



Cluster switches

Install and maintain

NetApp
March 06, 2026

Table of Contents

- Cluster switches 1
 - Broadcom-supported BES-53248 1
 - Get started 1
 - Install the hardware 5
 - Configure the software 8
 - Upgrade the switch 106
 - Migrate the switches 150
 - Replace the switches 187
 - Cisco Nexus 9336C-FX2 or 9336C-FX2-T 214
 - Get started 214
 - Install the hardware 219
 - Configure the software 231
 - Migrate the switches 303
 - Replace the switches 358
 - NVIDIA SN2100 390
 - Get started 390
 - Install the hardware 393
 - Configure the software 402
 - Migrate the switches 478
 - Replace the switches 535

Cluster switches

Broadcom-supported BES-53248

Get started

Installation and setup workflow for BES-53248 switches

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

Follow these workflow steps to install and set up your to BES-53248 switches.

1

Review the configuration requirements

Review the configuration requirements for the BES-53248 cluster switch.

2

Review the components and part numbers

Review the components and part numbers for the BES-53248 cluster switch.

3

Review the required documentation

Review specific switch and controller documentation to set up your BES-53248 switches and the ONTAP cluster.

4

Install the hardware

Install the switch hardware.

5

Configure the software

Configure the switch software.

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-52348 switch machine types, see [Components and part numbers for BES-53248 cluster switches](#). See [What additional information do I need to install my equipment that is not in HWU?](#)[^] for more information about switch installation requirements.

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. See [What additional information do I need to install my equipment that is not in HWU?](#) for more information about switch installation requirements.

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.

What's next

After you've reviewed the configuration requirements, you can confirm your [components and part numbers](#).

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
.
.
.
```

What's next

After you've confirmed your components and part numbers, you can review the [required documentation](#).

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.

Document title	Description
<i>EFOS SNMP Reference Guide v3.4.3</i>	Provides examples of how to use the BES-53248 switch in a typical network.
<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 the hardware

Hardware install workflow for BES-53248 switches

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

1 Install the switch hardware

Install and configure the BES-53248 switch hardware.

2 Review cabling and configuration

Review the cabling and configuration considerations for the BES-53248 cluster switch.

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?

After you've installed the hardware for the switch, you can [review cabling and configuration](#) requirements.

Review cabling and configuration considerations

Before configuring your Broadcom BES-53248 switch, review the following considerations.

Cluster port switch assignments

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

Switch ports	Ports usage
0-16	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. See [What additional information do I need to install my equipment that is not in HWU?](#) for more information about switch installation requirements.

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.
- If speeds in a 4-port group are different, the switch ports will not operate correctly.

FEC requirements

- For 25G ports with copper cables, see the following table for details.

If the Controller side is `auto`, the switch side is set to FEC 25G.

FAS2820 FEC			Switch FEC			link status
write	read		write	read		
	requested_fec	negotiated_fec		Configured FEC Mode	Physical FEC Status	
fc	FC-FEC/BASE-R	none	No FEC	FEC Disabled	FEC Disabled	UP
fc	FC-FEC/BASE-R	FC-FEC/BASE-R	FEC 25G	FEC 25G	CL-74	UP
auto	RS-FEC	none	FEC 25G	FEC 25G	CL74	UP
auto	RS-FEC	none	No FEC	FEC Disabled	FEC Disabled	UP
none	none	none	No FEC	FEC Disabled	FEC Disabled	UP
none	none	none	FEC 25G	FEC 25G	CL74	UP
rs	RS-FEC	none	FEC 25G	FEC 25G	CL74	UP
rs	RS-FEC	none	No FEC	FEC Disabled	FEC Disabled	UP

- For 25G ports with Fiber/Optical cables, see the following table for details:

FAS2820 FEC			Switch FEC			link status
write	read		write	read		
	requested_fec	negotiated_fec		Configured FEC Mode	Physical FEC Status	
fc	FC-FEC/BASE-R	none	No FEC	FEC Disabled	FEC Disabled	DOWN
fc	FC-FEC/BASE-R	FC-FEC/BASE-R	FEC 25G	FEC 25G	CL-74	UP
auto	RS-FEC	none	FEC 25G	FEC 25G	CL74	DOWN
auto	RS-FEC	none	No FEC	FEC Disabled	FEC Disabled	DOWN
none	none	none	No FEC	FEC Disabled	FEC Disabled	UP
none	none	none	FEC 25G	FEC 25G	CL74	DOWN
rs	RS-FEC	none	FEC 25G	FEC 25G	CL74	DOWN
rs	RS-FEC	none	No FEC	FEC Disabled	FEC Disabled	DOWN

Bootarg implementation

Use the following command to set the 25G port FEC to either `auto` or `fc`, as required:

```
systemsshell -node <node> -command sudo sysctl
dev.ice.<X>.requested_fec=<auto/fc>
```

- When set to **auto**:
 - The `auto` setting propagates the setting to hardware immediately and no reboot is required.
 - If `bootarg.cpk_fec_fc_eXx` already exists, it is deleted from the bootarg storage.
 - After a reboot, the `auto` setting remains in place since `auto` is the default FEC setting.
- When set to **fc**:
 - The `FC-FEC` setting propagates the setting to the hardware immediately and no reboot is required.
 - A new `bootarg.cpk_fec_fc_eXx` is created with the value set to "true".
 - After a reboot, `FC-FEC` setting remains in place for the driver code to use.

Configure the software

Software install workflow for BES-53248 switches

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

1

Configure the switch

Configure the BES-53248 cluster switch.

2

Install the EFOS software

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

3

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.

4

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.

5

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.

6

Reset the switch to factory defaults

Erase the BES-53248 cluster switch settings.

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 IPv4 or IPv6 management address for the switch's service port:

IPv4

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 <ip-address> <netmask> <gateway>
```

IPv6

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 ipv6 <address>
(cs2) # serviceport ipv6 <gateway>
```

7. Verify the results using the command:

```
show serviceport
```

```
(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?

After you've configured your switches, you can [install the EFOS software](#).

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 four steps to reconfigure the switch:
 - a. [Install licenses](#)
 - b. [Install the RCF file](#)
 - c. [Enable SSH](#)
 - d. [Configure switch health 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.

Steps

1. Connect to a console that is attached to the serial port of the switch.
2. Boot the switch into ONIE installation mode.

During boot, select ONIE when you see the prompt.


```
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 five steps to reconfigure the switch:
 - a. [Configure switch](#)
 - b. [Install licenses](#)
 - c. [Install the RCF file](#)
 - d. [Enable SSH](#)

e. [Configure switch health monitoring](#)

9. Repeat steps 1 to 8 on the partner switch.

What's next

After you've installed the EFOS software, you can [install your licenses](#).

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 -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
e0b         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
e0b         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
```

Current Vserver Port	Logical Current Home	Is Interface	Status Admin/Oper	Network Address/Mask	Node
Cluster					
cluster1-01	cluster1-01_e0a	clus1 true	up/up	169.254.3.4/23	
cluster1-01	cluster1-01_e0b	clus2 true	up/up	169.254.3.5/23	
cluster1-02	cluster1-02_e0a	clus1 true	up/up	169.254.3.8/23	
cluster1-02	cluster1-02_e0b	clus2 true	up/up	169.254.3.9/23	
cluster1-03	cluster1-03_e0a	clus1 true	up/up	169.254.1.3/23	
cluster1-03	cluster1-03_e0b	clus2 true	up/up	169.254.1.1/23	
cluster1-04	cluster1-04_e0a	clus1 true	up/up	169.254.1.6/23	
cluster1-04	cluster1-04_e0b	clus2 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 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 -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_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 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 -ipSpace Cluster
```

Show example

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

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 -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          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 -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
          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	Date	LIF
Loss		LIF
-----	-----	-----
-----	-----	-----
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?

After you've installed the RCF and license file, you can [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



To redeem a transaction key for a port license key file, go to the [License Portal for Broadcom Supported Ethernet Switches](#) page. See the Knowledge Base article [How to add Additional Port Licensing for the Broadcom BES-53248 Switch](#) for further details.

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

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

After you've installed the licenses, you can [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 -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
e0b         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
e0b         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
```

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

```

- Verify that the cluster LIFs 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::*>
```

- 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: HTTP, HTTPS, FTP, TFTP, SFTP, or SCP.

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

```
(cs2)# copy tftp://172.19.2.1/BES-53248-RCF-v1.9-Cluster-HA.txt
nvram:reference-config
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: HTTP, HTTPS, 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	Mode	Admin	Physical	Physical	Link	Link
Mode	Type	Mode	Mode	Status	Status	Trap
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
```

```
(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 -ipSpace Cluster
```

Show example

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

After you've installed the RCF, you can [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
```

- If SSH is not disabled, disable it as follows:

```
no ip ssh server enable
```

```
no ip scp server enable
```



- For EFOS 3.12 and later, console access is required as active SSH sessions are lost when SSH is disabled.
- For EFOS 3.11 and earlier, current SSH sessions are kept open after disabling the SSH server.



Make sure that you disable SSH before you modify the keys, otherwise, a warning is reported on the switch.

2. In config mode, 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
```

3. In config mode, set AAA authorization for ONTAP log collection:

```
aaa authorization commands "noCmdAuthList" none
```

Show example

```
(switch) (Config)# aaa authorization commands "noCmdAuthList" none
switch (Config)# exit
```

4. Re-enable SSH/SCP.

Show example

```
(switch)# ip ssh server enable
switch)# ip scp server enable
switch)# ip ssh pubkey-auth
```

5. Save these changes to the startup-config:

```
write memory
```

Show example

```
(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!
```

6. 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!

7. Reboot the switch:

```
reload
```

8. 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?

After you've enabled SSH, you can [configure switch health monitoring](#).

Reset the BES-53248 cluster switch to factory defaults

To reset the BES-53248 cluster switch to factory defaults, you must erase the BES-53248 switch settings.

