9000	auto/10000	healthy
e0b	Cluster	Cluster	up	9000	auto/10000	healthy

```
Node: node2
```

Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Oper	Health Status
e0a	Cluster	Cluster	up	9000	auto/10000	healthy
e0b	Cluster	Cluster	up	9000	auto/10000	healthy

```
4 entries were displayed.
```

5. Verify that all cluster LIFs are up and operational:

```
network interface show -vserver Cluster
```

Each cluster LIF should display true for Is Home and have a Status Admin/Oper of up/up

Show example

```
cluster1::*> network interface show -vserver Cluster
```

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
e0a	node1_clus1	up/up	169.254.209.69/16	node1
	true			
e0b	node1_clus2	up/up	169.254.49.125/16	node1
	true			
e0a	node2_clus1	up/up	169.254.47.194/16	node2
	true			
e0b	node2_clus2	up/up	169.254.19.183/16	node2
	true			

4 entries were displayed.

6. Verify that auto-revert is enabled on all cluster LIFs:

```
network interface show -vserver Cluster -fields auto-revert
```

Show example

```
cluster1::*> network interface show -vserver Cluster -fields auto-revert
```

Vserver	Logical	Auto-revert
Interface		

Cluster		
	node1_clus1	true
	node1_clus2	true
	node2_clus1	true
	node2_clus2	true

4 entries were displayed.

7. Disconnect the cable from cluster port e0a on node1, and then connect e0a to port 1 on cluster switch cs1, using the appropriate cabling supported by the 92300YC switches.

The [Hardware Universe - Switches](#) contains more information about cabling.

8. Disconnect the cable from cluster port e0a on node2, and then connect e0a to port 2 on cluster switch cs1, using the appropriate cabling supported by the 92300YC switches.
9. Enable all node-facing ports on cluster switch cs1.

Show example

The following example shows that ports 1/1 through 1/64 are enabled on switch cs1:

```
cs1# config
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# interface e1/1-64
cs1(config-if-range)# no shutdown
```

10. Verify that all cluster LIFs are up, operational, and display as true for Is Home:

```
network interface show -vserver Cluster
```

Show example

The following example shows that all of the LIFs are up on node1 and node2 and that Is Home results are true:

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current	
Current Is	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	-----
Cluster					
true	node1_clus1	up/up	169.254.209.69/16	node1	e0a
true	node1_clus2	up/up	169.254.49.125/16	node1	e0b
true	node2_clus1	up/up	169.254.47.194/16	node2	e0a
true	node2_clus2	up/up	169.254.19.183/16	node2	e0b

4 entries were displayed.

11. Display information about the status of the nodes in the cluster:

```
cluster show
```

Show example

The following example displays information about the health and eligibility of the nodes in the cluster:

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
node1	true	true	false
node2	true	true	false

2 entries were displayed.

12. Disconnect the cable from cluster port e0b on node1, and then connect e0b to port 1 on cluster switch cs2, using the appropriate cabling supported by the 92300YC switches.
13. Disconnect the cable from cluster port e0b on node2, and then connect e0b to port 2 on cluster switch cs2, using the appropriate cabling supported by the 92300YC switches.
14. Enable all node-facing ports on cluster switch cs2.

Show example

The following example shows that ports 1/1 through 1/64 are enabled on switch cs2:

```
cs2# config  
Enter configuration commands, one per line. End with CNTL/Z.  
cs2(config)# interface e1/1-64  
cs2(config-if-range)# no shutdown
```

Step 3: Verify the configuration

1. Verify that all cluster ports are up:

```
network port show -ipSpace Cluster
```

Show example

The following example shows that all of the cluster ports are up on node1 and node2:

```
cluster1::*> network port show -ipspace Cluster

Node: node1

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
e0a      Cluster      Cluster      up    9000  auto/10000
healthy  false
e0b      Cluster      Cluster      up    9000  auto/10000
healthy  false

Node: node2

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
e0a      Cluster      Cluster      up    9000  auto/10000
healthy  false
e0b      Cluster      Cluster      up    9000  auto/10000
healthy  false

4 entries were displayed.
```

2. Verify that all interfaces display true for Is Home:

```
network interface show -vserver Cluster
```



This might take several minutes to complete.

Show example

The following example shows that all LIFs are up on node1 and node2 and that Is Home results are true:

```
cluster1::*> network interface show -vserver Cluster
```

```
Current Is Logical Status Network Current
Vserver Interface Admin/Oper Address/Mask Node Port
Home
-----
Cluster
true node1_clus1 up/up 169.254.209.69/16 node1 e0a
true node1_clus2 up/up 169.254.49.125/16 node1 e0b
true node2_clus1 up/up 169.254.47.194/16 node2 e0a
true node2_clus2 up/up 169.254.19.183/16 node2 e0b
true
4 entries were displayed.
```

3. Verify that both nodes each have one connection to each switch:

```
show cdp neighbors
```

Show example

The following example shows the appropriate results for both switches:

```
(cs1)# 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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0a	Eth1/1	133	H	FAS2980
node2 e0a	Eth1/2	133	H	FAS2980
cs2 (FDO220329V5) Eth1/65	Eth1/65	175	R S I s	N9K-C92300YC
cs2 (FDO220329V5) Eth1/66	Eth1/66	175	R S I s	N9K-C92300YC

```
Total entries displayed: 4
```

```
(cs2)# 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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0b	Eth1/1	133	H	FAS2980
node2 e0b	Eth1/2	133	H	FAS2980
cs1 (FDO220329KU) Eth1/65	Eth1/65	175	R S I s	N9K-C92300YC
cs1 (FDO220329KU) Eth1/66	Eth1/66	175	R S I s	N9K-C92300YC

```
Total entries displayed: 4
```

4. Display information about the discovered network devices in your cluster:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol       Port   Device (LLDP: ChassisID)  Interface
Platform
-----
node2          /cdp
               e0a    cs1                        0/2          N9K-
C92300YC
               e0b    cs2                        0/2          N9K-
C92300YC
node1          /cdp
               e0a    cs1                        0/1          N9K-
C92300YC
               e0b    cs2                        0/1          N9K-
C92300YC

4 entries were displayed.
```

5. Verify that the settings are disabled:

```
network options switchless-cluster show
```



It might take several minutes for the command to complete. Wait for the '3 minute lifetime to expire' announcement.

Show example

The false output in the following example shows that the configuration settings are disabled:

```
cluster1::*> network options switchless-cluster show
Enable Switchless Cluster: false
```

6. Verify the status of the node members in the cluster:

```
cluster show
```


Show example

The following example shows information about the health and eligibility of the nodes in the cluster:

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
node1	true	true	false
node2	true	true	false

7. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```
cluster1::~*> cluster ping-cluster -node local
```

```
Host is node2
```

```
Getting addresses from network interface table...
```

```
Cluster node1_clus1 169.254.209.69 node1 e0a
```

```
Cluster node1_clus2 169.254.49.125 node1 e0b
```

```
Cluster node2_clus1 169.254.47.194 node2 e0a
```

```
Cluster node2_clus2 169.254.19.183 node2 e0b
```

```
Local = 169.254.47.194 169.254.19.183
```

```
Remote = 169.254.209.69 169.254.49.125
```

```
Cluster Vserver Id = 4294967293
```

```
Ping status:
```

```
Basic connectivity succeeds on 4 path(s)
```

```
Basic connectivity fails on 0 path(s)
```

```
Detected 9000 byte MTU on 4 path(s):
```

```
Local 169.254.47.194 to Remote 169.254.209.69
```

```
Local 169.254.47.194 to Remote 169.254.49.125
```

```
Local 169.254.19.183 to Remote 169.254.209.69
```

```
Local 169.254.19.183 to Remote 169.254.49.125
```

```
Larger than PMTU communication succeeds on 4 path(s)
```

```
RPC status:
```

```
2 paths up, 0 paths down (tcp check)
```

```
2 paths up, 0 paths down (udp check)
```

8. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Show example

```
cluster1::~*> system node autosupport invoke -node * -type all  
-message MAINT=END
```

9. Change the privilege level back to admin:

```
set -privilege admin
```

What's next?

[Configure switch health monitoring.](#)

Migrate from a Cisco switch to a Cisco Nexus 92300YC switch

You can migrate nondisruptively older Cisco cluster switches for an ONTAP cluster to Cisco Nexus 92300YC cluster network switches.



After your migration completes, you might need to install the required configuration file to support the Cluster Switch Health Monitor (CSHM) for 92300YC cluster switches. See [Install the Cluster Switch Health Monitor \(CSHM\)](#).

Review requirements

What you'll need

- A fully functional existing cluster.
- 10 GbE and 40 GbE connectivity from nodes to Nexus 92300YC cluster switches.
- All cluster ports are in the up state to ensure nondisruptive operations.
- Proper version of NX-OS and reference configuration file (RCF) installed on the Nexus 92300YC cluster switches.
- A redundant and fully functional NetApp cluster using both older Cisco switches.
- Management connectivity and console access to both the older Cisco switches and the new switches.
- All cluster LIFs in the up state with the cluster LIFs are on their home ports.
- ISL ports enabled and cabled between the older Cisco switches and between the new switches.

Migrate the switch

About the examples

The examples in this procedure use the following switch and node nomenclature:

- The existing Cisco Nexus 5596UP cluster switches are c1 and c2.
- The new Nexus 92300YC cluster switches are cs1 and cs2.
- The nodes are node1 and node2.
- The cluster LIFs are node1_clus1 and node1_clus2 on node 1, and node2_clus1 and node2_clus2 on node 2 respectively.
- Switch c2 is replaced by switch cs2 first and then switch c1 is replaced by switch cs1.
 - A temporary ISL is built on cs1 connecting c1 to cs1.
 - Cabling between the nodes and c2 are then disconnected from c2 and reconnected to cs2.
 - Cabling between the nodes and c1 are then disconnected from c1 and reconnected to cs1.
 - The temporary ISL between c1 and cs1 is then removed.

Ports used for connections

- Some of the ports are configured on Nexus 92300YC switches to run at 10 GbE or 40 GbE.
- The cluster switches use the following ports for connections to nodes:
 - Ports e1/1-48 (10/25 GbE), e1/49-64 (40/100 GbE): Nexus 92300YC
 - Ports e1/1-40 (10 GbE): Nexus 5596UP
 - Ports e1/1-32 (10 GbE): Nexus 5020

- Ports e1/1-12, e2/1-6 (10 GbE): Nexus 5010 with expansion module
- The cluster switches use the following Inter-Switch Link (ISL) ports:
 - Ports e1/65-66 (100 GbE): Nexus 92300YC
 - Ports e1/41-48 (10 GbE): Nexus 5596UP
 - Ports e1/33-40 (10 GbE): Nexus 5020
 - Ports e1/13-20 (10 GbE): Nexus 5010
- [Hardware Universe - Switches](#) contains information about supported cabling for all cluster switches.
- The ONTAP and NX-OS versions supported in this procedure are on the [Cisco Ethernet Switches](#) page.

Step 1: Prepare for migration

1. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (*>) appears.

2. 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
```

where x is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

Show example

The following command suppresses automatic case creation for two hours:

```
cluster1::*> system node autosupport invoke -node * -type all  
-message MAINT=2h
```

3. Verify that auto-revert is enabled on all cluster LIFs:

```
network interface show -vserver Cluster -fields auto-revert
```

Show example

```
cluster1::*> network interface show -vserver Cluster -fields auto-revert
```

Vserver	Logical Interface	Auto-revert
Cluster	node1_clus1	true
	node1_clus2	true
	node2_clus1	true
	node2_clus2	true

4 entries were displayed.

4. Determine the administrative or operational status for each cluster interface:

Each port should display up for Link and healthy for Health Status.

a. Display the network port attributes:

```
network port show -ipspace Cluster
```

Show example

```
cluster1::*> network port show -ipspace Cluster

Node: node1

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper  Speed (Mbps)
Status      Status
-----
e0a         Cluster    Cluster          up   9000  auto/10000
healthy     false
e0b         Cluster    Cluster          up   9000  auto/10000
healthy     false

Node: node2

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper  Speed (Mbps)
Status      Status
-----
e0a         Cluster    Cluster          up   9000  auto/10000
healthy     false
e0b         Cluster    Cluster          up   9000  auto/10000
healthy     false

4 entries were displayed.
```

- b. Display information about the logical interfaces and their designated home nodes:

```
network interface show -vserver Cluster
```

Each LIF should display up/up for Status Admin/Oper and true for Is Home.

Show example

```
cluster1::*> network interface show -vserver Cluster
```

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
e0a	node1_clus1	up/up	169.254.209.69/16	node1
e0b	true			
e0a	node1_clus2	up/up	169.254.49.125/16	node1
e0b	true			
e0a	node2_clus1	up/up	169.254.47.194/16	node2
e0b	true			
e0a	node2_clus2	up/up	169.254.19.183/16	node2
e0b	true			

4 entries were displayed.

5. Verify that the cluster ports on each node are connected to existing cluster switches in the following way (from the nodes' perspective) using the command:

```
network device-discovery show -protocol cdp
```


Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port  Device (LLDP: ChassisID)  Interface
Platform
-----
node2         /cdp
              e0a   c1                        0/2          N5K-
C5596UP
              e0b   c2                        0/2          N5K-
C5596UP
node1         /cdp
              e0a   c1                        0/1          N5K-
C5596UP
              e0b   c2                        0/1          N5K-
C5596UP

4 entries were displayed.
```

6. Verify that the cluster ports and switches are connected in the following way (from the switches' perspective) using the command:

```
show cdp neighbors
```

Show example

```
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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0a	Eth1/1	124	H	FAS2750
node2 e0a	Eth1/2	124	H	FAS2750
c2 (FOX2025GEFC) Eth1/41	Eth1/41	179	S I s	N5K-C5596UP
c2 (FOX2025GEFC) Eth1/42	Eth1/42	175	S I s	N5K-C5596UP
c2 (FOX2025GEFC) Eth1/43	Eth1/43	179	S I s	N5K-C5596UP
c2 (FOX2025GEFC) Eth1/44	Eth1/44	175	S I s	N5K-C5596UP
c2 (FOX2025GEFC) Eth1/45	Eth1/45	179	S I s	N5K-C5596UP
c2 (FOX2025GEFC) Eth1/46	Eth1/46	179	S I s	N5K-C5596UP
c2 (FOX2025GEFC) Eth1/47	Eth1/47	175	S I s	N5K-C5596UP
c2 (FOX2025GEFC) Eth1/48	Eth1/48	179	S I s	N5K-C5596UP