About this task

- You must be connected to the switch using the serial console.
- This task resets the configuration of the management network.

Steps

1. Change to the elevated command prompt.

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

2. Erase the startup configuration.

```
erase startup-config
```

```
(cs2)# erase startup-config
```

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

3. Reboot the switch.

```
(cs2)# reload
```

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



If the system asks whether to save the unsaved or changed configuration before reloading the switch, select **No**.

4. Wait for the switch to reload, and then log in to the switch.

The default user is “admin”, and no password is set. A prompt similar to the following is displayed:

```
(Routing) >
```

Upgrade the switch

Upgrade workflow 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 backup image means that when you reboot, that image establishes the running EFOS version, completing the update.

```
(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     3.10.0.3     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            3.10.0.3
```

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
```

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

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

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

```
network interface show -vserver 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

- Disable auto-revert on the cluster LIFs.

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

- 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 -vserver 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.

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 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 -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/100000
healthy false
e0b Cluster Cluster up 9000 auto/100000
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/100000
healthy false
e0b Cluster Cluster up 9000 auto/100000
healthy false
8 entries were displayed.

Node: cluster1-03

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: cluster1-04
```

```
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
```

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

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 -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_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. However, when you upgrade to RCF 1.12 or later, the modifications are no longer needed because all interfaces are now pre-configured.
- 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: HTTP, HTTPS, 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
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: HTTP, HTTPS, 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 -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/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
-----
-----
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
-----
-----
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 -vserver 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 -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
          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	Date	LIF
Loss		LIF
-----	-----	-----
-----	-----	-----
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?

After you've confirmed the upgrade of your EFOS software or RCF, you can [configure switch health monitoring](#).

Migrate the 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. See [What additional information do I need to install my equipment that is not in HWU?](#) for more information about switch installation requirements.

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/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface
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  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			

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
```

```
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
```

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 interface 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	Port	Interface	Admin/Oper	Address/Mask	Node
	Home				

Cluster					
e0b	false	node1_clus1	up/up	169.254.209.69/16	node1
e0b	true	node1_clus2	up/up	169.254.49.125/16	node1
e0b	false	node2_clus1	up/up	169.254.47.194/16	node2
e0b	true	node2_clus2	up/up	169.254.19.183/16	node2

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
```

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

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	Status	Speed(Mbps)	Health
------	---------	-----------	--------	------	-----	------------	--------	-------------	--------

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 interface 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. On switch cs2, shut down and restart all cluster ports to trigger an auto-revert of all cluster LIFs that are not on their home ports.

```
cs2> enable
cs2# configure
cs2(config)# interface 0/1-0/16
cs2(config-if-range)# shutdown

(Wait for 5-10 seconds before re-enabling the ports)

cs2(config-if-range)# no shutdown

(After executing the no shutdown command, the nodes detect the change
and begin to auto-revert the cluster LIFs to their home ports)

cs2(config-if-range)# exit
cs2(config)# exit
cs2#
```

3. Verify that the cluster LIFs have reverted to their home ports (this might take a minute):

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

If any of the cluster LIFs have not reverted to their home port, manually revert them. You must connect to each node-mgmt LIF or SP/BMC system console of the local node that owns the LIF:

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

4. Verify that the cluster is healthy:

```
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
```

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
node1	3/5/2022 19:21:20 -06:00	node1_clus2
node1	3/5/2022 19:21:20 -06:00	node2_clus2
node2	3/5/2022 19:21:18 -06:00	node2_clus2
node2	3/5/2022 19:21:18 -06:00	node1_clus1
node2	3/5/2022 19:21:20 -06:00	node2_clus2
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)

```

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
```

```

cluster1::*> system node autosupport invoke -node * -type all -message
MAINT=END

```

What's next?

After you've migrated your switches, you can [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. See [What additional information do I need to install my equipment that is not in HWU?](#) for more information about switch installation requirements.

- 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					
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					

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
e0b	Cluster	Cluster	up	9000	auto/10000		healthy

```
Node: node2
```

```
Ignore
```

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

Step 3: Verify the configuration

1. Enable auto-revert on the cluster LIFs.

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

2. On switch cs2, shut down and restart all cluster ports to trigger an auto-revert of all cluster LIFs that are not on their home ports.

```
cs2> enable
cs2# configure
cs2(config)# interface 0/1-0/16
cs2(config-if-range)# shutdown

(Wait for 5-10 seconds before re-enabling the ports)

cs2(config-if-range)# no shutdown

(After executing the no shutdown command, the nodes detect the change
and begin to auto-revert the cluster LIFs to their home ports)

cs2(config-if-range)# exit
cs2(config)# exit
cs2#
```

3. Verify that the cluster LIFs have reverted to their home ports (this might take a minute):

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

If any of the cluster LIFs have not reverted to their home port, manually revert them. You must connect to each node-mgmt LIF or SP/BMC system console of the local node that owns the LIF:

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

4. 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
```

	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
true					

5. 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
```

6. 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
```

7. 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
```

8. 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

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	Date	LIF
Loss		LIF
node1	3/5/2022 19:21:18 -06:00	node1_clus2
node2	3/5/2022 19:21:18 -06:00	node2_clus2
node1	3/5/2022 19:21:20 -06:00	node1_clus2
node2	3/5/2022 19:21:20 -06:00	node2_clus2
node1	3/5/2022 19:21:18 -06:00	node2_clus1
node2	3/5/2022 19:21:18 -06:00	node1_clus1
node1	3/5/2022 19:21:20 -06:00	node2_clus2
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)

```

10. Change the privilege level back to admin:

```
set -privilege admin
```

11. 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?

After you've migrated your switches, you can [configure switch health monitoring](#).

Replace the 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/ Ports	Port Timeout	Port Speed	Port Active
0/55	actor/long partner/long	100G Full	True	
0/56	actor/long partner/long	100G Full	True	

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
	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

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

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

	Cluster			
e0a	node1_clus1	up/up	169.254.209.69/16	node1
e0b	node1_clus2	up/up	169.254.49.125/16	node1
	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?

After you've replaced your switches, you can [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.

Before you begin

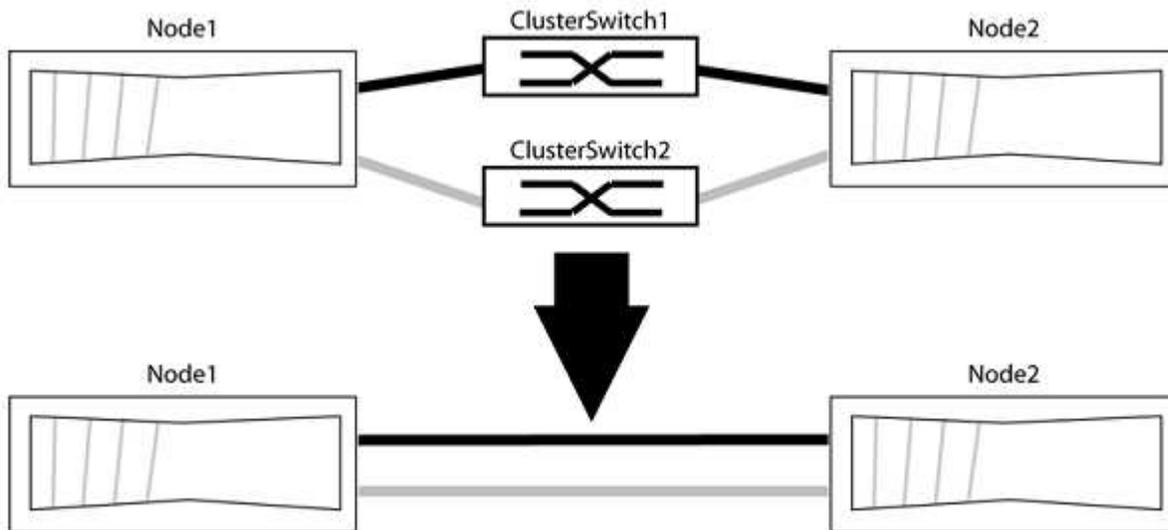
Make sure you have the following:

- 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

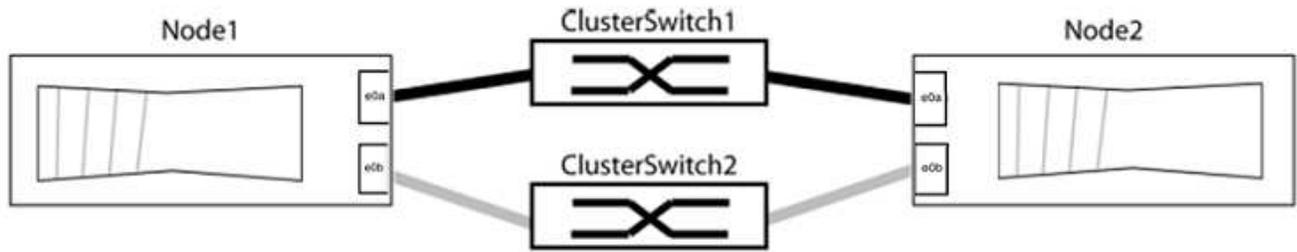
```
cluster::*> system node autosupport invoke -node * -type all
-message MAINT=2h
```

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.