```
Total entries displayed: 10
```

```
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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0b	Eth1/1	124	H	FAS2750
node2 e0b	Eth1/2	124	H	FAS2750
c1 (FOX2025GEEX) Eth1/41	Eth1/41	175	S I s	N5K-C5596UP
c1 (FOX2025GEEX) Eth1/42	Eth1/42	175	S I s	N5K-C5596UP
c1 (FOX2025GEEX) Eth1/43	Eth1/43	175	S I s	N5K-C5596UP
c1 (FOX2025GEEX) Eth1/44	Eth1/44	175	S I s	N5K-C5596UP
c1 (FOX2025GEEX) Eth1/45	Eth1/45	175	S I s	N5K-C5596UP
c1 (FOX2025GEEX) Eth1/46	Eth1/46	175	S I s	N5K-C5596UP
c1 (FOX2025GEEX) Eth1/47	Eth1/47	176	S I s	N5K-C5596UP
c1 (FOX2025GEEX) Eth1/48	Eth1/48	176	S I s	N5K-C5596UP

7. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
    Local 169.254.19.183 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 2: Configure cables and ports

1. Configure a temporary ISL on cs1 on ports e1/41-48, between c1 and cs1.

Show example

The following example shows how the new ISL is configured on c1 and cs1:

```
cs1# configure
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# interface e1/41-48
cs1(config-if-range)# description temporary ISL between Nexus 5596UP
and Nexus 92300YC
cs1(config-if-range)# no lldp transmit
cs1(config-if-range)# no lldp receive
cs1(config-if-range)# switchport mode trunk
cs1(config-if-range)# no spanning-tree bpduguard enable
cs1(config-if-range)# channel-group 101 mode active
cs1(config-if-range)# exit
cs1(config)# interface port-channel 101
cs1(config-if)# switchport mode trunk
cs1(config-if)# spanning-tree port type network
cs1(config-if)# exit
cs1(config)# exit
```

2. Remove ISL cables from ports e1/41-48 from c2 and connect the cables to ports e1/41-48 on cs1.
3. Verify that the ISL ports and port-channel are operational connecting c1 and cs1:

```
show port-channel summary
```

Show example

The following example shows the Cisco show port-channel summary command being used to verify the ISL ports are operational on c1 and cs1:

```
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
        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      Eth1/41(P)  Eth1/42(P)
Eth1/43(P)
                                     Eth1/44(P)  Eth1/45(P)
Eth1/46(P)
                                     Eth1/47(P)  Eth1/48(P)
```

```
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      Eth1/65(P)  Eth1/66(P)
101   Po101(SU)   Eth       LACP      Eth1/41(P)  Eth1/42(P)
Eth1/43(P)
                                     Eth1/44(P)  Eth1/45(P)
Eth1/46(P)
                                     Eth1/47(P)  Eth1/48(P)
```


4. For node1, disconnect the cable from e1/1 on c2, and then connect the cable to e1/1 on cs2, using appropriate cabling supported by Nexus 92300YC.
5. For node2, disconnect the cable from e1/2 on c2, and then connect the cable to e1/2 on cs2, using appropriate cabling supported by Nexus 92300YC.
6. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp

Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface
Platform
-----
node2      /cdp
           e0a    c1                          0/2          N5K-
C5596UP
           e0b    cs2                         0/2          N9K-
C92300YC
node1      /cdp
           e0a    c1                          0/1          N5K-
C5596UP
           e0b    cs2                         0/1          N9K-
C92300YC

4 entries were displayed.
```

7. For node1, disconnect the cable from e1/1 on c1, and then connect the cable to e1/1 on cs1, using appropriate cabling supported by Nexus 92300YC.
8. For node2, disconnect the cable from e1/2 on c1, and then connect the cable to e1/2 on cs1, using appropriate cabling supported by Nexus 92300YC.
9. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/          Local  Discovered
Protocol      Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
node2         /cdp
              e0a    cs1                       0/2           N9K-
C92300YC
              e0b    cs2                       0/2           N9K-
C92300YC
node1         /cdp
              e0a    cs1                       0/1           N9K-
C92300YC
              e0b    cs2                       0/1           N9K-
C92300YC
4 entries were displayed.
```

10. Delete the temporary ISL between cs1 and c1.

Show example

```
cs1(config)# no interface port-channel 10
cs1(config)# interface e1/41-48
cs1(config-if-range)# lldp transmit
cs1(config-if-range)# lldp receive
cs1(config-if-range)# no switchport mode trunk
cs1(config-if-range)# no channel-group
cs1(config-if-range)# description 10GbE Node Port
cs1(config-if-range)# spanning-tree bpduguard enable
cs1(config-if-range)# exit
cs1(config)# exit
```

Step 3: Complete the migration

1. Verify the final configuration of the cluster:

```
network port show -ipspace Cluster
```

Each port should display up for Link and healthy for Health Status.

Show example

```
cluster1::*> network port show -ipspace Cluster
```

```
Node: node1
```

```
Ignore
```

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
e0a	Cluster	Cluster		up	9000	auto/10000	
healthy	false						
e0b	Cluster	Cluster		up	9000	auto/10000	
healthy	false						

```
Node: node2
```

```
Ignore
```

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
e0a	Cluster	Cluster		up	9000	auto/10000	
healthy	false						
e0b	Cluster	Cluster		up	9000	auto/10000	
healthy	false						

```
4 entries were displayed.
```

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			
-----	-----	-----	-----	-----
-----	-----			
Cluster				
	node1_clus1	up/up	169.254.209.69/16	node1
e0a	true			

```

node1_clus2 up/up 169.254.49.125/16 node1
e0b true
node2_clus1 up/up 169.254.47.194/16 node2
e0a true
node2_clus2 up/up 169.254.19.183/16 node2
e0b true

```

4 entries were displayed.

```
cluster1::*> network device-discovery show -protocol cdp
```

```

Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface
Platform
-----
node2      /cdp
           e0a    cs1                        0/2          N9K-
C92300YC
           e0b    cs2                        0/2          N9K-
C92300YC
node1      /cdp
           e0a    cs1                        0/1          N9K-
C92300YC
           e0b    cs2                        0/1          N9K-
C92300YC

```

4 entries were displayed.

```
cs1# 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
node1          Eth1/1        124     H           FAS2750
e0a
node2          Eth1/2        124     H           FAS2750
e0a
cs2 (FDO220329V5)  Eth1/65      179     R S I s     N9K-C92300YC
Eth1/65

```

```
cs2 (FDO220329V5)      Eth1/66      179      R S I s      N9K-C92300YC
Eth1/66
```

```
cs2# 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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0b	Eth1/1	124	H	FAS2750
node2 e0b	Eth1/2	124	H	FAS2750
cs1 (FDO220329KU) Eth1/65	Eth1/65	179	R S I s	N9K-C92300YC
cs1 (FDO220329KU) Eth1/66	Eth1/66	179	R S I s	N9K-C92300YC

```
Total entries displayed: 4
```

2. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
    Local 169.254.19.183 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)

```

3. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Show example

```

cluster1::*> system node autosupport invoke -node * -type all
-message MAINT=END

```

4. Change the privilege level back to admin:

```
set -privilege admin
```

What's next?

[Configure switch health monitoring.](#)

Replace switches

Replace a Cisco Nexus 92300YC switch

Replacing a defective Nexus 92300YC switch in a cluster network is a nondisruptive procedure (NDU).

Review requirements

What you'll need

Before performing the switch replacement, ensure that:

- In the existing cluster and network infrastructure:
 - The existing cluster is verified as completely functional, with at least one fully connected cluster switch.
 - All cluster ports are up.
 - All cluster logical interfaces (LIFs) are up and on their home ports.
 - The ONTAP cluster ping-cluster -node node1 command must indicate that basic connectivity and larger than PMTU communication are successful on all paths.
- For the Nexus 92300YC replacement switch:
 - Management network connectivity on the replacement switch are functional.
 - Console access to the replacement switch are in place.
 - The node connections are ports 1/1 through 1/64.
 - All Inter-Switch Link (ISL) ports are disabled on ports 1/65 and 1/66.
 - The desired reference configuration file (RCF) and NX-OS operating system image switch are loaded onto the switch.
 - Initial customization of the switch are complete, as detailed in: [Configure the Cisco Nexus 92300YC switch](#).

Any previous site customizations, such as STP, SNMP, and SSH, are copied to the new switch.

Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when replacing your switch:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport before and after maintenance to disable case creation for the duration of the maintenance. See this Knowledge Base article [SU92: How to suppress automatic case creation during scheduled maintenance windows](#) for further details.
- Enable session logging for any CLI sessions. For instructions on how to enable session logging, review the "Logging Session Output" section in this Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

Replace the switch

About the examples

The examples in this procedure use the following switch and node nomenclature:

- The names of the existing Nexus 92300YC switches are cs1 and cs2.
- The name of the new Nexus 92300YC switch is newcs2.
- The node names are node1 and node2.
- The cluster ports on each node are named e0a and e0b.
- The cluster LIF names are node1_clus1 and node1_clus2 for node1, and node2_clus1 and node2_clus2 for node2.
- The prompt for changes to all cluster nodes is cluster1::<*>

About this task

You must execute the command for migrating a cluster LIF from the node where the cluster LIF is hosted.

The following procedure is based on the following cluster network topology:

Show topology

```
cluster1::*> network port show -ipspace Cluster
```

```
Node: node1
```

```
Ignore
```

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							

```
-----  
-----  
e0a      Cluster      Cluster      up    9000  auto/10000 healthy  
false  
e0b      Cluster      Cluster      up    9000  auto/10000 healthy  
false
```

```
Node: node2
```

```
Ignore
```

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							

```
-----  
-----  
e0a      Cluster      Cluster      up    9000  auto/10000 healthy  
false  
e0b      Cluster      Cluster      up    9000  auto/10000 healthy  
false
```

```
4 entries were displayed.
```

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current	
Current Is					
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
Cluster					
	node1_clus1	up/up	169.254.209.69/16	node1	e0a
true	node1_clus2	up/up	169.254.49.125/16	node1	e0b

```

true
      node2_clus1  up/up    169.254.47.194/16  node2      e0a
true
      node2_clus2  up/up    169.254.19.183/16  node2      e0b
true
4 entries were displayed.

```

```
cluster1::*> network device-discovery show -protocol cdp
```

Node/	Local	Discovered		
Protocol	Port	Device (LLDP: ChassisID)	Interface	Platform
node2	/cdp			
	e0a	cs1	Eth1/2	N9K-
C92300YC				
	e0b	cs2	Eth1/2	N9K-
C92300YC				
node1	/cdp			
	e0a	cs1	Eth1/1	N9K-
C92300YC				
	e0b	cs2	Eth1/1	N9K-
C92300YC				

```
4 entries were displayed.
```

```
cs1# 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
node1	Eth1/1	144	H	FAS2980	e0a
node2	Eth1/2	145	H	FAS2980	e0a
cs2 (FDO220329V5)	Eth1/65	176	R S I s	N9K-C92300YC	
Eth1/65					
cs2 (FDO220329V5)	Eth1/66	176	R S I s	N9K-C92300YC	
Eth1/66					

```
Total entries displayed: 4
```

```
cs2# 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 ID	Local Intrfce	Hldtme	Capability	Platform	Port
node1	Eth1/1	139	H	FAS2980	e0b
node2	Eth1/2	124	H	FAS2980	e0b
cs1 (FDO220329KU) Eth1/65	Eth1/65	178	R S I s	N9K-C92300YC	
cs1 (FDO220329KU) Eth1/66	Eth1/66	178	R S I s	N9K-C92300YC	

```
Total entries displayed: 4
```

Step 1: Prepare for replacement

1. Install the appropriate RCF and image on the switch, newcs2, and make any necessary site preparations.

If necessary, verify, download, and install the appropriate versions of the RCF and NX-OS software for the new switch. If you have verified that the new switch is correctly set up and does not need updates to the RCF and NX-OS software, continue to step 2.

- a. Go to the *NetApp Cluster and Management Network Switches Reference Configuration File Description Page* on the NetApp Support Site.
 - b. Click the link for the *Cluster Network and Management Network Compatibility Matrix*, and then note the required switch software version.
 - c. Click your browser's back arrow to return to the **Description** page, click **CONTINUE**, accept the license agreement, and then go to the **Download** page.
 - d. Follow the steps on the Download page to download the correct RCF and NX-OS files for the version of ONTAP software you are installing.
2. On the new switch, log in as admin and shut down all of the ports that will be connected to the node cluster interfaces (ports 1/1 to 1/64).

If the switch that you are replacing is not functional and is powered down, go to Step 4. The LIFs on the cluster nodes should have already failed over to the other cluster port for each node.

Show example

```
newcs2# config  
Enter configuration commands, one per line. End with CNTL/Z.  
newcs2(config)# interface e1/1-64  
newcs2(config-if-range)# shutdown
```

3. Verify that all cluster LIFs have auto-revert enabled:

```
network interface show -vserver Cluster -fields auto-revert
```

Show example

```
cluster1::> network interface show -vserver Cluster -fields auto-revert
```

Vserver	Logical Interface	Auto-revert
Cluster	node1_clus1	true
Cluster	node1_clus2	true
Cluster	node2_clus1	true
Cluster	node2_clus2	true

4 entries were displayed.

4. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

Step 2: Configure cables and ports

1. Shut down the ISL ports 1/65 and 1/66 on the Nexus 92300YC switch cs1:

Show example

```

cs1# configure
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# interface e1/65-66
cs1(config-if-range)# shutdown
cs1(config-if-range)#

```

2. Remove all of the cables from the Nexus 92300YC cs2 switch, and then connect them to the same ports on the Nexus 92300YC newcs2 switch.
3. Bring up the ISLs ports 1/65 and 1/66 between the cs1 and newcs2 switches, and then verify the port channel operation status.

Port-Channel should indicate Po1(SU) and Member Ports should indicate Eth1/65(P) and Eth1/66(P).

Show example

This example enables ISL ports 1/65 and 1/66 and displays the port channel summary on switch cs1:

```
cs1# configure
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# int e1/65-66
cs1(config-if-range)# no shutdown

cs1(config-if-range)# 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      Eth1/65(P)  Eth1/66(P)

cs1(config-if-range)#
```

4. Verify that port e0b is up on all nodes:

```
network port show ipspace Cluster
```


Show example

The output should be similar to the following:

```
cluster1::*> network port show -ipSpace Cluster

Node: node1

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper  Speed (Mbps)
Status      Status
-----
e0a         Cluster    Cluster          up   9000  auto/10000
healthy    false
e0b         Cluster    Cluster          up   9000  auto/10000
healthy    false

Node: node2

Ignore

Health      Health
Port        IPspace    Broadcast Domain Link MTU  Admin/Oper  Speed (Mbps)
Status      Status
-----
e0a         Cluster    Cluster          up   9000  auto/10000
healthy    false
e0b         Cluster    Cluster          up   9000  auto/auto   -
false

4 entries were displayed.
```

5. On the same node you used in the previous step, revert the cluster LIF associated with the port in the previous step by using the network interface revert command.