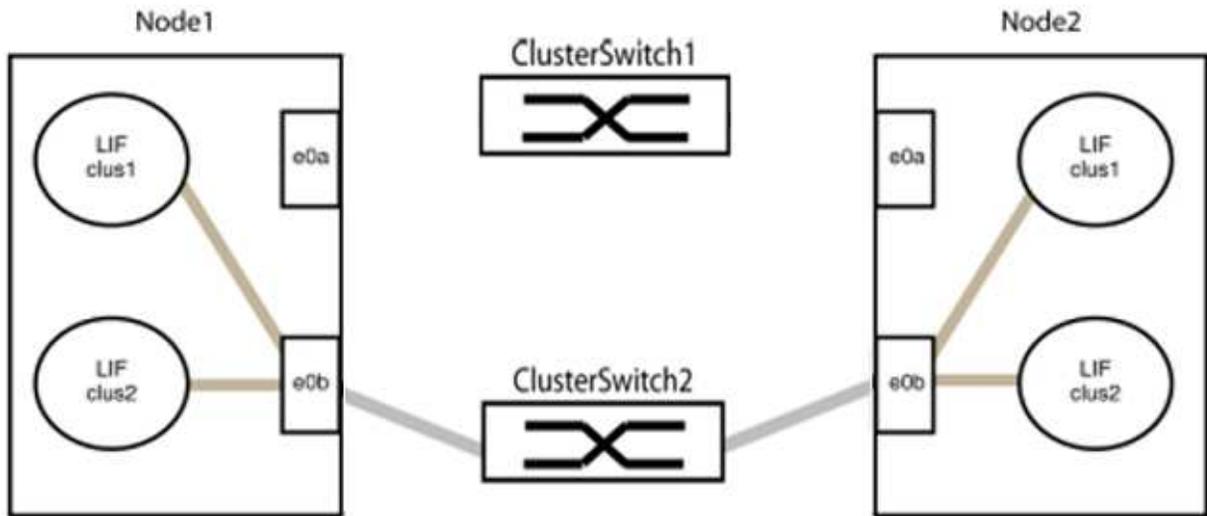
8. Set up the switchless configuration for the ports in group 1.



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**.

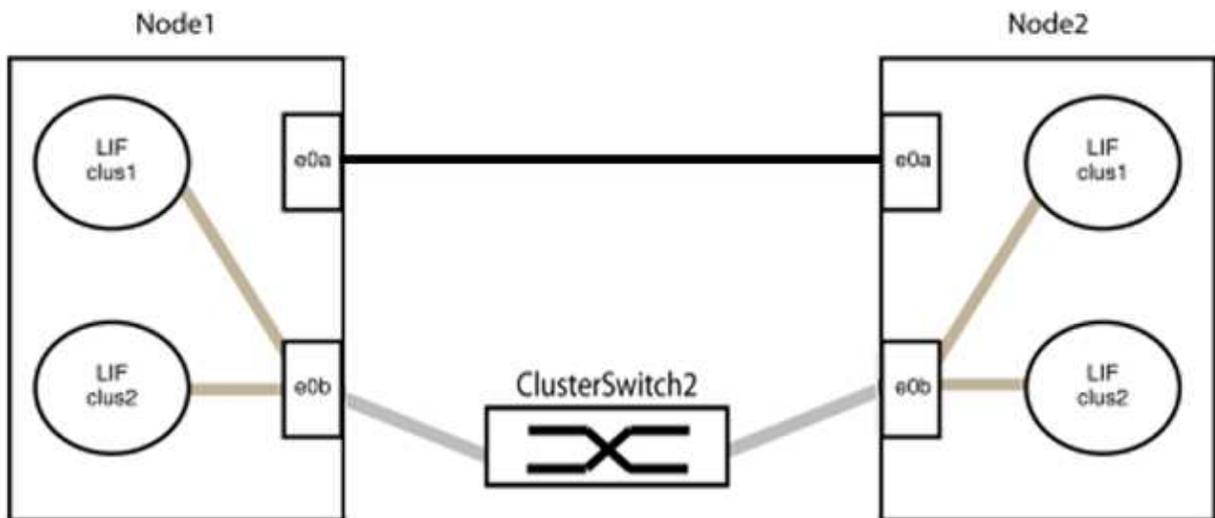
a. Disconnect all the cables from the ports in group1 at the same time.

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
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.

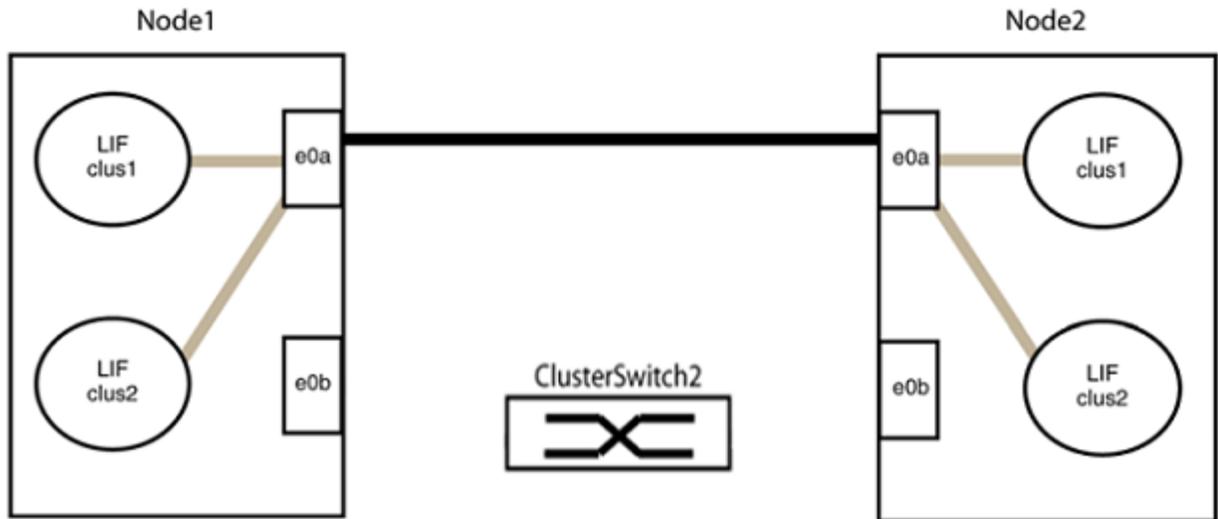
11. Set up the switchless configuration for the ports in group 2.



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
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)

```

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
```

What's next?

After you've replaced your switches, you can [configure switch health monitoring](#).

Cisco Nexus 9336C-FX2 or 9336C-FX2-T

Get started

Installation and setup workflow for Cisco Nexus 9336C-FX2 and 9336C-FX2-T cluster switches

The Cisco Nexus 9336C-FX2 and 9336C-FX2-T cluster switches are 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.

Cisco Nexus 9336C-FX2 (36 ports) is high-port density cluster/storage/data switch. Cisco Nexus 9336C-FX2-T (12 ports) is a low-port-density, high performance switch that supports 10/25/40/100GbE cluster configurations.

Follow these workflow steps to install and setup your to Cisco 9336C-FX2 and 9336C-FX2-T switches.

1

Review the configuration requirements

Review the configuration requirements for the 9336C-FX2 and 9336C-FX2-T cluster switches.

2

Review the components and part numbers

Review the components and part numbers for the 9336C-FX2 and 9336C-FX2-T cluster switches.

3

Review the required documentation

Review specific switch and controller documentation to set up your 9336C-FX2 and 9336C-FX2-T switches and the ONTAP cluster.

4

Review the Smart Call Home requirements

Review the requirements for the Cisco Smart Call Home feature, used to monitor the hardware and software components on your network.

5

Install the hardware

Install the switch hardware.

6

Configure the software

Configure the switch software.

Configuration requirements for Cisco Nexus 9336C-FX2 and 9336C-FX2-T cluster switches

For Cisco Nexus 9336C-FX2 and 9336C-FX2-T switch installation and maintenance, make sure to review the configuration and network requirements.

ONTAP support

ONTAP 9.9.1 and later

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).

ONTAP 9.10.1 and later

In addition, beginning with ONTAP 9.10.1, you can use Cisco Nexus 9336C-FX2-T 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

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](#).

What's next

After you've reviewed the configuration requirements, you can confirm your [components and part numbers](#).

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
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 for 9336C-FX2-T ports only

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.

What's next

After you've confirmed your components and part numbers, you can review the [required documentation](#).

Documentation requirements for Cisco Nexus 9336C-FX2 and 9336C-FX2-T switches

For Cisco Nexus 9336C-FX2 and 9336C-FX2-T switches 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 and 9336C-FX2-T 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 [ONTAP 9](#).

Name	Description
<i>Controller-specific Installation and Setup Instructions</i>	Describes how to install NetApp hardware.

Name	Description
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 Cisco Nexus 9336C-FX2 and 9336C-FX2-T switches in a four-post NetApp cabinet.

Smart Call Home requirements

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.

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.

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.

Before you can use Smart Call Home, be aware of the following requirements:

- 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 the hardware

Hardware install workflow for Cisco Nexus 9336C-FX2 and 9336C-FX2-T switches

To install and configure the hardware for 9336C-FX2 and 9336C-FX2-T cluster switches,

follow these steps:

1

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.

2

Install the switch

Install the 9336C-FX2 and 9336C-FX2-T switches.

3

Install the switch in a NetApp cabinet

Install the 9336C-FX2 and 9336C-FX2-T switches and pass-through panel in a NetApp cabinet as required.

4

Review cabling and configuration

Review support for NVIDIA Ethernet ports, 25GbE FEC requirements, and information on TCAM resources.

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

Cluster switch A		Cluster switch B	
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
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

Cluster switch A		Cluster switch B	
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	
12		12	
13		13	
14		14	
15		15	

Cluster switch A		Cluster switch B	
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
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

Cluster switch A		Cluster switch B	
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	
7		7	
8		8	
9		9	
10		10	

Cluster switch A		Cluster switch B	
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.

What's next

After you've completed your cabling worksheets, you can [install the switch](#).

Install 9336C-FX2 and 9336C-FX2-T cluster switches

Follow this procedure to set up and configure the Cisco Nexus 9336C-FX2 and 9336C-FX2-T switches.

Before you begin

Make sure you have the following:

- 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?

Optionally, you can [install a Cisco Nexus 9336C-FX2 switch in a NetApp cabinet](#). Otherwise, go to [review cabling and configuration](#).