7. Verify that all physical cluster ports are up:

```
network port show ipspace Cluster
```

Show example

```
cluster1::*> network port show -ipspace Cluster

Node: node1

Ignore

Health      Health
Port        IPspace    Broadcast Domain  Link  MTU  Admin/Oper
Status      Status
-----
-----
e0a         Cluster    Cluster           up    9000 auto/10000
healthy    false
e0b         Cluster    Cluster           up    9000 auto/10000
healthy    false

Node: node2

Ignore

Health      Health
Port        IPspace    Broadcast Domain  Link  MTU  Admin/Oper
Status      Status
-----
-----
e0a         Cluster    Cluster           up    9000 auto/10000
healthy    false
e0b         Cluster    Cluster           up    9000 auto/10000
healthy    false

4 entries were displayed.
```

Step 3: Complete the procedure

1. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```
cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

2. Confirm the following cluster network configuration:

```
network port show
```

Show example

```
cluster1::*> network port show -ipspace Cluster
```

```
Node: node1
```

```
Ignore
```

				Speed (Mbps)			Health
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status

```
-----  
-----
```

e0a	Cluster	Cluster		up	9000	auto/10000	
healthy	false						
e0b	Cluster	Cluster		up	9000	auto/10000	
healthy	false						

```
Node: node2
```

```
Ignore
```

				Speed (Mbps)			Health
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status

```
-----  
-----
```

e0a	Cluster	Cluster		up	9000	auto/10000	
healthy	false						
e0b	Cluster	Cluster		up	9000	auto/10000	
healthy	false						

```
4 entries were displayed.
```

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

```
-----  
-----
```

Cluster				
	node1_clus1	up/up	169.254.209.69/16	node1
e0a	true			
	node1_clus2	up/up	169.254.49.125/16	node1

```

e0b      true
          node2_clus1  up/up    169.254.47.194/16  node2
e0a      true
          node2_clus2  up/up    169.254.19.183/16  node2
e0b      true

```

4 entries were displayed.

```
cluster1::> network device-discovery show -protocol cdp
```

```

Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface
Platform
-----
node2      /cdp
          e0a    cs1                        0/2          N9K-
C92300YC
          e0b    newcs2                    0/2          N9K-
C92300YC
node1      /cdp
          e0a    cs1                        0/1          N9K-
C92300YC
          e0b    newcs2                    0/1          N9K-
C92300YC

```

4 entries were displayed.

```
cs1# 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
node1          Eth1/1        144     H           FAS2980
e0a
node2          Eth1/2        145     H           FAS2980
e0a
newcs2 (FDO296348FU)  Eth1/65      176     R S I s     N9K-C92300YC
Eth1/65
newcs2 (FDO296348FU)  Eth1/66      176     R S I s     N9K-C92300YC

```

```
Eth1/66
```

```
Total entries displayed: 4
```

```
cs2# 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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0b	Eth1/1	139	H	FAS2980
node2 e0b	Eth1/2	124	H	FAS2980
cs1 (FDO220329KU) Eth1/65	Eth1/65	178	R S I s	N9K-C92300YC
cs1 (FDO220329KU) Eth1/66	Eth1/66	178	R S I s	N9K-C92300YC

```
Total entries displayed: 4
```

What's next?

[Configure switch health monitoring.](#)

Replace Cisco Nexus 92300YC cluster switches with switchless connections

You can migrate from a cluster with a switched cluster network to one where two nodes are directly connected for ONTAP 9.3 and later.

Review requirements

Guidelines

Review the following guidelines:

- Migrating to a two-node switchless cluster configuration is a nondisruptive operation. Most systems have two dedicated cluster interconnect ports on each node, but you can also use this procedure for systems with a larger number of dedicated cluster interconnect ports on each node, such as four, six or eight.
- You cannot use the switchless cluster interconnect feature with more than two nodes.
- If you have an existing two-node cluster that uses cluster interconnect switches and is running ONTAP 9.3 or later, you can replace the switches with direct, back-to-back connections between the nodes.


What you'll need

- A healthy cluster that consists of two nodes connected by cluster switches. The nodes must be running the same ONTAP release.
- Each node with the required number of dedicated cluster ports, which provide redundant cluster interconnect connections to support your system configuration. For example, there are two redundant ports for a system with two dedicated cluster interconnect ports on each node.

Migrate the switches

About this task

The following procedure removes the cluster switches in a two-node cluster and replaces each connection to the switch with a direct connection to the partner node.



About the examples

The examples in the following procedure show nodes that are using "e0a" and "e0b" as cluster ports. Your nodes might be using different cluster ports as they vary by system.

Step 1: Prepare for migration

1. Change the privilege level to advanced, entering `y` when prompted to continue:

```
set -privilege advanced
```

The advanced prompt `*>` appears.

2. ONTAP 9.3 and later supports automatic detection of switchless clusters, which is enabled by default.

You can verify that detection of switchless clusters is enabled by running the advanced privilege command:

```
network options detect-switchless-cluster show
```

Show example

The following example output shows if the option is enabled.

```
cluster::*> network options detect-switchless-cluster show
(network options detect-switchless-cluster show)
Enable Switchless Cluster Detection: true
```

If "Enable Switchless Cluster Detection" is `false`, contact NetApp support.

3. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message
MAINT=<number_of_hours>h
```

where `h` is the duration of the maintenance window in hours. The message notifies technical support of this maintenance task so that they can suppress automatic case creation during the maintenance window.

In the following example, the command suppresses automatic case creation for two hours:

Show example


```
cluster::*> system node autosupport invoke -node * -type all
-message MAINT=2h
```

Step 2: Configure ports and cabling

1. Organize the cluster ports on each switch into groups so that the cluster ports in group1 go to cluster switch1 and the cluster ports in group2 go to cluster switch2. These groups are required later in the procedure.
2. Identify the cluster ports and verify link status and health:

```
network port show -ipSpace Cluster
```

In the following example for nodes with cluster ports "e0a" and "e0b", one group is identified as "node1:e0a" and "node2:e0a" and the other group as "node1:e0b" and "node2:e0b". Your nodes might be using different cluster ports because they vary by system.



Verify that the ports have a value of `up` for the “Link” column and a value of `healthy` for the “Health Status” column.

Show example

```

cluster::> network port show -ipspace Cluster
Node: node1

Ignore
Speed (Mbps) Health
Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status
Status
-----
-----
e0a Cluster Cluster up 9000 auto/10000 healthy
false
e0b Cluster Cluster up 9000 auto/10000 healthy
false

Node: node2

Ignore
Speed (Mbps) Health
Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status
Status
-----
-----
e0a Cluster Cluster up 9000 auto/10000 healthy
false
e0b Cluster Cluster up 9000 auto/10000 healthy
false
4 entries were displayed.

```

3. Confirm that all the cluster LIFs are on their home ports.

Verify that the “is-home” column is true for each of the cluster LIFs:

```
network interface show -vserver Cluster -fields is-home
```

Show example

```
cluster::*> net int show -vserver Cluster -fields is-home
(network interface show)
vserver  lif          is-home
-----  -
Cluster  node1_clus1  true
Cluster  node1_clus2  true
Cluster  node2_clus1  true
Cluster  node2_clus2  true
4 entries were displayed.
```

If there are cluster LIFs that are not on their home ports, revert those LIFs to their home ports:

```
network interface revert -vserver Cluster -lif *
```

4. Disable auto-revert for the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

5. Verify that all ports listed in the previous step are connected to a network switch:

```
network device-discovery show -port cluster_port
```

The “Discovered Device” column should be the name of the cluster switch that the port is connected to.

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to cluster switches "cs1" and "cs2".

```
cluster::> network device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol  Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1/cdp
          e0a    cs1                        0/11       BES-53248
          e0b    cs2                        0/12       BES-53248
node2/cdp
          e0a    cs1                        0/9        BES-53248
          e0b    cs2                        0/9        BES-53248
4 entries were displayed.
```

6. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

7. Verify that the cluster is healthy:

```
cluster ring show
```

All units must be either master or secondary.


8. Set up the switchless configuration for the ports in group 1.



To avoid potential networking issues, you must disconnect the ports from group1 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.


a. Disconnect all the cables from the ports in group1 at the same time.

In the following example, the cables are disconnected from port "e0a" on each node, and cluster traffic continues through the switch and port "e0b" on each node:



b. Cable the ports in group1 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2:



9. The switchless cluster network option transitions from *false* to *true*. This might take up to 45 seconds. Confirm that the switchless option is set to *true*:

```
network options switchless-cluster show
```

The following example shows that the switchless cluster is enabled:

```
cluster::*> network options switchless-cluster show
Enable Switchless Cluster: true
```

10. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```



Before proceeding to the next step, you must wait at least two minutes to confirm a working back-to-back connection on group 1.


11. Set up the switchless configuration for the ports in group 2.



To avoid potential networking issues, you must disconnect the ports from group2 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.


- a. Disconnect all the cables from the ports in group2 at the same time.

In the following example, the cables are disconnected from port "e0b" on each node, and cluster traffic continues through the direct connection between the "e0a" ports:



b. Cable the ports in group2 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2 and "e0b" on node1 is connected to "e0b" on node2:



Step 3: Verify the configuration

1. Verify that the ports on both nodes are correctly connected:

```
network device-discovery show -port cluster_port
```

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to the corresponding port on the cluster partner:

```
cluster::> net device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1/cdp
          e0a    node2                      e0a        AFF-A300
          e0b    node2                      e0b        AFF-A300
node1/lldp
          e0a    node2 (00:a0:98:da:16:44)  e0a        -
          e0b    node2 (00:a0:98:da:16:44)  e0b        -
node2/cdp
          e0a    node1                      e0a        AFF-A300
          e0b    node1                      e0b        AFF-A300
node2/lldp
          e0a    node1 (00:a0:98:da:87:49)  e0a        -
          e0b    node1 (00:a0:98:da:87:49)  e0b        -
8 entries were displayed.
```

2. Re-enable auto-revert for the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert true
```

3. Verify that all LIFs are home. This might take a few seconds.

```
network interface show -vserver Cluster -lif lif_name
```

Show example

The LIFs have been reverted if the “Is Home” column is `true`, as shown for `node1_clus2` and `node2_clus2` in the following example:

```
cluster::> network interface show -vserver Cluster -fields curr-
port,is-home
vserver  lif                curr-port  is-home
-----  -
Cluster  node1_clus1            e0a        true
Cluster  node1_clus2            e0b        true
Cluster  node2_clus1            e0a        true
Cluster  node2_clus2            e0b        true
4 entries were displayed.
```

If any cluster LIFS have not returned to their home ports, revert them manually from the local node:

```
network interface revert -vserver Cluster -lif lif_name
```

4. Check the cluster status of the nodes from the system console of either node:

```
cluster show
```

Show example

The following example shows `epsilon` on both nodes to be `false`:

```
Node  Health  Eligibility  Epsilon
-----  -
node1 true    true        false
node2 true    true        false
2 entries were displayed.
```

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

6. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

For more information, see [NetApp KB Article 1010449: How to suppress automatic case creation during scheduled maintenance windows](#).

7. Change the privilege level back to admin:

```
set -privilege admin
```

NetApp CN1610

Overview of installation and configuration for NetApp CN1610 switches

The CN1610 is a high bandwidth, managed Layer 2 switch that provides 16 10-Gigabit Small Form-Factor Pluggable Plus (SFP+) ports.

The switch includes redundant power supplies and fan trays that support hot swapping for high availability. This 1U switch can be installed in a standard 19-inch NetApp 42U system cabinet or third-party cabinet.

The switch supports local management through the console port or remote management by using Telnet or SSH through a network connection. The CN1610 includes a dedicated 1-Gigabit Ethernet RJ45 management port for out-of-band switch management. You can manage the switch by entering commands into the command-line interface (CLI) or by using an SNMP-based network management system (NMS).

Install and configure workflow for NetApp CN1610 switches

To install and configure a NetApp CN1610 switch on systems running ONTAP, follow these steps:

1. [Install hardware](#)
2. [Install FASTPATH software](#)
3. [Install Reference Configuration file](#)

If the switches are running ONTAP 8.3.1 or later, follow the instructions in [Install FASTPATH and RCFs on switches running ONTAP 8.3.1 and later](#).

4. [Configure switch](#)

Documentation requirements for NetApp CN1610 switches

For NetApp CN1610 switch installation and maintenance, be sure to review all the recommended documentation.

Document title	Description
1G Installation Guide	An overview of the CN1601 switch hardware and software features and installation process.
10G Installation Guide	An overview of the CN1610 switch hardware and software features and describes the features to install the switch and access the CLI.
CN1601 and CN1610 Switch Setup and Configuration Guide	Details how to configure the switch hardware and software for your cluster environment.
CN1601 Switch Administrator's Guide	Provides examples of how to use the CN1601 switch in a typical network. <ul style="list-style-type: none">• Administrator's Guide• Administrator's Guide, Version 1.1.x.x• Administrator's Guide, Version 1.2.x.x
CN1610 Network Switch CLI Command Reference	Provides detailed information about the command-line interface (CLI) commands you use to configure the CN1601 software. <ul style="list-style-type: none">• Command Reference• Command Reference, Version 1.1.x.x• Command Reference, Version 1.2.x.x

Install and configure

Install the hardware for the NetApp CN1610 switch

To install the NetApp CN1610 switch hardware, use the instructions in one of the following guides.

- [1G Installation Guide](#).

An overview of the CN1601 switch hardware and software features and installation process.

- [10G Installation Guide](#)

An overview of the CN1610 switch hardware and software features and describes the features to install the switch and access the CLI.

Install FASTPATH software

When you install the FASTPATH software on your NetApp switches, you must begin the upgrade with the second switch, `cs2`.

Review requirements

What you'll need

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs and no defective cluster network interface cards (NICs) or similar issues).
- Fully functional port connections on the cluster switch.
- All cluster ports set up.
- All cluster logical interfaces (LIFs) set up (must not have been migrated).
- A successful communication path: The ONTAP (privilege: advanced) `cluster ping-cluster -node node1` command must indicate that larger than PMTU communication is successful on all paths.
- A supported version of FASTPATH and ONTAP.

Make sure you consult the switch compatibility table on the [NetApp CN1601 and CN1610 Switches](#) page for the supported FASTPATH and ONTAP versions.

Install FASTPATH

The following procedure uses the clustered Data ONTAP 8.2 syntax. As a result, the cluster Vserver, LIF names, and CLI output are different than those in Data ONTAP 8.3.

There can be command dependencies between command syntax in the RCF and FASTPATH versions.

About the examples

The examples in this procedure use the following switch and node nomenclature:

- The two NetApp switches are `cs1` and `cs2`.
- The two cluster LIFs are `clus1` and `clus2`.

- The Vservers are vs1 and vs2.
- The `cluster::*>` prompt indicates the name of the cluster.
- The cluster ports on each node are named e1a and e2a.

[Hardware Universe](#) has more information about the actual cluster ports that are supported on your platform.

- The supported Inter-Switch Links (ISLs) are ports 0/13 through 0/16.
- The supported node connections are ports 0/1 through 0/12.

Step 1: Migrate cluster

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all - message MAINT=xh
```

x is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Log into the switch as admin. There is no password by default. At the `(cs2) #` prompt, enter the `enable` command. Again, there is no password by default. This gives you access to Privileged EXEC mode, which allows you to configure the network interface.

Show example

```
(cs2) # enable
Password (Enter)
(cs2) #
```

3. On the console of each node, migrate clus2 to port e1a:

```
network interface migrate
```

Show example

```
cluster::*> network interface migrate -vserver vs1 -lif clus2
-destnode node1 -dest-port e1a
cluster::*> network interface migrate -vserver vs2 -lif clus2
-destnode node2 -dest-port e1a
```

4. On the console of each node, verify that the migration took place:

```
network interface show
```

The following example shows that clus2 has migrated to port e1a on both nodes:

Show example

```
cluster::*> network interface show -role cluster
```

Vserver	Logical Interface	Status Admin/Open	Network Address/Mask	Current Node	Current Port	Is Home
vs1						
	clus1	up/up	10.10.10.1/16	node1	e1a	true
	clus2	up/up	10.10.10.2/16	node1	e1a	
false						
vs2						
	clus1	up/up	10.10.10.1/16	node2	e1a	true
	clus2	up/up	10.10.10.2/16	node2	e1a	
false						

Step 2: Install FASTPATH software

1. Shut down cluster port e2a on both nodes:

```
network port modify
```

Show example

The following example shows port e2a being shut down on both nodes:

```
cluster::*> network port modify -node node1 -port e2a -up-admin  
false  
cluster::*> network port modify -node node2 -port e2a -up-admin  
false
```

2. Verify that port e2a is shut down on both nodes:

```
network port show
```

Show example

```
cluster::*> network port show -role cluster
```

(Mbps)					Auto-Negot	Duplex	Speed
Node	Port	Role	Link	MTU	Admin/Oper	Admin/Oper	Admin/Oper
-----	----	-----	----	-----	-----	-----	
node1							
	e1a	cluster	up	9000	true/true	full/full	auto/10000
	e2a	cluster	down	9000	true/true	full/full	auto/10000
node2							
	e1a	cluster	up	9000	true/true	full/full	auto/10000
	e2a	cluster	down	9000	true/true	full/full	auto/10000

3. Shut down the Inter-Switch Link (ISL) ports on cs1, the active NetApp switch:

Show example

```
(cs1) # configure
(cs1)(config) # interface 0/13-0/16
(cs1)(Interface 0/13-0/16) # shutdown
(cs1)(Interface 0/13-0/16) # exit
(cs1)(config) # exit
```

4. Back up the current active image on cs2.

Show example

```
(cs2) # show bootvar

Image Descriptions .

  active:
  backup:

Images currently available on Flash

-----
--
  unit          active      backup      current-active      next-
active
-----
--

   1           1.1.0.3      1.1.0.1      1.1.0.3             1.1.0.3

(cs2) # copy active backup
Copying active to backup
Copy operation successful

(cs2) #
```

5. Download the image file to the switch.

Copying the image file to the active image means that when you reboot, that image establishes the running FASTPATH version. The previous image remains available as a backup.

Show example

```
(cs2) # copy tftp://10.0.0.1/NetApp_CN1610_1.1.0.5.stk active

Mode..... TFTP
Set Server IP..... 10.0.0.1
Path..... ./
Filename..... NetApp_CN1610_1.1.0.5.stk
Data Type..... Code
Destination Filename..... active

Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y
TFTP Code transfer starting...

File transfer operation completed successfully.
```

6. Verify the running version of the FASTPATH software.

```
show version
```

Show example

```
(cs2) # show version

Switch: 1

System Description..... Broadcom Scorpion 56820
                        Development System - 16 TENGIG,
                        1.1.0.3, Linux 2.6.21.7
Machine Type.....      Broadcom Scorpion 56820
                        Development System - 16TENGIG
Machine Model.....     BCM-56820
Serial Number.....     10611100004
FRU Number.....
Part Number.....       BCM56820
Maintenance Level..... A
Manufacturer.....      0xbc00
Burned In MAC Address..... 00:A0:98:4B:A9:AA
Software Version.....  1.1.0.3
Operating System.....  Linux 2.6.21.7
Network Processing Device..... BCM56820_B0
Additional Packages..... FASTPATH QOS
                        FASTPATH IPv6 Management
```

7. View the boot images for the active and backup configuration.

```
show bootvar
```

Show example

```
(cs2) # show bootvar
```

```
Image Descriptions
```

```
active :
```

```
backup :
```

```
Images currently available on Flash
```

```
-----  
--  
unit          active      backup      current-active      next-  
active  
-----  
--  
1             1.1.0.3      1.1.0.3      1.1.0.3              1.1.0.5
```

8. Reboot the switch.

```
reload
```

Show example

```
(cs2) # reload
```

```
Are you sure you would like to reset the system? (y/n) y
```

```
System will now restart!
```

Step 3: Validate installation

1. Log in again, and verify the new version of the FASTPATH software.

```
show version
```


Show example

```
(cs2) # show version

Switch: 1

System Description..... Broadcom Scorpion 56820
                          Development System - 16
TENGIG,
                          1.1.0.5, Linux 2.6.21.7
Machine Type.....        Broadcom Scorpion 56820
                          Development System - 16TENGIG
Machine Model.....        BCM-56820
Serial Number.....        10611100004
FRU Number.....
Part Number.....          BCM56820
Maintenance Level.....    A
Manufacturer.....          0xbc00
Burned In MAC Address..... 00:A0:98:4B:A9:AA
Software Version.....      1.1.0.5
Operating System.....      Linux 2.6.21.7
Network Processing Device..... BCM56820_B0
Additional Packages.....   FASTPATH QOS
                          FASTPATH IPv6 Management
```

2. Bring up the ISL ports on cs1, the active switch.

```
configure
```

Show example

```
(cs1) # configure
(cs1) (config) # interface 0/13-0/16
(cs1) (Interface 0/13-0/16) # no shutdown
(cs1) (Interface 0/13-0/16) # exit
(cs1) (config) # exit
```

3. Verify that the ISLs are operational:

```
show port-channel 3/1
```

The Link State field should indicate Up.

Show example

```
(cs2) # show port-channel 3/1

Local Interface..... 3/1
Channel Name..... ISL-LAG
Link State..... Up
Admin Mode..... Enabled
Type..... Static
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr      Device/      Port      Port
Ports   Timeout     Speed     Active
-----
0/13    actor/long   10G Full   True
        partner/long
0/14    actor/long   10G Full   True
        partner/long
0/15    actor/long   10G Full   True
        partner/long
0/16    actor/long   10G Full   True
        partner/long
```

4. Copy the running-config file to the startup-config file when you are satisfied with the software versions and switch settings.

Show example

```
(cs2) # write memory

This operation may take a few minutes.
Management interfaces will not be available during this time.

Are you sure you want to save? (y/n) y

Config file 'startup-config' created successfully .

Configuration Saved!
```

5. Enable the second cluster port, e2a, on each node:

```
network port modify
```

Show example

```
cluster::*> network port modify -node node1 -port e2a -up-admin true
cluster::*> **network port modify -node node2 -port e2a -up-admin
true**
```

6. Revert clus2 that is associated with port e2a:

```
network interface revert
```

The LIF might revert automatically, depending on your version of ONTAP software.

Show example

```
cluster::*> network interface revert -vserver Cluster -lif n1_clus2
cluster::*> network interface revert -vserver Cluster -lif n2_clus2
```

7. Verify that the LIF is now home (true) on both nodes:

```
network interface show -role cluster
```

Show example

```
cluster::*> network interface show -role cluster
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vs1	clus1	up/up	10.10.10.1/24	node1	e1a	true
	clus2	up/up	10.10.10.2/24	node1	e2a	true
vs2	clus1	up/up	10.10.10.1/24	node2	e1a	true
	clus2	up/up	10.10.10.2/24	node2	e2a	true

8. View the status of the nodes:

```
cluster show
```

Show example

```
cluster::> cluster show
```

Node	Health	Eligibility
node1	true	true
node2	true	true

9. Repeat the previous steps to install the FASTPATH software on the other switch, cs1.
10. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Install a Reference Configuration File on a CN1610 switch

Follow this procedure to install a Reference Configuration File (RCF).

Before installing an RCF, you must first migrate the cluster LIFs away from switch cs2. After the RCF is installed and validated, the LIFs can be migrated back.

Review requirements

What you'll need

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs and no defective cluster network interface cards (NICs) or similar issues).
- Fully functional port connections on the cluster switch.
- All cluster ports set up.
- All cluster logical interfaces (LIFs) set up.
- A successful communication path: The ONTAP (privilege: advanced) `cluster ping-cluster -node node1` command must indicate that larger than PMTU communication is successful on all paths.
- A supported version of RCF and ONTAP.

Make sure you consult the switch compatibility table on the [NetApp CN1601 and CN1610 Switches](#) page for the supported RCF and ONTAP versions.

Install the RCF

The following procedure uses the clustered Data ONTAP 8.2 syntax. As a result, the cluster Vserver, LIF names, and CLI output are different than those in Data ONTAP 8.3.

There can be command dependencies between command syntax in the RCF and FASTPATH versions.



In RCF version 1.2, support for Telnet has been explicitly disabled because of security concerns. To avoid connectivity issues while installing RCF 1.2, verify that Secure Shell (SSH) is enabled. The [NetApp CN1610 Switch Administrator's Guide](#) has more information about SSH.

About the examples

The examples in this procedure use the following switch and node nomenclature:

- The two NetApp switches are cs1 and cs2.
- The two cluster LIFs are clus1 and clus2.
- The Vservers are vs1 and vs2.
- The `cluster: :*>` prompt indicates the name of the cluster.
- The cluster ports on each node are named e1a and e2a.

[Hardware Universe](#) has more information about the actual cluster ports that are supported on your platform.

- The supported Inter-Switch Links (ISLs) are ports 0/13 through 0/16.
- The supported node connections are ports 0/1 through 0/12.
- A supported version of FASTPATH, RCF, and ONTAP.

Make sure you consult the switch compatibility table on the [NetApp CN1601 and CN1610 Switches](#) page for the supported FASTPATH, RCF, and ONTAP versions.

Step 1: Migrate cluster

1. Save your current switch configuration information:

```
write memory
```

Show example

The following example shows the current switch configuration being saved to the startup configuration (`startup-config`) file on switch cs2:

```
(cs2) # write memory
This operation may take a few minutes.
Management interfaces will not be available during this time.

Are you sure you want to save? (y/n) y

Config file 'startup-config' created successfully.

Configuration Saved!
```

2. On the console of each node, migrate clus2 to port e1a:

```
network interface migrate
```

Show example

```
cluster::*> network interface migrate -vserver vs1 -lif clus2
-source-node node1 -destnode node1 -dest-port e1a

cluster::*> network interface migrate -vserver vs2 -lif clus2
-source-node node2 -destnode node2 -dest-port e1a
```

3. On the console of each node, verify that the migration occurred:

```
network interface show -role cluster
```

Show example

The following example shows that clus2 has migrated to port e1a on both nodes:

```
cluster::*> network port show -role cluster
      clus1      up/up      10.10.10.1/16   node2   e1a      true
      clus2      up/up      10.10.10.2/16   node2   e1a
false
```

4. Shut down port e2a on both nodes:

```
network port modify
```

Show example

The following example shows port e2a being shut down on both nodes:

```
cluster::*> network port modify -node node1 -port e2a -up-admin
false
cluster::*> network port modify -node node2 -port e2a -up-admin
false
```

5. Verify that port e2a is shut down on both nodes:

```
network port show
```

Show example

```
cluster::*> network port show -role cluster
```

(Mbps)					Auto-Negot	Duplex	Speed
Node	Port	Role	Link	MTU	Admin/Oper	Admin/Oper	Admin/Oper
-----	-----	-----	----	-----	-----	-----	-----
node1							
	e1a	cluster	up	9000	true/true	full/full	auto/10000
	e2a	cluster	down	9000	true/true	full/full	auto/10000
node2							
	e1a	cluster	up	9000	true/true	full/full	auto/10000
	e2a	cluster	down	9000	true/true	full/full	auto/10000

6. Shut down the ISL ports on cs1, the active NetApp switch.

Show example

```
(cs1) # configure
(cs1) (config) # interface 0/13-0/16
(cs1) (interface 0/13-0/16) # shutdown
(cs1) (interface 0/13-0/16) # exit
(cs1) (config) # exit
```

Step 2: Install RCF

1. Copy the RCF to the switch.



You must set the `.scr` extension as part of the file name before invoking the script. This extension is the extension for the FASTPATH operating system.

The switch will validate the script automatically as it is downloaded to the switch, and the output will go to the console.

Show example

```
(cs2) # copy tftp://10.10.0.1/CN1610_CS_RCF_v1.1.txt nvram:script
CN1610_CS_RCF_v1.1.scr

[the script is now displayed line by line]
Configuration script validated.
File transfer operation completed successfully.
```

2. Verify that the script was downloaded and saved with the file name that you gave it.

Show example

```
(cs2) # script list
Configuration Script Name          Size(Bytes)
-----
running-config.scr                6960
CN1610_CS_RCF_v1.1.scr            2199

2 configuration script(s) found.
6038 Kbytes free.
```

3. Validate the script.



The script is validated during the download to verify that each line is a valid switch command line.

Show example

```
(cs2) # script validate CN1610_CS_RCF_v1.1.scr
[the script is now displayed line by line]
Configuration script 'CN1610_CS_RCF_v1.1.scr' validated.
```

4. Apply the script to the switch.

Show example

```
(cs2) #script apply CN1610_CS_RCF_v1.1.scr  
  
Are you sure you want to apply the configuration script? (y/n) y  
[the script is now displayed line by line]...  
  
Configuration script 'CN1610_CS_RCF_v1.1.scr' applied.
```

5. Verify that your changes have been implemented on the switch.

```
(cs2) # show running-config
```

The example displays the `running-config` file on the switch. You must compare the file to the RCF to verify that the parameters that you set are as you expect.

6. Save the changes.
7. Set the `running-config` file to be the standard one.

Show example

```
(cs2) # write memory  
This operation may take a few minutes.  
Management interfaces will not be available during this time.  
  