Install Cisco Nexus 9336C-FX2 and 9336C-FX2-T switches in a NetApp cabinet

Depending on your configuration, you might need to install the Cisco Nexus 9336C-FX2 and 9336C-FX2-T switch and pass-through panel in a NetApp cabinet. Standard brackets are included with the switch.

Before you begin

Make sure you have the following:

- 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](#).

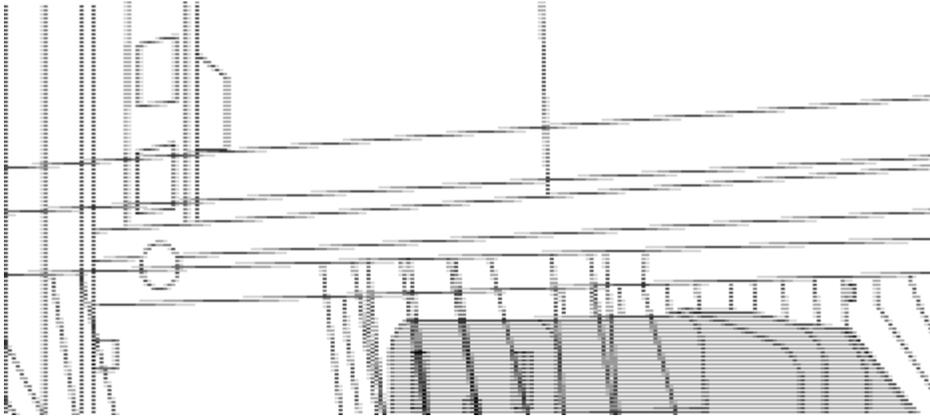
Steps

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.

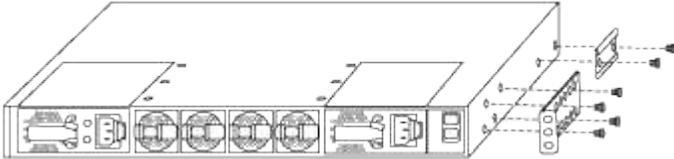
In this procedure, the blanking panel is installed in U40.

- 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.

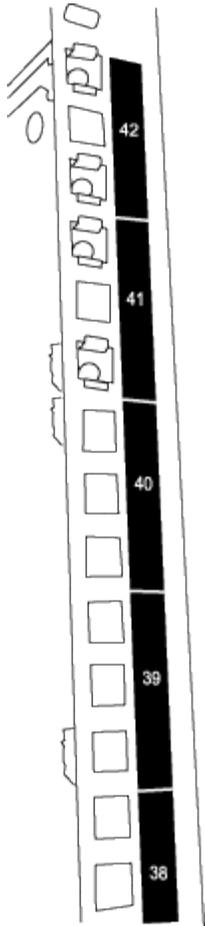


(1) Female connector of the jumper cord.

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.



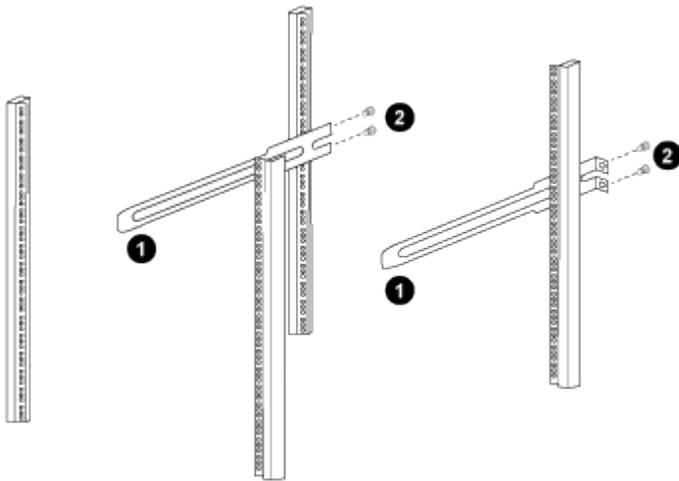
- b. Repeat step 2a with the other front rack-mount bracket on the other side of the switch.
 - c. Install the rear rack-mount bracket on the switch chassis.
 - d. Repeat step 2c with the other rear rack-mount bracket on the other side of the switch.
3. Install the clip nuts in the square hole locations for all four IEA posts.



The two 9336C-FX2 and 9336C-FX2-T 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.

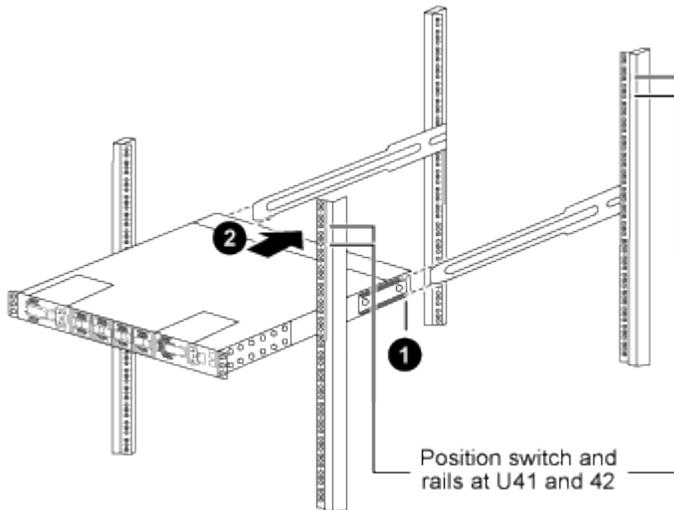
(2) Tighten the screws of the slider rails to the cabinet posts.

- 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.

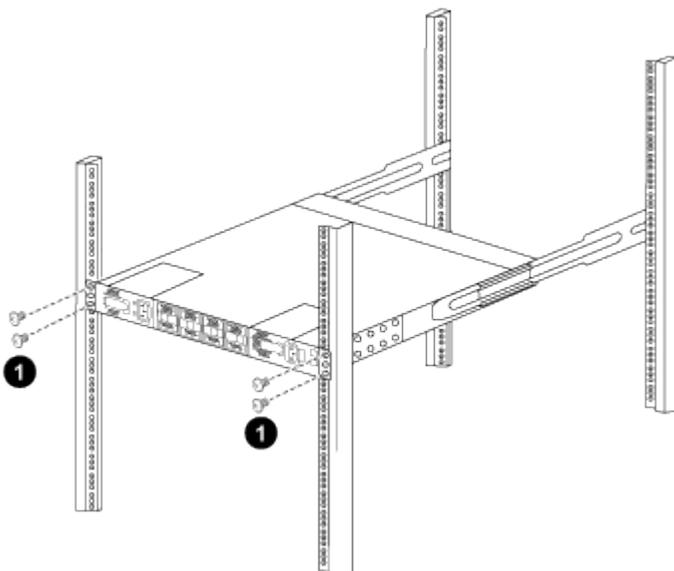
- a. Position the back of the switch at RU41.



(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 and 9336C-FX2-T switches 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?

After you've installed the switches in the NetApp cabinet, you can [configure the Cisco Nexus 9336C-FX2 and 9336C-FX2-T switches](#).

Review cabling and configuration considerations

Before configuring your 9336C-FX2 and 9336C-FX2-T switches, 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. See [What additional information do I need to install my equipment that is not in HWU?](#) for more information about switch installation requirements.

25GbE FEC requirements

FAS2820 e0a/e0b ports

FAS2820 e0a and e0b ports require FEC configuration changes to link up with 9336C-FX2 and 9336C-FX2-T 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 and 9336C-FX2-T switches, 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 the software

Software install workflow for Cisco Nexus 9336C-FX2 and 9336C-FX2-T cluster switches

To install and configure the software for Cisco Nexus 9336C-FX2 and 9336C-FX2-T switches and to install or upgrade the Reference Configuration File (RCF), follow these steps:

1

Configure the switch

Configure the 9336C-FX2 and 9336C-FX2-T cluster switches.

2

Prepare to install the NX-OS software and RCF

The Cisco NX-OS software and reference configuration files (RCFs) must be installed on Cisco 9336C-FX2 and 9336C-FX2-T cluster switches.

3

Install or upgrade the NX-OS software

Download and install or upgrade the NX-OS software on the Cisco 9336C-FX2 and 9336C-FX2-T cluster switches.

4

Install or upgrade the RCF

Install or upgrade the RCF after setting up the Cisco 9336C-FX2 and 9336C-FX2-T switches for the first time. You can also use this procedure to upgrade your RCF version.

5

Verify SSH configuration

Verify that SSH is enabled on the switches to use the Ethernet Switch Health Monitor (CSHM) and log collection features.

6

Reset the switch to factory defaults

Erase the 9336C-FX2 and 9336C-FX2-T cluster switches settings.

Configure the 9336C-FX2 and 9336C-FX2-T cluster switches

Follow this procedure to configure the Cisco Nexus 9336C-FX2 and 9336C-FX2-T switches.

Before you begin

Make sure you have the following:

- 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.

Prompt	Response
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.	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: 1px solid #ccc; 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.
Configure default interface layer (L3/L2)	Respond with L2 . The default is L2.

Prompt	Response
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>

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?

After you've configured your switches, you can [prepare to install the NX-OS software and 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.

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

4 entries were displayed.

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?

After you've prepared to install the NX-OS software and RCF, you can [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 and 9336C-FX2-T cluster switches.

Before you begin, complete the procedure in [Prepare to install NX-OS and RCF](#).

Review requirements

Before you begin

Make sure you do the following:

- Run the `show install all impact nxos bootflash:<image_name>.bin` command on the switch to review the impact of installing or upgrading the new NX-OS software image. It verifies the image integrity, checks for necessary reboots, evaluates hardware compatibility, and confirms sufficient space.
- Review the release notes for the target NX-OS software version to check for any specific requirements.
- Verify that you have a current backup of the switch configuration.
- Verify that you have 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 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      N9K-
C9336C-FX2
              e0b    cs2                      Ethernet1/7      N9K-
C9336C-FX2
cluster1-02/cdp
              e0a    cs1                      Ethernet1/8      N9K-
C9336C-FX2
              e0b    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 -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
e0b         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
e0b         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
            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
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. If you are using NX-OS 10.5(3)F or later, you do not need to perform a separate EPLD image upgrade because the EPLD image is included in the NX-OS image. Skip this step and continue with verifying cluster port health.