Are you sure you want to save? (y/n) y  
  
Config file 'startup-config' created successfully.
```

8. Reboot the switch and verify that the `running-config` file is correct.

After the reboot completes, you must log in, view the `running-config` file, and then look for the description on interface 3/64, which is the version label for the RCF.

Show example

```
(cs2) # reload

The system has unsaved changes.
Would you like to save them now? (y/n) y

Config file 'startup-config' created successfully.
Configuration Saved!
System will now restart!
```

9. Bring up the ISL ports on cs1, the active switch.

Show example

```
(cs1) # configure
(cs1) (config)# interface 0/13-0/16
(cs1) (Interface 0/13-0/16)# no shutdown
(cs1) (Interface 0/13-0/16)# exit
(cs1) (config)# exit
```

10. Verify that the ISLs are operational:

```
show port-channel 3/1
```

The Link State field should indicate Up.

Show example

```
(cs2) # show port-channel 3/1

Local Interface..... 3/1
Channel Name..... ISL-LAG
Link State..... Up
Admin Mode..... Enabled
Type..... Static
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr      Device/      Port      Port
Ports   Timeout     Speed     Active
-----
0/13    actor/long   10G Full   True
        partner/long
0/14    actor/long   10G Full   True
        partner/long
0/15    actor/long   10G Full   True
        partner/long
0/16    actor/long   10G Full   True
        partner/long
```

11. Bring up cluster port e2a on both nodes:

```
network port modify
```

Show example

The following example shows port e2a being brought up on node1 and node2:

```
cluster::*> network port modify -node node1 -port e2a -up-admin true
cluster::*> network port modify -node node2 -port e2a -up-admin true
```

Step 3: Validate installation

1. Verify that port e2a is up on both nodes:

```
network port show -role cluster
```

Show example

```
cluster::*> network port show -role cluster
```

Node	Port	Role	Link	MTU	Auto-Negot Admin/Oper	Duplex Admin/Oper	Speed (Mbps) Admin/Oper
-----	-----	-----	-----	-----	-----	-----	-----
node1							
	e1a	cluster	up	9000	true/true	full/full	auto/10000
	e2a	cluster	up	9000	true/true	full/full	auto/10000
node2							
	e1a	cluster	up	9000	true/true	full/full	auto/10000
	e2a	cluster	up	9000	true/true	full/full	auto/10000

2. On both nodes, revert clus2 that is associated with port e2a:

```
network interface revert
```

The LIF might revert automatically, depending on your version of ONTAP.

Show example

```
cluster::*> network interface revert -vserver node1 -lif clus2  
cluster::*> network interface revert -vserver node2 -lif clus2
```

3. Verify that the LIF is now home (true) on both nodes:

```
network interface show -role cluster
```

Show example

```
cluster::*> network interface show -role cluster
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
-----	-----	-----	-----	-----	-----	-----
vs1						
	clus1	up/up	10.10.10.1/24	node1	e1a	true
	clus2	up/up	10.10.10.2/24	node1	e2a	true
vs2						
	clus1	up/up	10.10.10.1/24	node2	e1a	true
	clus2	up/up	10.10.10.2/24	node2	e2a	true

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs and no defective cluster network interface cards (NICs) or similar issues).
- Fully functional port connections on the cluster switch.
- All cluster ports set up.
- All cluster logical interfaces (LIFs) set up (must not have been migrated).
- A successful communication path: The ONTAP (privilege: advanced) `cluster ping-cluster -node node1` command must indicate that larger than PMTU communication is successful on all paths.
- A supported version of FASTPATH, RCF, and ONTAP.

Make sure you consult the switch compatibility table on the [NetApp CN1601 and CN1610 Switches](#) page for the supported FASTPATH, RCF, and ONTAP versions.

Install the FASTPATH software

The following procedure uses the clustered Data ONTAP 8.2 syntax. As a result, the cluster Vserver, LIF names, and CLI output are different than those in Data ONTAP 8.3.

There can be command dependencies between command syntax in the RCF and FASTPATH versions.



In RCF version 1.2, support for Telnet has been explicitly disabled because of security concerns. To avoid connectivity issues while installing RCF 1.2, verify that Secure Shell (SSH) is enabled. The [NetApp CN1610 Switch Administrator's Guide](#) has more information about SSH.

About the examples

The examples in this procedure use the following switch and node nomenclature:

- The two NetApp switch names are cs1 and cs2.
- The cluster logical interface (LIF) names are node1_clus1 and node1_clus2 for node1, and node2_clus1 and node2_clus2 for node2. (You can have up to 24 nodes in a cluster.)
- The storage virtual machine (SVM) name is Cluster.
- The `cluster1::*>` prompt indicates the name of the cluster.
- The cluster ports on each node are named e0a and e0b.

[Hardware Universe](#) has more information about the actual cluster ports that are supported on your platform.

- The supported Inter-Switch Links (ISLs) are ports 0/13 through 0/16.
- The supported node connections are ports 0/1 through 0/12.

Step 1: Migrate cluster

1. Display information about the network ports on the cluster:

```
network port show -ipspace cluster
```

Show example

The following example shows the type of output from the command:

```
cluster1::> network port show -ipspace cluster
```

						Speed
(Mbps)						
Node	Port	IPspace	Broadcast Domain	Link	MTU	
Admin/Oper						
-----	-----	-----	-----	-----	-----	-----
node1						
	e0a	Cluster	Cluster	up	9000	
auto/10000						
	e0b	Cluster	Cluster	up	9000	
auto/10000						
node2						
	e0a	Cluster	Cluster	up	9000	
auto/10000						
	e0b	Cluster	Cluster	up	9000	
auto/10000						

4 entries were displayed.

2. Display information about the LIFs on the cluster:

```
network interface show -role cluster
```

Show example

The following example shows the logical interfaces on the cluster. In this example the `-role` parameter displays information about the LIFs that are associated with cluster ports:

```
cluster1::> network interface show -role cluster
(network interface show)

Current Is      Logical      Status      Network      Current
Vserver Is      Interface   Admin/Oper  Address/Mask  Node
Port   Home
-----
Cluster
e0a    true      node1_clus1 up/up       10.254.66.82/16  node1
e0b    true      node1_clus2 up/up       10.254.206.128/16 node1
e0a    true      node2_clus1 up/up       10.254.48.152/16  node2
e0b    true      node2_clus2 up/up       10.254.42.74/16   node2
4 entries were displayed.
```

3. On each respective node, using a node management LIF, migrate `node1_clus2` to `e0a` on `node1` and `node2_clus2` to `e0a` on `node2`:

```
network interface migrate
```

You must enter the commands on the controller consoles that own the respective cluster LIFs.

Show example

```
cluster1::> network interface migrate -vserver Cluster -lif
node1_clus2 -destination-node node1 -destination-port e0a
cluster1::> network interface migrate -vserver Cluster -lif
node2_clus2 -destination-node node2 -destination-port e0a
```



For this command, the name of the cluster is case-sensitive and the command should be run on each node. It is not possible to run this command in the general cluster LIF.

4. Verify that the migration took place by using the `network interface show` command on a node.

Show example

The following example shows that clus2 has migrated to port e0a on nodes node1 and node2:

```
cluster1::> **network interface show -role cluster**
          Logical      Status      Network      Current
Current Is
Vserver   Interface  Admin/Oper  Address/Mask  Node
Port      Home
-----
-----
Cluster
          node1_clus1  up/up      10.254.66.82/16  node1
e0a       true
          node1_clus2  up/up      10.254.206.128/16 node1
e0a       false
          node2_clus1  up/up      10.254.48.152/16  node2
e0a       true
          node2_clus2  up/up      10.254.42.74/16  node2
e0a       false
4 entries were displayed.
```

5. Change the privilege level to advanced, entering y when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (*>) appears.

6. Shut down cluster port e0b on both nodes:

```
network port modify -node node_name -port port_name -up-admin false
```

You must enter the commands on the controller consoles that own the respective cluster LIFs.

Show example

The following example shows the commands to shut down port e0b on all nodes:

```
cluster1::*> network port modify -node node1 -port e0b -up-admin
false
cluster1::*> network port modify -node node2 -port e0b -up-admin
false
```

7. Verify that port e0b is shut down on both nodes:

network port show

Show example

```
cluster1::*> network port show -role cluster
```

(Mbps)	Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed
	node1	e0a	Cluster	Cluster	up	9000	
		auto/10000					
		e0b	Cluster	Cluster	down	9000	
		auto/10000					
	node2	e0a	Cluster	Cluster	up	9000	
		auto/10000					
		e0b	Cluster	Cluster	down	9000	
		auto/10000					

4 entries were displayed.

8. Shut down the Inter-Switch Link (ISL) ports on cs1.

Show example

```
(cs1) #configure
(cs1) (Config)#interface 0/13-0/16
(cs1) (Interface 0/13-0/16)#shutdown
(cs1) (Interface 0/13-0/16)#exit
(cs1) (Config)#exit
```

9. Back up the current active image on cs2.

Show example

```
(cs2) # show bootvar
```

```
Image Descriptions
```

```
active :
```

```
backup :
```

```
Images currently available on Flash
```

```
-----  
unit      active      backup      current-active      next-active  
-----  
1         1.1.0.5      1.1.0.3      1.1.0.5              1.1.0.5
```

```
(cs2) # copy active backup
```

```
Copying active to backup
```

```
Copy operation successful
```

Step 2: Install the FASTPATH software and RCF

1. Verify the running version of the FASTPATH software.

Show example

```
(cs2) # show version

Switch: 1

System Description..... NetApp CN1610,
1.1.0.5, Linux
                               2.6.21.7
Machine Type..... NetApp CN1610
Machine Model..... CN1610
Serial Number..... 20211200106
Burned In MAC Address..... 00:A0:98:21:83:69
Software Version..... 1.1.0.5
Operating System..... Linux 2.6.21.7
Network Processing Device..... BCM56820_B0
Part Number..... 111-00893

--More-- or (q)uit

Additional Packages..... FASTPATH QOS
                               FASTPATH IPv6

Management
```

2. Download the image file to the switch.

Copying the image file to the active image means that when you reboot, that image establishes the running FASTPATH version. The previous image remains available as a backup.

Show example

```
(cs2) #copy
sftp://root@10.22.201.50//tftpboot/NetApp_CN1610_1.2.0.7.stk active
Remote Password:*****

Mode..... SFTP
Set Server IP..... 10.22.201.50
Path..... /tftpboot/
Filename.....
NetApp_CN1610_1.2.0.7.stk
Data Type..... Code
Destination Filename..... active

Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y
SFTP Code transfer starting...

File transfer operation completed successfully.
```

3. Confirm the current and next-active boot image versions:

```
show bootvar
```

Show example

```
(cs2) #show bootvar

Image Descriptions

active :
backup :

Images currently available on Flash

-----
unit      active      backup      current-active      next-active
-----
1         1.1.0.8      1.1.0.8      1.1.0.8              1.2.0.7
```

4. Install the compatible RCF for the new image version to the switch.

If the RCF version is already correct, bring up the ISL ports.

Show example

```
(cs2) #copy tftp://10.22.201.50//CN1610_CS_RCF_v1.2.txt nvram:script
CN1610_CS_RCF_v1.2.scr

Mode..... TFTP
Set Server IP..... 10.22.201.50
Path..... /
Filename.....
CN1610_CS_RCF_v1.2.txt
Data Type..... Config Script
Destination Filename.....
CN1610_CS_RCF_v1.2.scr

File with same name already exists.
WARNING:Continuing with this command will overwrite the existing
file.

Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y

Validating configuration script...
[the script is now displayed line by line]

Configuration script validated.
File transfer operation completed successfully.
```



The `.scr` extension must be set as part of the file name before invoking the script. This extension is for the FASTPATH operating system.

The switch validates the script automatically as it is downloaded to the switch. The output goes to the console.

5. Verify that the script was downloaded and saved to the file name you gave it.

Show example

```
(cs2) #script list

Configuration Script Name          Size(Bytes)
-----
CN1610_CS_RCF_v1.2.scr            2191

1 configuration script(s) found.
2541 Kbytes free.
```

6. Apply the script to the switch.

Show example

```
(cs2) #script apply CN1610_CS_RCF_v1.2.scr

Are you sure you want to apply the configuration script? (y/n) y
[the script is now displayed line by line]...

Configuration script 'CN1610_CS_RCF_v1.2.scr' applied.
```

7. Verify that the changes have been applied to the switch, and then save them:

```
show running-config
```

Show example

```
(cs2) #show running-config
```

8. Save the running configuration so it becomes the startup configuration when you reboot the switch.

Show example

```
(cs2) #write memory
This operation may take a few minutes.
Management interfaces will not be available during this time.

Are you sure you want to save? (y/n) y

Config file 'startup-config' created successfully.

Configuration Saved!
```

9. Reboot the switch.

Show example

```
(cs2) #reload

The system has unsaved changes.
Would you like to save them now? (y/n) y

Config file 'startup-config' created successfully.
Configuration Saved!
System will now restart!
```

Step 3: Validate installation

1. Log in again, and then verify that the switch is running the new version of the FASTPATH software.

Show example

```
(cs2) #show version

Switch: 1

System Description..... NetApp CN1610,
1.2.0.7, Linux
                               3.8.13-4ce360e8
Machine Type..... NetApp CN1610
Machine Model..... CN1610
Serial Number..... 20211200106
Burned In MAC Address..... 00:A0:98:21:83:69
Software Version..... 1.2.0.7
Operating System..... Linux 3.8.13-
4ce360e8
Network Processing Device..... BCM56820_B0
Part Number..... 111-00893
CPLD version..... 0x5

Additional Packages..... FASTPATH QOS
                               FASTPATH IPv6

Management
```

After the reboot completes, you must log in to verify the image version, view the running configuration, and look for the description on interface 3/64, which is the version label for the RCF.

2. Bring up the ISL ports on cs1, the active switch.

Show example

```
(cs1) #configure
(cs1) (Config) #interface 0/13-0/16
(cs1) (Interface 0/13-0/16) #no shutdown
(cs1) (Interface 0/13-0/16) #exit
(cs1) (Config) #exit
```

3. Verify that the ISLs are operational:

```
show port-channel 3/1
```

The Link State field should indicate Up.

Show example

```
(cs1) #show port-channel 3/1

Local Interface..... 3/1
Channel Name..... ISL-LAG
Link State..... Up
Admin Mode..... Enabled
Type..... Static
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr      Device/      Port      Port
Ports   Timeout     Speed     Active
-----
0/13    actor/long   10G Full   True
        partner/long
0/14    actor/long   10G Full   True
        partner/long
0/15    actor/long   10G Full   False
        partner/long
0/16    actor/long   10G Full   True
        partner/long
```

4. Bring up cluster port e0b on all nodes:

```
network port modify
```

You must enter the commands on the controller consoles that own the respective cluster LIFs.

Show example

The following example shows port e0b being brought up on node1 and node2:

```
cluster1::*> network port modify -node node1 -port e0b -up-admin
true
cluster1::*> network port modify -node node2 -port e0b -up-admin
true
```

5. Verify that the port e0b is up on all nodes:

```
network port show -ip-space cluster
```

Show example

```
cluster1::*> network port show -ip-space cluster
```

(Mbps)					Speed
Node	Port	IPspace	Broadcast Domain	Link	MTU
Admin/Oper					

node1					
	e0a	Cluster	Cluster	up	9000
auto/10000					
	e0b	Cluster	Cluster	up	9000
auto/10000					
node2					
	e0a	Cluster	Cluster	up	9000
auto/10000					
	e0b	Cluster	Cluster	up	9000
auto/10000					

4 entries were displayed.

6. Verify that the LIF is now home (`true`) on both nodes:

```
network interface show -role cluster
```

Show example

```
cluster1::*> network interface show -role cluster
```

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
e0a	node1_clus1	up/up	169.254.66.82/16	node1
	true			
e0b	node1_clus2	up/up	169.254.206.128/16	node1
	true			
e0a	node2_clus1	up/up	169.254.48.152/16	node2
	true			
e0b	node2_clus2	up/up	169.254.42.74/16	node2
	true			

4 entries were displayed.

7. Show the status of the node members:

```
cluster show
```

Show example

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon

node1	true	true	false
node2	true	true	false

2 entries were displayed.

8. Return to the admin privilege level:

```
set -privilege admin
```

9. Repeat the previous steps to install the FASTPATH software and RCF on the other switch, cs1.

Configure the hardware for the NetApp CN1610 switch

To configure the switch hardware and software for your cluster environment, refer to the

Migrate switches

Migrate from a switchless cluster environment to a switched NetApp CN1610 cluster environment

If you have an existing two-node switchless cluster environment, you can migrate to a two-node switched cluster environment using CN1610 cluster network switches that enables you to scale beyond two nodes.

Review requirements

What you'll need

For a two-node switchless configuration, ensure that:

- The two-node switchless configuration is properly set up and functioning.
- The nodes are running ONTAP 8.2 or later.
- All cluster ports are in the `up` state.
- All cluster logical interfaces (LIFs) are in the `up` state and on their home ports.

For the CN1610 cluster switch configuration:

- The CN1610 cluster switch infrastructure are fully functional on both switches.
- Both switches have management network connectivity.
- There is console access to the cluster switches.
- CN1610 node-to-node switch and switch-to-switch connections use twinax or fiber cables.

The [Hardware Universe](#) contains more information about cabling.

- Inter-Switch Link (ISL) cables are connected to ports 13 through 16 on both CN1610 switches.
- Initial customization of both the CN1610 switches are completed.

Any previous site customization, such as SMTP, SNMP, and SSH should be copied to the new switches.

Related information

- [Hardware Universe](#)
- [NetApp CN1601 and CN1610 description page](#)
- [CN1601 and CN1610 Switch Setup and Configuration Guide](#)
- [NetApp KB Article 1010449: How to suppress automatic case creation during scheduled maintenance windows](#)

Migrate the switches

About the examples

The examples in this procedure use the following cluster switch and node nomenclature:

- The names of the CN1610 switches are `cs1` and `cs2`.

- The names of the LIFs are clus1 and clus2.
- The names of the nodes are node1 and node2.
- The `cluster::*>` prompt indicates the name of the cluster.
- The cluster ports used in this procedure are e1a and e2a.

The [Hardware Universe](#) contains the latest information about the actual cluster ports for your platforms.

Step 1: Prepare for migration

1. Change the privilege level to advanced, entering `y` when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (`*>`) appears.

2. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=xh
```

`x` is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

Show example

The following command suppresses automatic case creation for two hours:

```
cluster::*> system node autosupport invoke -node * -type all  
-message MAINT=2h
```

Step 2: Configure ports

1. Disable all of the node-facing ports (not ISL ports) on both the new cluster switches cs1 and cs2.

You must not disable the ISL ports.

Show example

The following example shows that node-facing ports 1 through 12 are disabled on switch cs1:

```
(cs1)> enable
(cs1)# configure
(cs1) (Config)# interface 0/1-0/12
(cs1) (Interface 0/1-0/12)# shutdown
(cs1) (Interface 0/1-0/12)# exit
(cs1) (Config)# exit
```

The following example shows that node-facing ports 1 through 12 are disabled on switch cs2:

```
(c2)> enable
(cs2)# configure
(cs2) (Config)# interface 0/1-0/12
(cs2) (Interface 0/1-0/12)# shutdown
(cs2) (Interface 0/1-0/12)# exit
(cs2) (Config)# exit
```

2. Verify that the ISL and the physical ports on the ISL between the two CN1610 cluster switches cs1 and cs2 are up:

```
show port-channel
```

Show example

The following example shows that the ISL ports are up on switch cs1:

```
(cs1)# show port-channel 3/1
Local Interface..... 3/1
Channel Name..... ISL-LAG
Link State..... Up
Admin Mode..... Enabled
Type..... Static
Load Balance Option..... 7
(Enhanced hashing mode)

Mbr      Device/      Port      Port
Ports   Timeout     Speed     Active
-----
0/13    actor/long   10G Full  True
        partner/long
0/14    actor/long   10G Full  True
        partner/long
0/15    actor/long   10G Full  True
        partner/long
0/16    actor/long   10G Full  True
        partner/long
```

The following example shows that the ISL ports are up on switch cs2:


```
(cs2)# show port-channel 3/1
Local Interface..... 3/1
Channel Name..... ISL-LAG
Link State..... Up
Admin Mode..... Enabled
Type..... Static
Load Balance Option..... 7
(Enhanced hashing mode)
```

Mbr	Device/	Port	Port
Ports	Timeout	Speed	Active
-----	-----	-----	-----
0/13	actor/long	10G Full	True
	partner/long		
0/14	actor/long	10G Full	True
	partner/long		
0/15	actor/long	10G Full	True
	partner/long		
0/16	actor/long	10G Full	True
	partner/long		

3. Display the list of neighboring devices:

```
show isdp neighbors
```

This command provides information about the devices that are connected to the system.

Show example

The following example lists the neighboring devices on switch cs1:

```
(cs1)# show isdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route
Bridge,
                S - Switch, H - Host, I - IGMP, r - Repeater
Device ID      Intf      Holdtime  Capability  Platform
Port ID
-----
cs2            0/13     11        S           CN1610
0/13
cs2            0/14     11        S           CN1610
0/14
cs2            0/15     11        S           CN1610
0/15
cs2            0/16     11        S           CN1610
0/16
```

The following example lists the neighboring devices on switch cs2:

```
(cs2)# show isdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route
Bridge,
                S - Switch, H - Host, I - IGMP, r - Repeater
Device ID      Intf      Holdtime  Capability  Platform
Port ID
-----
cs1            0/13     11        S           CN1610
0/13
cs1            0/14     11        S           CN1610
0/14
cs1            0/15     11        S           CN1610
0/15
cs1            0/16     11        S           CN1610
0/16
```

4. Display the list of cluster ports:

```
network port show
```

Show example

The following example shows the available cluster ports:

```
cluster::*> network port show -ipspace Cluster
```

```
Node: node1
```

```
Ignore
```

Health						Speed(Mbps)	Health
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
e0a	Cluster	Cluster		up	9000	auto/10000	healthy
e0b	Cluster	Cluster		up	9000	auto/10000	healthy
e0c	Cluster	Cluster		up	9000	auto/10000	healthy
e0d	Cluster	Cluster		up	9000	auto/10000	healthy
e4a	Cluster	Cluster		up	9000	auto/10000	healthy
e4b	Cluster	Cluster		up	9000	auto/10000	healthy

```
Node: node2
```

```
Ignore
```

Health						Speed(Mbps)	Health
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
-----	-----	-----	-----	-----	-----	-----	-----
e0a	Cluster	Cluster		up	9000	auto/10000	healthy
e0b	Cluster	Cluster		up	9000	auto/10000	healthy
e0c	Cluster	Cluster		up	9000	auto/10000	healthy
e0d	Cluster	Cluster		up	9000	auto/10000	healthy
e4a	Cluster	Cluster		up	9000	auto/10000	healthy
e4b	Cluster	Cluster		up	9000	auto/10000	healthy

```
12 entries were displayed.
```

5. Verify that each cluster port is connected to the corresponding port on its partner cluster node:

```
run * cdpd show-neighbors
```

Show example

The following example shows that cluster ports e1a and e2a are connected to the same port on their cluster partner node:

```
cluster::*> run * cdpd show-neighbors
2 entries were acted on.

Node: node1
Local Remote Remote Remote Remote Hold
Remote
Port Device Interface Platform Time
Capability
-----
-----
e1a node2 e1a FAS3270 137
H
e2a node2 e2a FAS3270 137
H

Node: node2
Local Remote Remote Remote Remote Hold
Remote
Port Device Interface Platform Time
Capability
-----
-----
e1a node1 e1a FAS3270 161
H
e2a node1 e2a FAS3270 161
H
```

6. Verify that all of the cluster LIFs are up and operational:

```
network interface show -vserver Cluster
```

Each cluster LIF should display `true` in the “Is Home” column.

Show example

```
cluster::*> network interface show -vserver Cluster
          Logical   Status   Network   Current
Current Is
Vserver   Interface  Admin/Oper Address/Mask Node      Port
Home
-----
node1
true      clus1      up/up    10.10.10.1/16 node1     e1a
true      clus2      up/up    10.10.10.2/16 node1     e2a
node2
true      clus1      up/up    10.10.11.1/16 node2     e1a
true      clus2      up/up    10.10.11.2/16 node2     e2a

4 entries were displayed.
```



The following modification and migration commands in steps 10 through 13 must be done from the local node.

7. Verify that all cluster ports are up:

```
network port show -ipSPACE Cluster
```

Show example

```
cluster::*> network port show -ipSpace Cluster
```

(Mbps)					Auto-Negot	Duplex	Speed
Node	Port	Role	Link	MTU	Admin/Oper	Admin/Oper	
Admin/Oper							

node1							
	e1a	clus1	up	9000	true/true	full/full	
auto/10000							
	e2a	clus2	up	9000	true/true	full/full	
auto/10000							
node2							
	e1a	clus1	up	9000	true/true	full/full	
auto/10000							
	e2a	clus2	up	9000	true/true	full/full	
auto/10000							

4 entries were displayed.

8. Set the `-auto-revert` parameter to `false` on cluster LIFs `clus1` and `clus2` on both nodes:

```
network interface modify
```

Show example

```
cluster::*> network interface modify -vserver node1 -lif clus1 -auto-revert false
cluster::*> network interface modify -vserver node1 -lif clus2 -auto-revert false
cluster::*> network interface modify -vserver node2 -lif clus1 -auto-revert false
cluster::*> network interface modify -vserver node2 -lif clus2 -auto-revert false
```



For release 8.3 and later, use the following command: `network interface modify -vserver Cluster -lif * -auto-revert false`

9. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```



```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

10. Migrate clus1 to port e2a on the console of each node:

```
network interface migrate
```

Show example

The following example shows the process for migrating clus1 to port e2a on node1 and node2:

```

cluster::*> network interface migrate -vserver node1 -lif clus1
-source-node node1 -dest-node node1 -dest-port e2a
cluster::*> network interface migrate -vserver node2 -lif clus1
-source-node node2 -dest-node node2 -dest-port e2a

```



For release 8.3 and later, use the following command: `network interface migrate -vserver Cluster -lif clus1 -destination-node node1 -destination-port e2a`

11. Verify that the migration took place:

```
network interface show -vserver Cluster
```

Show example

The following example verifies that clus1 is migrated to port e2a on node1 and node2:

```
cluster::*> network interface show -vserver Cluster
      Logical   Status   Network   Current
Current Is
Vserver   Interface  Admin/Oper Address/Mask  Node      Port
Home
-----
-----
node1
false     clus1      up/up     10.10.10.1/16 node1     e2a
true      clus2      up/up     10.10.10.2/16 node1     e2a
node2
false     clus1      up/up     10.10.11.1/16 node2     e2a
true      clus2      up/up     10.10.11.2/16 node2     e2a

4 entries were displayed.
```

12. Shut down cluster port e1a on both nodes:

```
network port modify
```

Show example

The following example shows how to shut down the port e1a on node1 and node2:

```
cluster::*> network port modify -node node1 -port e1a -up-admin
false
cluster::*> network port modify -node node2 -port e1a -up-admin
false
```

13. Verify the port status:

```
network port show
```

Show example

The following example shows that port e1a is down on node1 and node2:

```
cluster::*> network port show -role cluster
                                     Auto-Negot Duplex      Speed
(Mbps)
Node  Port  Role      Link  MTU Admin/Oper Admin/Oper
Admin/Oper
-----
node1
      e1a  clus1    down  9000 true/true  full/full
      auto/10000
      e2a  clus2    up    9000 true/true  full/full
      auto/10000
node2
      e1a  clus1    down  9000 true/true  full/full
      auto/10000
      e2a  clus2    up    9000 true/true  full/full
      auto/10000

4 entries were displayed.
```

14. Disconnect the cable from cluster port e1a on node1, and then connect e1a to port 1 on cluster switch cs1, using the appropriate cabling supported by the CN1610 switches.

The [Hardware Universe](#) contains more information about cabling.

15. Disconnect the cable from cluster port e1a on node2, and then connect e1a to port 2 on cluster switch cs1, using the appropriate cabling supported by the CN1610 switches.
16. Enable all of the node-facing ports on cluster switch cs1.