Show example



```
cs2# show version module 1 epld
```

EPLD	Device	Version
MI	FPGA	0x7
IO	FPGA	0x17
MI	FPGA2	0x2
GEM	FPGA	0x2

```
cs2# install epld bootflash:n9000-epld.9.3.5.img module all
```

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

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 -ipSpace Cluster
```

Show example

```
cluster1::*> network port show -ipspace 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
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
          e0b    cs2                      Ethernet1/7
N9K-C9336C-FX2
cluster01-2/cdp
          e0a    cs1                      Ethernet1/8
N9K-C9336C-FX2
          e0b    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. Verify the connectivity of the remote cluster interfaces before enabling auto-revert on the cluster LIFs:

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)

```

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 -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_clus1 up/up      169.254.3.4/23
cluster1-01 e0b true
          cluster1-01_clus2 up/up      169.254.3.5/23
cluster1-01 e0b true
          cluster1-02_clus1 up/up      169.254.3.8/23
cluster1-02 e0b true
          cluster1-02_clus2 up/up      169.254.3.9/23
cluster1-02 e0b 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?

After you've installed or upgraded the NX-OS software, you can [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 and 9336C-FX2-T switches 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. See [Cisco Ethernet Switches](#) for more information.

For specific port and VLAN usage details, refer to the banner and important notes section in your RCF.

RCF configuration	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.

Available RCFs

The following table lists the available RCFs for 9336C-FX2 and 9336C-FX2-T switches. Choose the applicable RCF version for your configuration. See [Cisco Ethernet Switches](#) for more information.

RCF name
Cluster-HA-Breakout RCF 1.xx
Cluster-HA-Storage RCF 1.xx
Storage RCF 1.xx
MultiCluster-HA RCF 1.xx

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 9000 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?

After you've reviewed the install RCF or upgrade RCF procedure, you can [install the RCF](#) or [upgrade your RCF](#) as required.

Install the Reference Configuration File (RCF)

You install the Reference Configuration File (RCF) after setting up the Nexus 9336C-FX2 and 9336C-FX2-T switches 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.

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. Record any custom additions between the current `running-config` file and the RCF file in use.
7. After you verify that the RCF versions and switch settings are correct, copy the `running-config` file to the `startup-config` file.

```
cs1# copy running-config startup-config
[#####] 100% Copy complete
```

8. Save basic configuration details to the `write_erase.cfg` file on the bootflash.

Make sure to configure the following:



- Username and password
- Management IP address
- Default gateway
- Switch name

```
cs1# show run | i "username admin password" > bootflash:write_erase.cfg
```

```
cs1# show run | section "vrf context management" >> bootflash:write_erase.cfg
```

```
cs1# show run | section "interface mgmt0" >> bootflash:write_erase.cfg
```

```
cs1# show run | section "switchname" >> bootflash:write_erase.cfg
```

9. When installing RCF version 1.12 and later, run the following commands:

```
cs1# echo "hardware access-list tcam region ing-racl 1024" >>
bootflash:write_erase.cfg
```

```
cs1# echo "hardware access-list tcam region egr-racl 1024" >>
bootflash:write_erase.cfg
```

```
cs1# echo "hardware access-list tcam region ing-l2-qos 1280" >>
bootflash:write_erase.cfg
```

See the Knowledge Base article [How to clear configuration on a Cisco interconnect switch while retaining remote connectivity](#) for further details.

10. Verify that the `write_erase.cfg` file is populated as expected:

```
show file bootflash:write_erase.cfg
```

11. Repeat steps 1 through 10 on switch `cs2`.

12. 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?

After you've installed the RCF, you can [verify the 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. 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. 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::*>
```

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 -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
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::*>
```

5. 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> enable

cs1# configure

cs1(config)# interface eth1/1/1-2,eth1/7-8

cs1(config-if-range)# shutdown
```

```
cs1(config-if-range)# exit
```

```
cs1# exit
```



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 -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		

```
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. Beginning with NX-OS 10.2, 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.

Make sure to configure the following:



- Username and password
- Management IP address
- Default gateway
- Switch name

```
cs1# show run | i "username admin password" >> bootflash:write_erase.cfg
```

```
cs1# show run | section "vrf context management" >> bootflash:write_erase.cfg
```

```
cs1# show run | section "interface mgmt0" >> bootflash:write_erase.cfg
```

```
cs1# show run | section "switchname" >> bootflash:write_erase.cfg
```

6. When upgrading to RCF version 1.12 and later, run the following commands:

```
cs1# echo "hardware access-list tcam region ing-racl 1024" >>
bootflash:write_erase.cfg
```

```
cs1# echo "hardware access-list tcam region egr-racl 1024" >>
bootflash:write_erase.cfg
```

```
cs1# echo "hardware access-list tcam region ing-l2-qos 1280" >>
bootflash:write_erase.cfg
```

See the Knowledge Base article [How to clear configuration on a Cisco interconnect switch while retaining remote connectivity](#) for further details.

7. Verify that the `write_erase.cfg` file is populated as expected:

```
show file bootflash:write_erase.cfg
```

8. Issue the write erase command to erase the current saved configuration:

```
cs1# write erase
```

```
Warning: This command will erase the startup-configuration.
```

```
Do you wish to proceed anyway? (y/n) [n] y
```

9. Copy the previously saved basic configuration into the startup configuration.

```
cs1# copy bootflash:write_erase.cfg startup-config
```

10. Perform a reboot of the switch:

```
switch# reload
```

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

11. 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.

12. 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.

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)...
```

13. 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](#).

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
```



Make sure to read thoroughly the **Installation notes**, **Important Notes**, and **banner** sections of your RCF. You must read and follow these instructions to ensure the proper configuration and operation of the switch.

14. 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.

15. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
16. 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](#).

```
cs1# copy running-config startup-config

[] 100% Copy complete
```

17. 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
```

18. 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
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: cluster1-02

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: cluster1-03

Ignore

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

```

19. 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::*>

```

20. Repeat steps 1 to 19 on switch cs2.

21. Enable auto-revert on the cluster LIFs.

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

22. Perform a reboot of switch cs2.

```
cs2# reload
```

This command will reboot the system. (y/n)? [n] **y**

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 -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_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?

After you've upgraded your RCF, you can [verify the 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?

After you've verified your SSH configuration, you can [configure switch health monitoring](#).

Reset 9336C-FX2 and 9336C-FX2-T cluster switches to factory defaults

To reset the 9336C-FX2 and 9336C-FX2-T cluster switches to factory defaults, you must

erase the 9336C-FX2 and 9336C-FX2-T switch settings.

About this task

- You must be connected to the switch using the serial console.
- This task resets the configuration of the management network.

Steps

1. Erase the existing configuration:

```
write erase
```

```
(cs2)# write erase
```

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

2. Reload the switch software:

```
reload
```

```
(cs2)# reload
```

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

The system reboots and enters the configuration wizard. During the boot, if you receive the prompt “Abort Auto Provisioning and continue with normal setup? (yes/no)[n]”, you should respond **yes** to proceed.

What's next

After you've reset your switches, you can [reconfigure](#) them as needed.

Migrate the switches

Migrate from NetApp CN1610 cluster switches to Cisco 9336C-FX2 and 9336C-FX2-T cluster switches

You can migrate NetApp CN1610 cluster switches for an ONTAP cluster to Cisco 9336C-FX2 and 9336C-FX2-T 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 and 9336C-FX2-T cluster switches. You must 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
- Cisco 9336C-FX2-T

For details of supported ports and their configurations, see the [Hardware Universe](#). See [What additional information do I need to install my equipment that is not in HWU?](#) for more information about switch installation requirements.