Show example

The following example shows that ports 1 through 12 are enabled on switch cs1:

```
(cs1)# configure
(cs1)(Config)# interface 0/1-0/12
(cs1)(Interface 0/1-0/12)# no shutdown
(cs1)(Interface 0/1-0/12)# exit
(cs1)(Config)# exit
```

17. Enable the first cluster port e1a on each node:

```
network port modify
```

Show example

The following example shows how to enable the port e1a on node1 and node2:

```
cluster::*> network port modify -node node1 -port e1a -up-admin true
cluster::*> network port modify -node node2 -port e1a -up-admin true
```

18. Verify that all of the cluster ports are up:

```
network port show -ipSpace Cluster
```

Show example

The following example shows that all of the cluster ports are up on node1 and node2:

```
cluster::*> network port show -ipSpace Cluster
                                     Auto-Negot Duplex      Speed
(Mbps)
Node  Port  Role      Link  MTU Admin/Oper Admin/Oper
Admin/Oper
-----
-----
node1
   e1a  clus1    up    9000 true/true  full/full
auto/10000
   e2a  clus2    up    9000 true/true  full/full
auto/10000
node2
   e1a  clus1    up    9000 true/true  full/full
auto/10000
   e2a  clus2    up    9000 true/true  full/full
auto/10000

4 entries were displayed.
```

19. Revert clus1 (which was previously migrated) to e1a on both nodes:

```
network interface revert
```

Show example

The following example shows how to revert clus1 to the port e1a on node1 and node2:

```
cluster::*> network interface revert -vserver node1 -lif clus1
cluster::*> network interface revert -vserver node2 -lif clus1
```



For release 8.3 and later, use the following command: `network interface revert -vserver Cluster -lif <nodename_clus<N>>`

20. Verify that all of the cluster LIFs are up, operational, and display as true in the "Is Home" column:

```
network interface show -vserver Cluster
```

Show example

The following example shows that all of the LIFs are up on node1 and node2 and that the "Is Home" column results are true:

```
cluster::*> network interface show -vserver Cluster
      Logical      Status      Network      Current
Current Is
Vserver  Interface  Admin/Oper  Address/Mask  Node      Port
Home
-----
node1
      clus1      up/up      10.10.10.1/16  node1      e1a
true
      clus2      up/up      10.10.10.2/16  node1      e2a
true
node2
      clus1      up/up      10.10.11.1/16  node2      e1a
true
      clus2      up/up      10.10.11.2/16  node2      e2a
true

4 entries were displayed.
```

21. Display information about the status of the nodes in the cluster:

```
cluster show
```

Show example

The following example displays information about the health and eligibility of the nodes in the cluster:

```
cluster::*> cluster show
Node           Health Eligibility  Epsilon
-----
node1          true   true        false
node2          true   true        false
```

22. Migrate clus2 to port e1a on the console of each node:

```
network interface migrate
```

Show example

The following example shows the process for migrating clus2 to port e1a on node1 and node2:

```
cluster::*> network interface migrate -vserver node1 -lif clus2
-source-node node1 -dest-node node1 -dest-port e1a
cluster::*> network interface migrate -vserver node2 -lif clus2
-source-node node2 -dest-node node2 -dest-port e1a
```



For release 8.3 and later, use the following command: `network interface migrate -vserver Cluster -lif node1_clus2 -dest-node node1 -dest-port e1a`

23. Verify that the migration took place:

```
network interface show -vserver Cluster
```

Show example

The following example verifies that clus2 is migrated to port e1a on node1 and node2:

```
cluster::*> network interface show -vserver Cluster
          Logical   Status   Network   Current
Current Is
Vserver   Interface  Admin/Oper Address/Mask  Node   Port
Home
-----
node1
true      clus1      up/up     10.10.10.1/16  node1  e1a
false     clus2      up/up     10.10.10.2/16  node1  e1a
node2
true      clus1      up/up     10.10.11.1/16  node2  e1a
false     clus2      up/up     10.10.11.2/16  node2  e1a

4 entries were displayed.
```

24. Shut down cluster port e2a on both nodes:

```
network port modify
```

Show example

The following example shows how to shut down the port e2a on node1 and node2:

```
cluster::*> network port modify -node node1 -port e2a -up-admin
false
cluster::*> network port modify -node node2 -port e2a -up-admin
false
```

25. Verify the port status:

```
network port show
```

Show example

The following example shows that port e2a is down on node1 and node2:

```
cluster::*> network port show -role cluster

                                     Auto-Negot  Duplex      Speed
(Mbps)
Node  Port  Role      Link  MTU  Admin/Oper  Admin/Oper
Admin/Oper
-----
node1
      e1a   clus1    up    9000  true/true   full/full
auto/10000
      e2a   clus2    down  9000  true/true   full/full
auto/10000
node2
      e1a   clus1    up    9000  true/true   full/full
auto/10000
      e2a   clus2    down  9000  true/true   full/full
auto/10000

4 entries were displayed.
```

26. Disconnect the cable from cluster port e2a on node1, and then connect e2a to port 1 on cluster switch cs2, using the appropriate cabling supported by the CN1610 switches.
27. Disconnect the cable from cluster port e2a on node2, and then connect e2a to port 2 on cluster switch cs2, using the appropriate cabling supported by the CN1610 switches.
28. Enable all of the node-facing ports on cluster switch cs2.

Show example

The following example shows that ports 1 through 12 are enabled on switch cs2:

```
(cs2)# configure
(cs2) (Config)# interface 0/1-0/12
(cs2) (Interface 0/1-0/12)# no shutdown
(cs2) (Interface 0/1-0/12)# exit
(cs2) (Config)# exit
```

29. Enable the second cluster port e2a on each node.

Show example

The following example shows how to enable the port e2a on node1 and node2:

```
cluster::*> network port modify -node node1 -port e2a -up-admin true
cluster::*> network port modify -node node2 -port e2a -up-admin true
```

30. Verify that all of the cluster ports are up:

```
network port show -ipSPACE Cluster
```

Show example

The following example shows that all of the cluster ports are up on node1 and node2:

```
cluster::*> network port show -ipSPACE Cluster
                                     Auto-Negot Duplex      Speed
(Mbps)
Node  Port  Role      Link  MTU Admin/Oper  Admin/Oper
Admin/Oper
-----
node1
      e1a  clus1    up    9000 true/true   full/full
auto/10000
      e2a  clus2    up    9000 true/true   full/full
auto/10000
node2
      e1a  clus1    up    9000 true/true   full/full
auto/10000
      e2a  clus2    up    9000 true/true   full/full
auto/10000

4 entries were displayed.
```

31. Revert clus2 (which was previously migrated) to e2a on both nodes:

```
network interface revert
```

Show example

The following example shows how to revert clus2 to the port e2a on node1 and node2:

```
cluster::*> network interface revert -vserver node1 -lif clus2
cluster::*> network interface revert -vserver node2 -lif clus2
```



For release 8.3 and later, the commands are:

```
cluster::*> network interface revert -vserver Cluster -lif
node1_clus2 and
cluster::*> network interface revert -vserver Cluster -lif
node2_clus2
```

Step 3: Complete the configuration

1. Verify that all of the interfaces display `true` in the "Is Home" column:

```
network interface show -vserver Cluster
```

Show example

The following example shows that all of the LIFs are up on node1 and node2 and that the "Is Home" column results are `true`:

```
cluster::*> network interface show -vserver Cluster
```

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

node1				
e1a	clus1	up/up	10.10.10.1/16	node1
e2a	clus2	up/up	10.10.10.2/16	node1
node2				
e1a	clus1	up/up	10.10.11.1/16	node2
e2a	clus2	up/up	10.10.11.2/16	node2

2. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```
cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

3. Verify that both nodes have two connections to each switch:

```
show isdp neighbors
```

Show example

The following example shows the appropriate results for both switches:

```
(cs1)# show isdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route
Bridge,
                S - Switch, H - Host, I - IGMP, r - Repeater
Device ID      Intf      Holdtime  Capability  Platform
Port ID
-----
node1          0/1       132       H           FAS3270
e1a
node2          0/2       163       H           FAS3270
e1a
cs2            0/13      11        S           CN1610
0/13
cs2            0/14      11        S           CN1610
0/14
cs2            0/15      11        S           CN1610
0/15
cs2            0/16      11        S           CN1610
0/16

(cs2)# show isdp neighbors
Capability Codes: R - Router, T - Trans Bridge, B - Source Route
Bridge,
                S - Switch, H - Host, I - IGMP, r - Repeater
Device ID      Intf      Holdtime  Capability  Platform
Port ID
-----
node1          0/1       132       H           FAS3270
e2a
node2          0/2       163       H           FAS3270
e2a
cs1            0/13      11        S           CN1610
0/13
cs1            0/14      11        S           CN1610
0/14
cs1            0/15      11        S           CN1610
0/15
cs1            0/16      11        S           CN1610
0/16
```

4. Display information about the devices in your configuration:

```
network device discovery show
```

5. Disable the two-node switchless configuration settings on both nodes using the advanced privilege command:

```
network options detect-switchless modify
```

Show example

The following example shows how to disable the switchless configuration settings:

```
cluster::*> network options detect-switchless modify -enabled false
```



For release 9.2 and later, skip this step since the configuration is automatically converted.

6. Verify that the settings are disabled:

```
network options detect-switchless-cluster show
```

Show example

The false output in the following example shows that the configuration settings are disabled:

```
cluster::*> network options detect-switchless-cluster show
Enable Switchless Cluster Detection: false
```



For release 9.2 and later, wait until `Enable Switchless Cluster` is set to false. This can take up to three minutes.

7. Configure clusters `clus1` and `clus2` to auto revert on each node and confirm.

Show example

```
cluster::*> network interface modify -vserver node1 -lif clus1 -auto
-revert true
cluster::*> network interface modify -vserver node1 -lif clus2 -auto
-revert true
cluster::*> network interface modify -vserver node2 -lif clus1 -auto
-revert true
cluster::*> network interface modify -vserver node2 -lif clus2 -auto
-revert true
```



For release 8.3 and later, use the following command: `network interface modify -vserver Cluster -lif * -auto-revert true` to enable auto-revert on all nodes in the cluster.

- Verify the status of the node members in the cluster:

```
cluster show
```

Show example

The following example shows information about the health and eligibility of the nodes in the cluster:

```
cluster::*> cluster show
Node                Health  Eligibility  Epsilon
-----
node1                true    true         false
node2                true    true         false
```

- If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Show example

```
cluster::*> system node autosupport invoke -node * -type all
-messsage MAINT=END
```

- Change the privilege level back to admin:

```
set -privilege admin
```

Replace switches

Replace a NetApp CN1610 cluster switch

Follow these steps to replace a defective NetApp CN1610 switch in a cluster network. This is a non-disruptive procedure (NDU).

Review requirements

What you'll need

Before you perform the switch replacement, the following conditions must exist before you perform the switch replacement in the current environment and on the replacement switch for existing cluster and network infrastructure:

- The existing cluster must be verified as completely functional, with at least one fully connected cluster switch.
- All of the cluster ports must be **up**.
- All of the cluster logical interfaces (LIFs) must be up and must not have been migrated.
- The ONTAP cluster `ping-cluster -node node1` command must indicate that basic connectivity and larger than PMTU communication are successful on all of the paths.

Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when replacing your switch:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport before and after maintenance to disable case creation for the duration of the maintenance. See this Knowledge Base article [SU92: How to suppress automatic case creation during scheduled maintenance windows](#) for further details.
- Enable session logging for any CLI sessions. For instructions on how to enable session logging, review the "Logging Session Output" section in this Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

Replace the switch

About this task

You must execute the command for migrating a cluster LIF from the node where the cluster LIF is hosted.

The examples in this procedure use the following cluster switch and node nomenclature:

- The names of the two CN1610 cluster switches are `cs1` and `cs2`.
- The name of the CN1610 switch that is to be replaced (the defective switch) is `old_cs1`.
- The name of the new CN1610 switch (the replacement switch) is `new_cs1`.
- The name of the partner switch that is not being replaced is `cs2`.

Steps

1. Confirm that the startup configuration file matches the running configuration file. You must save these files locally for use during the replacement.

The configuration commands in the following example are for FASTPATH 1.2.0.7:

Show example

```
(old_cs1) > enable
(old_cs1) # show running-config
(old_cs1) # show startup-config
```

2. Create a copy of the running configuration file.

The command in the following example is for FASTPATH 1.2.0.7:

Show example

```
(old_cs1)# show running-config filename.scr  
Config script created successfully.
```



You can use any file name except `CN1610_CS_RCF_v1.2.scr`. The file name must have the `.scr` extension.

3. Save the running configuration file of the switch to an external host in preparation for the replacement.

Show example

```
(old_cs1)# copy nvram:script filename.scr  
scp://<Username>@<remote_IP_address>/path_to_file/filename.scr
```

4. Verify that the switch and ONTAP versions match in the compatibility matrix. See the [NetApp CN1601 and CN1610 Switches](#) page for details.
5. From the [Software Downloads page](#) on the NetApp Support Site, select NetApp Cluster Switches to download the appropriate RCF and FASTPATH versions.
6. Set up a Trivial File Transfer Protocol (TFTP) server with the FASTPATH, RCF, and saved configuration `.scr` file for use with the new switch.
7. Connect the serial port (the RJ-45 connector labeled “IOIOI” on the right side of the switch) to an available host with terminal emulation.
8. On the host, set the serial terminal connection settings:
 - a. 9600 baud
 - b. 8 data bits
 - c. 1 stop bit
 - d. parity: none
 - e. flow control: none
9. Connect the management port (the RJ-45 wrench port on the left side of the switch) to the same network where your TFTP server is located.
10. Prepare to connect to the network with the TFTP server.

If you are using Dynamic Host Configuration Protocol (DHCP), you do not have to configure an IP address for the switch at this time. The service port is set to use DHCP by default. The network management port is set to none for the IPv4 and IPv6 protocol settings. If your wrench port is connected to a network that has a DHCP server, then the server settings are configured automatically.

To set a static IP address, you should use the `serviceport protocol`, `network protocol`, and `serviceport ip` commands.

Show example

```
(new_cs1) # serviceport ip <ipaddr> <netmask> <gateway>
```

11. Optionally, if the TFTP server is on a laptop, then connect the CN1610 switch to the laptop by using a standard Ethernet cable, and then configure its network port in the same network with an alternate IP address.

You can use the `ping` command to verify the address. If you are unable to establish the connectivity, you should use a nonrouted network, and configure the service port using IP 192.168.x or 172.16.x. You can reconfigure the service port to the production management IP address at a later date.

12. Optionally, verify and install the appropriate versions of the RCF and FASTPATH software for the new switch. If you have verified that the new switch is correctly set up and does not require updates to the RCF and FASTPATH software, you should go to step 13.
 - a. Verify the new switch settings.

Show example

```
(new_cs1) > enable  
(new_cs1) # show version
```

- b. Download the RCF to the new switch.

Show example

```
(new_cs1)# copy tftp://<server_ip_address>/CN1610_CS_RCF_v1.2.txt
nvram:script CN1610_CS_RCF_v1.2.scr
Mode.      TFTP
Set Server IP.  172.22.201.50
Path.      /
Filename.....
CN1610_CS_RCF_v1.2.txt
Data Type..... Config Script
Destination Filename.....
CN1610_CS_RCF_v1.2.scr
File with same name already exists.
WARNING:Continuing with this command will overwrite the existing
file.

Management access will be blocked for the duration of the
transfer Are you sure you want to start? (y/n) y

File transfer in progress. Management access will be blocked for
the duration of the transfer. please wait...
Validating configuration script...
(the entire script is displayed line by line)
...
description "NetApp CN1610 Cluster Switch RCF v1.2 - 2015-01-13"
...
Configuration script validated.
File transfer operation completed successfully.
```

- c. Verify that the RCF is downloaded to the switch.

Show example

```
(new_cs1)# script list
Configuration Script Nam    Size(Bytes)
-----
CN1610_CS_RCF_v1.1.scr      2191
CN1610_CS_RCF_v1.2.scr      2240
latest_config.scr           2356

4 configuration script(s) found.
2039 Kbytes free.
```

13. Apply the RCF to the CN1610 switch.

Show example

```
(new_cs1)# script apply CN1610_CS_RCF_v1.2.scr
Are you sure you want to apply the configuration script? (y/n) y
...
(the entire script is displayed line by line)
...
description "NetApp CN1610 Cluster Switch RCF v1.2 - 2015-01-13"
...
Configuration script 'CN1610_CS_RCF_v1.2.scr' applied. Note that the
script output will go to the console.
After the script is applied, those settings will be active in the
running-config file. To save them to the startup-config file, you
must use the write memory command, or if you used the reload answer
yes when asked if you want to save the changes.
```

a. Save the running configuration file so that it becomes the startup configuration file when you reboot the switch.

Show example

```
(new_cs1)# write memory
This operation may take a few minutes.
Management interfaces will not be available during this time.

Are you sure you want to save? (y/n) y

Config file 'startup-config' created successfully.

Configuration Saved!
```

b. Download the image to the CN1610 switch.

Show example

```
(new_cs1)# copy
tftp://<server_ip_address>/NetApp_CN1610_1.2.0.7.stk active
Mode.      TFTP
Set Server IP.  tftp_server_ip_address
Path.      /
Filename.....
NetApp_CN1610_1.2.0.7.stk
Data Type.  Code
Destination Filename.  active

Management access will be blocked for the duration of the
transfer

Are you sure you want to start? (y/n) y

TFTP Code transfer starting...

File transfer operation completed successfully.
```

- c. Run the new active boot image by rebooting the switch.

The switch must be rebooted for the command in step 6 to reflect the new image. There are two possible views for a response that you might see after you enter the reload command.

Show example

```
(new_cs1)# reload
The system has unsaved changes.
Would you like to save them now? (y/n) y

Config file 'startup-config' created successfully.

Configuration Saved! System will now restart!
.
.
.
Cluster Interconnect Infrastructure

User:admin Password: (new_cs1) >*enable*
```

- d. Copy the saved configuration file from the old switch to the new switch.

Show example

```
(new_cs1)# copy tftp://<server_ip_address>/<filename>.scr  
nvrRam:script <filename>.scr
```

- e. Apply the previously saved configuration to the new switch.

Show example

```
(new_cs1)# script apply <filename>.scr  
Are you sure you want to apply the configuration script? (y/n) y  
  
The system has unsaved changes.  
Would you like to save them now? (y/n) y  
  
Config file 'startup-config' created successfully.  
  
Configuration Saved!
```

- f. Save the running configuration file to the startup configuration file.

Show example

```
(new_cs1)# write memory
```

14. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all - message MAINT=xh
```

x is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

15. On the new switch new_cs1, log in as the admin user, and shut down all of the ports that are connected to the node cluster interfaces (ports 1 through 12).

Show example

```
User:*admin*
Password:
(new_cs1)> enable
(new_cs1)#
(new_cs1)# config
(new_cs1)(config)# interface 0/1-0/12
(new_cs1)(interface 0/1-0/12)# shutdown
(new_cs1)(interface 0/1-0/12)# exit
(new_cs1)# write memory
```

16. Migrate the cluster LIFs from the ports that are connected to the old_cs1 switch.

You must migrate each cluster LIF from its current node's management interface.

Show example

```
cluster::> set -privilege advanced
cluster::> network interface migrate -vserver <vserver_name> -lif
<Cluster_LIF_to_be_moved> - sourcenode <current_node> -dest-node
<current_node> -dest-port <cluster_port_that_is_UP>
```

17. Verify that all of the cluster LIFs have been moved to the appropriate cluster port on each node.

Show example

```
cluster::> network interface show -role cluster
```

18. Shut down the cluster ports that are attached to the switch that you replaced.

Show example

```
cluster::*> network port modify -node <node_name> -port
<port_to_admin_down> -up-admin false
```

19. Verify the health of the cluster.

Show example

```
cluster::*> cluster show
```

20. Verify that the ports are down.

Show example

```
cluster::*> cluster ping-cluster -node <node_name>
```

21. On the switch cs2, shut down the ISL ports 13 through 16.

Show example

```
(cs2)# config  
(cs2)(config)# interface 0/13-0/16  
(cs2)(interface 0/13-0/16)# shutdown  
(cs2)# show port-channel 3/1
```

22. Verify whether the storage administrator is ready for the replacement of the switch.
23. Remove all of the cables from the old_cs1 switch, and then connect the cables to the same ports on the new_cs1 switch.
24. On the cs2 switch, bring up the ISL ports 13 through 16.

Show example

```
(cs2)# config  
(cs2)(config)# interface 0/13-0/16  
(cs2)(interface 0/13-0/16)# no shutdown
```

25. Bring up the ports on the new switch that are associated with the cluster nodes.

Show example

```
(new_cs1)# config  
(new_cs1)(config)# interface 0/1-0/12  
(new_cs1)(interface 0/13-0/16)# no shutdown
```

26. On a single node, bring up the cluster node port that is connected to the replaced switch, and then confirm that the link is up.

Show example

```
cluster::*> network port modify -node node1 -port
<port_to_be_onlined> -up-admin true
cluster::*> network port show -role cluster
```

27. Revert the cluster LIFs that are associated with the port in step 25 on the same node.

In this example, the LIFs on node1 are successfully reverted if the “Is Home” column is true.

Show example

```
cluster::*> network interface revert -vserver node1 -lif
<cluster_lif_to_be_reverted>
cluster::*> network interface show -role cluster
```

28. If the first node’s cluster LIF is up and is reverted to its home port, repeat steps 25 and 26 to bring up the cluster ports and to revert the cluster LIFs on the other nodes in the cluster.
29. Display information about the nodes in the cluster.

Show example

```
cluster::*> cluster show
```

30. Confirm that the startup configuration file and running configuration file are correct on the replaced switch. This configuration file should match the output in step 1.

Show example

```
(new_cs1)> enable
(new_cs1)# show running-config
(new_cs1)# show startup-config
```

31. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Replace NetApp CN1610 cluster switches with switchless connections

You can migrate from a cluster with a switched cluster network to one where two nodes are directly connected for ONTAP 9.3 and later.

Review requirements

Guidelines

Review the following guidelines:

- Migrating to a two-node switchless cluster configuration is a nondisruptive operation. Most systems have two dedicated cluster interconnect ports on each node, but you can also use this procedure for systems with a larger number of dedicated cluster interconnect ports on each node, such as four, six or eight.
- You cannot use the switchless cluster interconnect feature with more than two nodes.
- If you have an existing two-node cluster that uses cluster interconnect switches and is running ONTAP 9.3 or later, you can replace the switches with direct, back-to-back connections between the nodes.


What you'll need

- A healthy cluster that consists of two nodes connected by cluster switches. The nodes must be running the same ONTAP release.
- Each node with the required number of dedicated cluster ports, which provide redundant cluster interconnect connections to support your system configuration. For example, there are two redundant ports for a system with two dedicated cluster interconnect ports on each node.

Migrate the switches

About this task

The following procedure removes the cluster switches in a two-node cluster and replaces each connection to the switch with a direct connection to the partner node.



About the examples

The examples in the following procedure show nodes that are using "e0a" and "e0b" as cluster ports. Your nodes might be using different cluster ports as they vary by system.

Step 1: Prepare for migration

1. Change the privilege level to advanced, entering `y` when prompted to continue:

```
set -privilege advanced
```

The advanced prompt `*>` appears.

2. ONTAP 9.3 and later supports automatic detection of switchless clusters, which is enabled by default.

You can verify that detection of switchless clusters is enabled by running the advanced privilege command:

```
network options detect-switchless-cluster show
```

Show example

The following example output shows if the option is enabled.

```
cluster::*> network options detect-switchless-cluster show
(network options detect-switchless-cluster show)
Enable Switchless Cluster Detection: true
```

If "Enable Switchless Cluster Detection" is `false`, contact NetApp support.

3. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message
MAINT=<number_of_hours>h
```

where `h` is the duration of the maintenance window in hours. The message notifies technical support of this maintenance task so that they can suppress automatic case creation during the maintenance window.

In the following example, the command suppresses automatic case creation for two hours:

Show example


```
cluster::*> system node autosupport invoke -node * -type all
-message MAINT=2h
```

Step 2: Configure ports and cabling

1. Organize the cluster ports on each switch into groups so that the cluster ports in `group1` go to cluster switch1 and the cluster ports in `group2` go to cluster switch2. These groups are required later in the procedure.
2. Identify the cluster ports and verify link status and health:

```
network port show -ipSpace Cluster
```

In the following example for nodes with cluster ports "e0a" and "e0b", one group is identified as "node1:e0a" and "node2:e0a" and the other group as "node1:e0b" and "node2:e0b". Your nodes might be using different cluster ports because they vary by system.



Verify that the ports have a value of `up` for the "Link" column and a value of `healthy` for the "Health Status" column.

Show example

```
cluster::> network port show -ipspace Cluster
Node: node1

Ignore
Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status
Speed (Mbps) Health
Status
-----
-----
e0a Cluster Cluster up 9000 auto/10000 healthy
false
e0b Cluster Cluster up 9000 auto/10000 healthy
false

Node: node2

Ignore
Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status
Speed (Mbps) Health
Status
-----
-----
e0a Cluster Cluster up 9000 auto/10000 healthy
false
e0b Cluster Cluster up 9000 auto/10000 healthy
false
4 entries were displayed.
```

3. Confirm that all the cluster LIFs are on their home ports.

Verify that the “is-home” column is `true` for each of the cluster LIFs:

```
network interface show -vserver Cluster -fields is-home
```

Show example

```
cluster::*> net int show -vserver Cluster -fields is-home
(network interface show)
vserver  lif          is-home
-----  -
Cluster  node1_clus1  true
Cluster  node1_clus2  true
Cluster  node2_clus1  true
Cluster  node2_clus2  true
4 entries were displayed.
```

If there are cluster LIFs that are not on their home ports, revert those LIFs to their home ports:

```
network interface revert -vserver Cluster -lif *
```

4. Disable auto-revert for the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

5. Verify that all ports listed in the previous step are connected to a network switch:

```
network device-discovery show -port cluster_port
```

The “Discovered Device” column should be the name of the cluster switch that the port is connected to.

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to cluster switches "cs1" and "cs2".

```
cluster:::> network device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol  Port   Device (LLDP: ChassisID)  Interface  Platform
-----  -
node1/cdp
          e0a    cs1                       0/11       BES-53248
          e0b    cs2                       0/12       BES-53248
node2/cdp
          e0a    cs1                       0/9        BES-53248
          e0b    cs2                       0/9        BES-53248
4 entries were displayed.
```

6. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node2		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node1		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node2		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```



```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

7. Verify that the cluster is healthy:

```
cluster ring show
```

All units must be either master or secondary.


8. Set up the switchless configuration for the ports in group 1.



To avoid potential networking issues, you must disconnect the ports from group1 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.


a. Disconnect all the cables from the ports in group1 at the same time.

In the following example, the cables are disconnected from port "e0a" on each node, and cluster traffic continues through the switch and port "e0b" on each node:



b. Cable the ports in group1 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2:



9. The switchless cluster network option transitions from *false* to *true*. This might take up to 45 seconds. Confirm that the switchless option is set to *true*:

```
network options switchless-cluster show
```

The following example shows that the switchless cluster is enabled:

```
cluster::*> network options switchless-cluster show
Enable Switchless Cluster: true
```

10. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```



Before proceeding to the next step, you must wait at least two minutes to confirm a working back-to-back connection on group 1.


11. Set up the switchless configuration for the ports in group 2.



To avoid potential networking issues, you must disconnect the ports from group2 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.

- a. Disconnect all the cables from the ports in group2 at the same time.

In the following example, the cables are disconnected from port "e0b" on each node, and cluster traffic continues through the direct connection between the "e0a" ports:



b. Cable the ports in group2 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2 and "e0b" on node1 is connected to "e0b" on node2:



Step 3: Verify the configuration

1. Verify that the ports on both nodes are correctly connected:

```
network device-discovery show -port cluster_port
```

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to the corresponding port on the cluster partner:

```
cluster::> net device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1/cdp
          e0a    node2                      e0a        AFF-A300
          e0b    node2                      e0b        AFF-A300
node1/lldp
          e0a    node2 (00:a0:98:da:16:44)  e0a        -
          e0b    node2 (00:a0:98:da:16:44)  e0b        -
node2/cdp
          e0a    node1                      e0a        AFF-A300
          e0b    node1                      e0b        AFF-A300
node2/lldp
          e0a    node1 (00:a0:98:da:87:49)  e0a        -
          e0b    node1 (00:a0:98:da:87:49)  e0b        -
8 entries were displayed.
```

2. Re-enable auto-revert for the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert true
```

3. Verify that all LIFs are home. This might take a few seconds.

```
network interface show -vserver Cluster -lif lif_name
```

Show example

The LIFs have been reverted if the “Is Home” column is `true`, as shown for `node1_clus2` and `node2_clus2` in the following example:

```
cluster::> network interface show -vserver Cluster -fields curr-  
port,is-home  
vserver  lif                curr-port  is-home  
-----  -  
Cluster  node1_clus1  e0a       true  
Cluster  node1_clus2  e0b       true  
Cluster  node2_clus1  e0a       true  
Cluster  node2_clus2  e0b       true  
4 entries were displayed.
```

If any cluster LIFS have not returned to their home ports, revert them manually from the local node:

```
network interface revert -vserver Cluster -lif lif_name
```

4. Check the cluster status of the nodes from the system console of either node:

```
cluster show
```

Show example

The following example shows `epsilon` on both nodes to be `false`:

```
Node  Health  Eligibility  Epsilon  
-----  
node1 true    true        false  
node2 true    true        false  
2 entries were displayed.
```

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the `show` command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

Packet	Source	Destination
Node	LIF	LIF
Date		
Loss		
node1		
3/5/2022 19:21:18 -06:00	node1_clus2	node2-clus1
node		
3/5/2022 19:21:20 -06:00	node1_clus2	node2_clus2
node2		
3/5/2022 19:21:18 -06:00	node2_clus2	node1_clus1
node		
3/5/2022 19:21:20 -06:00	node2_clus2	node1_clus2
node		

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```



```

cluster1::*> cluster ping-cluster -node local
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

6. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

For more information, see [NetApp KB Article 1010449: How to suppress automatic case creation during scheduled maintenance windows](#).

7. Change the privilege level back to admin:

```
set -privilege admin
```

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

<https://www.netapp.com/company/legal/copyright/>

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf>

Privacy policy

<https://www.netapp.com/company/legal/privacy-policy/>

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.