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 and 9336C-FX2-T 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 and 9336C-FX2-T switches to run at 40GbE or 100GbE.
- You have planned, migrated, and documented 40GbE and 100GbE connectivity from nodes to Cisco 9336C-FX2 and 9336C-FX2-T 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
e3b	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
```

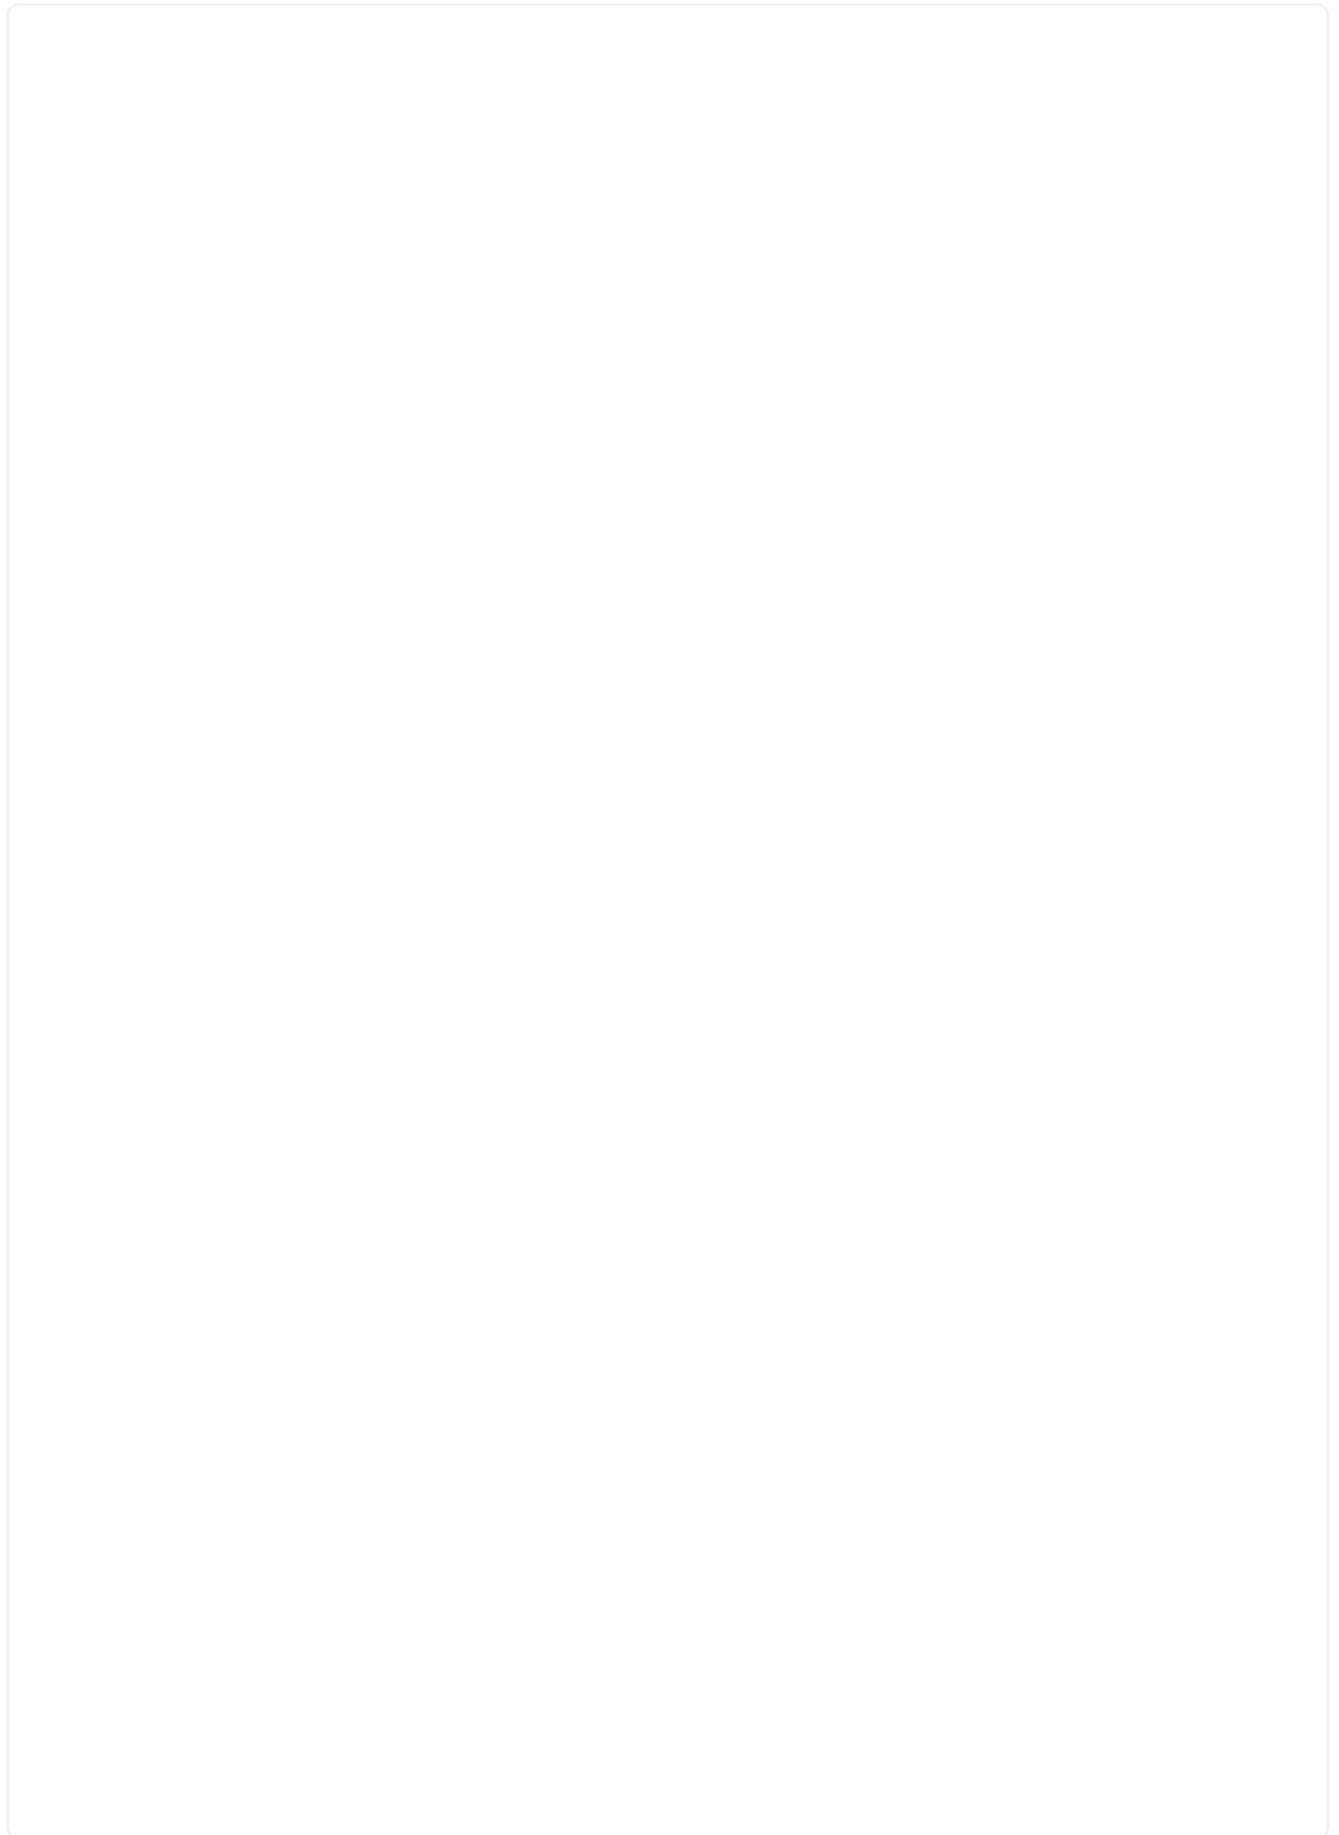
Node/	Local	Discovered		
Protocol	Port	Device (LLDP: ChassisID)	Interface	
Platform				

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 and 9336C-FX2-T.
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:

```
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 and 9336C-FX2-T.
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. On switch cs2, shut down and restart all cluster ports to trigger an auto-revert of all cluster LIFs that are not on their home ports.

```
cs2> enable
cs2# configure
cs2(config)# interface eth1/1-1/2
cs2(config-if-range)# shutdown

(Wait for 5-10 seconds before re-enabling the ports)

cs2(config-if-range)# no shutdown

(After executing the no shutdown command, the nodes detect the change
and begin to auto-revert the cluster LIFs to their home ports)

cs2(config-if-range)# exit
cs2(config)# exit
cs2#
```

3. Verify that the cluster LIFs have reverted to their home ports (this might take a minute):

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

If any of the cluster LIFs have not reverted to their home port, manually revert them. You must connect to each node-mgmt LIF or SP/BMC system console of the local node that owns the LIF:

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

4. Verify that the cluster is healthy:

```
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
```



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
node1		

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)

```

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?

After you've migrated your switches, you can [configure switch health monitoring](#).

Migrate from older Cisco switches to Cisco Nexus 9336C-FX2 and 9336C-FX2-T switches

You can perform a nondisruptive migration from older Cisco cluster switches to Cisco Nexus 9336C-FX2 and 9336C-FX2-T cluster network switches.

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. See [What additional information do I need to install my equipment that is not in HWU?](#) for more information about switch installation requirements.

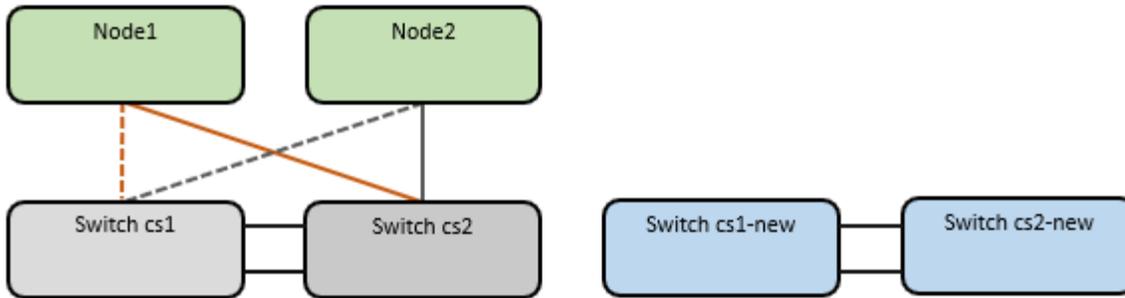


The command outputs might vary depending on the different releases of ONTAP.

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	node1_clus1	up/up	169.254.3.4/16	node1
e0a	node1_clus2	up/up	169.254.3.5/16	node1
e0a	node2_clus1	up/up	169.254.3.8/16	node2
e0a	node2_clus2	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:
- 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#
```

- 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	node1_clus1	up/up	169.254.3.4/16	node1
	true			
e0a	node1_clus2	up/up	169.254.3.5/16	node1
	false			
e0a	node2_clus1	up/up	169.254.3.8/16	node2
	true			
e0a	node2_clus2	up/up	169.254.3.9/16	node2
	false			

- Verify that the cluster is healthy:

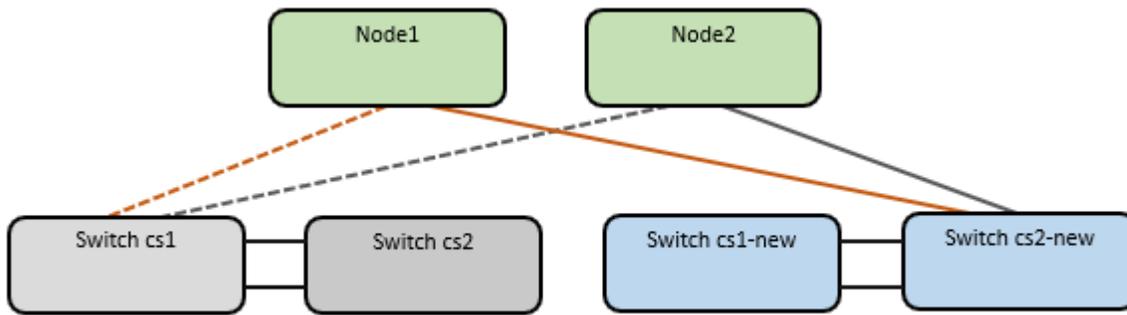
```
cluster show
```

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:

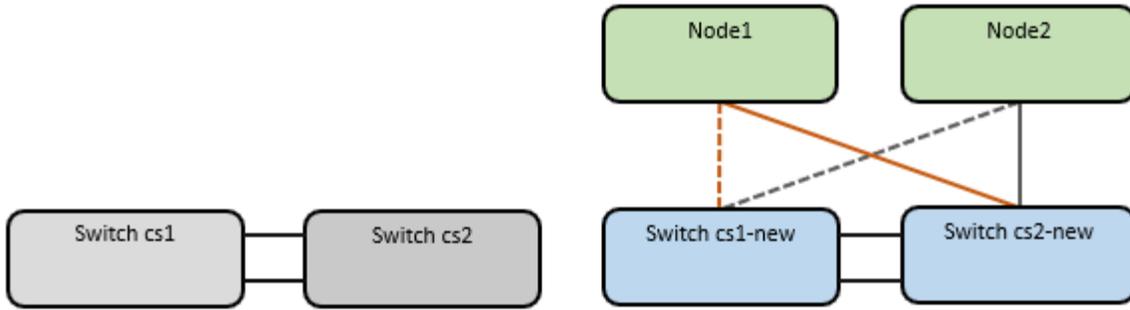
```
cluster show
```

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. On switch cs2, shut down and restart all cluster ports to trigger an auto-revert of all cluster LIFs that are not on their home ports.

```
cs2> enable
cs2# configure
cs2(config)# interface eth1/1-1/2
cs2(config-if-range)# shutdown
```

(Wait for 5-10 seconds before re-enabling the ports)

```
cs2(config-if-range)# no shutdown
```

(After executing the no shutdown command, the nodes detect the change and begin to auto-revert the cluster LIFs to their home ports)

```
cs2(config-if-range)# exit
cs2(config)# exit
cs2#
```

3. Verify that the cluster LIFs have reverted to their home ports (this might take a minute):

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

If any of the cluster LIFs have not reverted to their home port, manually revert them. You must connect to each node-mgmt LIF or SP/BMC system console of the local node that owns the LIF:

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

4. Verify that the cluster is healthy:

```
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
```



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
node1		

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)

```

6. 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?

After you've migrated the switches, you can [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 and 9336C-FX2-T 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. See [What additional information do I need to install my equipment that is not in HWU?](#) for more information about switch installation requirements.

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
-----
Admin/Oper    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) Health
-----
Admin/Oper    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. Disable auto-revert on all of 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

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 and operational:

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

Show example

The following example shows that all of the LIFs are up on node1 and node2:

```
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	e0b	
false					
node1_clus2	up/up	169.254.49.125/16	node1	e0b	
true					
node2_clus1	up/up	169.254.47.194/16	node2	e0b	
false					
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
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.
```

Step 3: Verify the configuration

1. Enable auto-revert on the cluster LIFs.

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

2. On switch cs2, shut down and restart all cluster ports to trigger an auto-revert of all cluster LIFs that are not on their home ports.

```
cs2> enable
cs2# configure
cs2(config)# interface eth1/1-1/2
cs2(config-if-range)# shutdown
```

(Wait for 5-10 seconds before re-enabling the ports)

```
cs2(config-if-range)# no shutdown
```

(After executing the no shutdown command, the nodes detect the change and begin to auto-revert the cluster LIFs to their home ports)

```
cs2(config-if-range)# exit
cs2(config)# exit
cs2#
```

3. Verify that the cluster LIFs have reverted to their home ports (this might take a minute):

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

If any of the cluster LIFs have not reverted to their home port, manually revert them. You must connect to each node-mgmt LIF or SP/BMC system console of the local node that owns the LIF:

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

4. 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.

5. 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
```

6. 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.
```

7. 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
```

8. 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

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
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)

```

10. Change the privilege level back to admin:

```
set -privilege admin
```

11. 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?

After you've migrated your switches, you can [configure switch health monitoring](#).

Replace the switches

Replace Cisco Nexus 9336C-FX2 and 9336C-FX2-T cluster switches

Follow these steps to replace defective Nexus 9336C-FX2 and 9336C-FX2-T switches 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/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface      Platform
-----
node2      /cdp
           e0a    cs1                       Eth1/2         N9K-
C9336C
           e0b    cs2                       Eth1/2         N9K-
C9336C
node1      /cdp
           e0a    cs1                       Eth1/1         N9K-
C9336C
           e0b    cs2                       Eth1/1         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
cs2            Eth1/35        176     R S I s     N9K-C9336C
Eth1/35
cs2 (FDO220329V5) 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	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.

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
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  
e0a      true  
         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?

After you've replaced your switches, you can [configure switch health monitoring](#).

Replace Cisco Nexus 9336C-FX2 and 9336C-FX2-T 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.

Before you begin

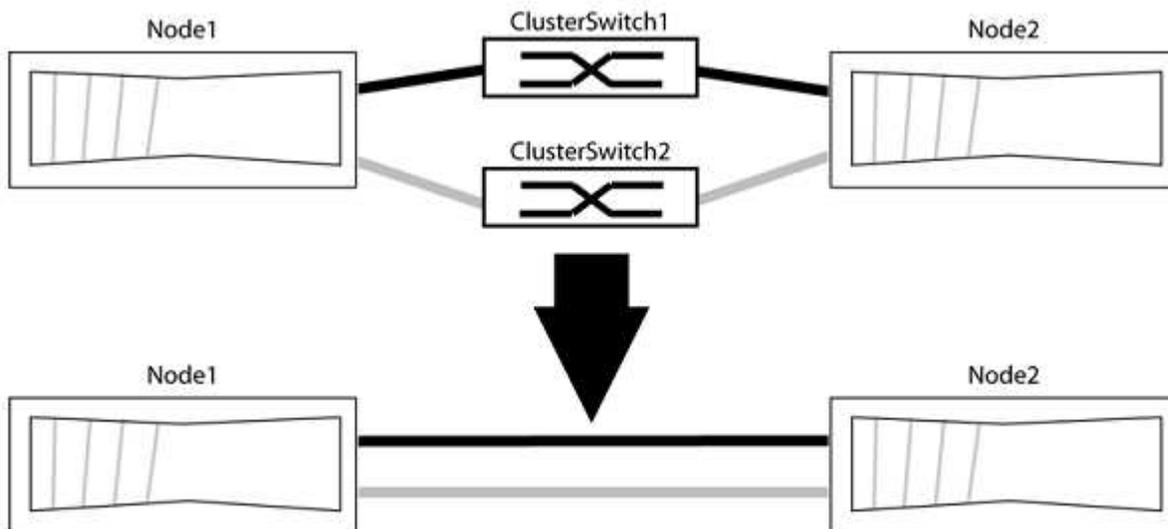
Make sure you have the following:

- 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

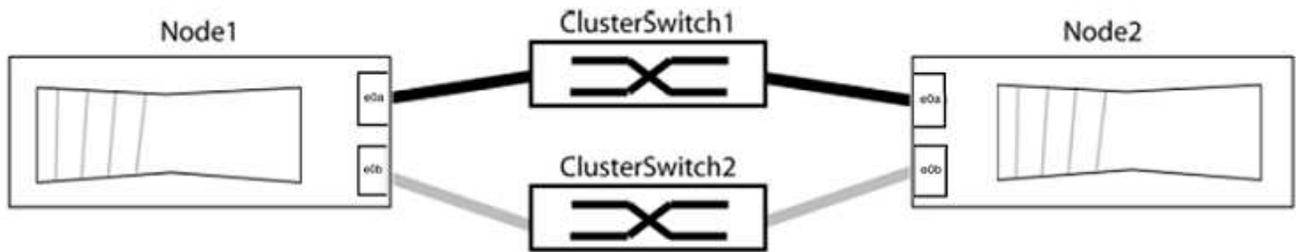
```
cluster::*> system node autosupport invoke -node * -type all
-message MAINT=2h
```

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.

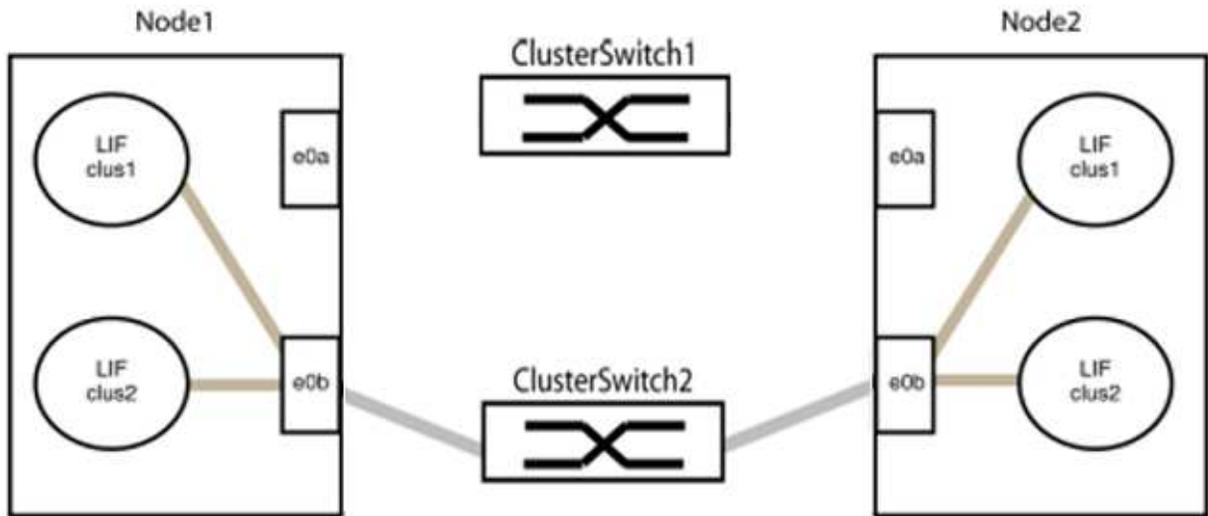
8. Set up the switchless configuration for the ports in group 1.



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**.

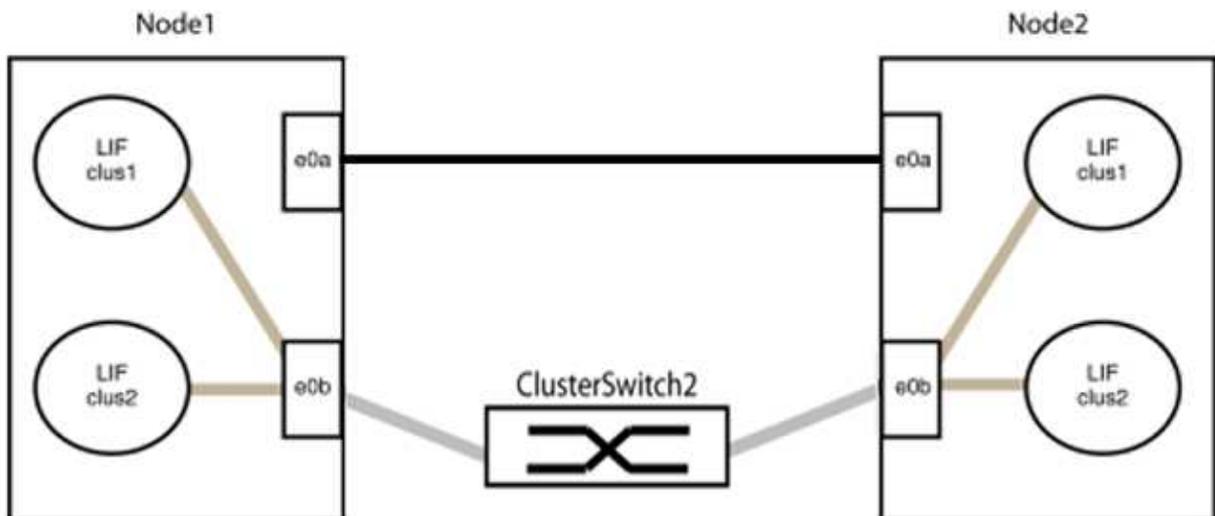
a. Disconnect all the cables from the ports in group1 at the same time.

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.

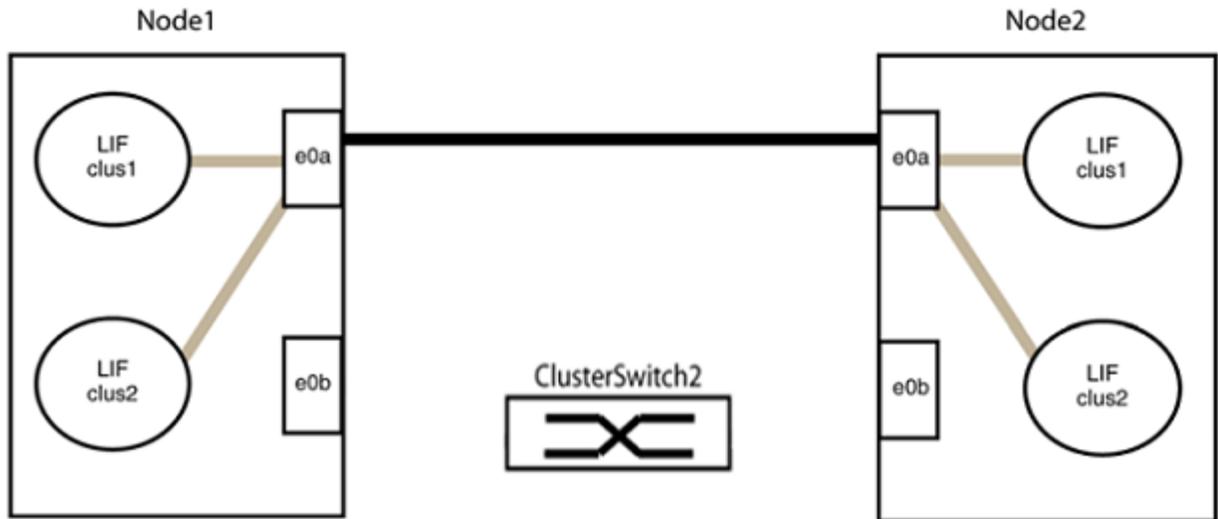
11. Set up the switchless configuration for the ports in group 2.



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**.

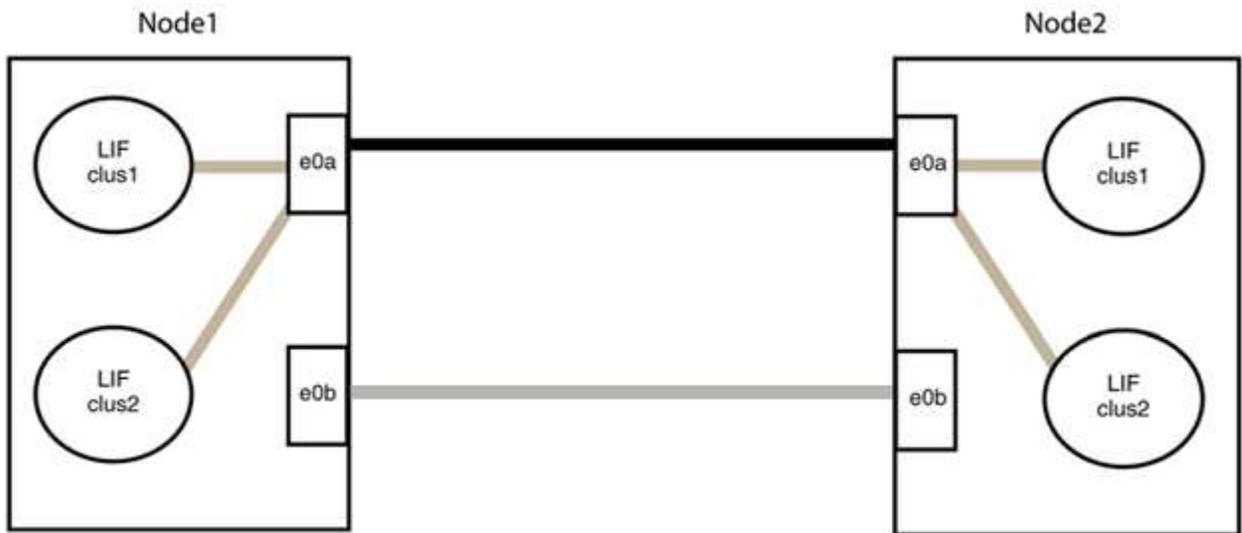
- a. Disconnect all the cables from the ports in group2 at the same time.

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
```

NVIDIA SN2100

Get started

Installation and setup workflow for NVIDIA SN2100 switches

The NVIDIA SN2100 is a cluster switch that allows you to build ONTAP clusters with more than two nodes.

Follow these workflow steps to install and set up your to NVIDIA SN2100 switches.

1**Review the configuration requirements**

Review the configuration requirements for the SN2100 cluster switch.

2**Review the components and part numbers**

Review the components and part numbers for the SN2100 cluster switch.

3**Review the required documentation**

Review specific switch and controller documentation to set up your SN2100 switches and the ONTAP cluster.

4**Install the hardware**

Install the switch hardware.

5**Configure the software**

Configure the switch software.

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.

What's next?

After you've reviewed the configuration requirements, you can confirm your [components and part numbers](#).

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



Refer to the NVIDIA documentation for details on [installing your SN2100 switch and rail kit](#).

What's next?

After you've confirmed your components and part numbers, you can review the [required documentation](#).

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.

Title	Description
NetApp Hardware Universe	Allows you to confirm supported hardware, such as storage switches and cables, for your platform model.

Install the hardware

Hardware install workflow for NVIDIA SN2100 switches

To install and configure the hardware for a SN2100 cluster switch, follow these steps:

1

Install the hardware

Install the switch hardware.

2

Review cabling and configuration considerations

Review requirements for optical connections, the QSA adapter, and the switchport speed.

3

Cable the NS224 shelves

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).

Install the hardware for the NVIDIA SN2100 switch

Installing the NVIDIA SN2100 switch hardware is a critical first step in building a reliable and scalable network infrastructure to support data protection, replication, and management workflows.

Follow these steps to set up and configure your SN2100 switches. Refer to NVIDIA's documentation for detailed instructions.

Steps

1. Review the [configuration requirements](#).
2. Follow the instructions in the [NVIDIA Switch Installation Guide](#).

What's next?

After you've installed your hardware, you can [review cabling and configuration](#) requirements.

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 B0	0x11 (QSFP28)	Molex	112-00576	93A2229911111
swp4 B0	0x11 (QSFP28)	Molex	112-00576	93A2229922222

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 module

When using QSFP+ (40GbE) to SFP+ (10GbE) adapters or QSFP28 (100GbE) to SFP28 (25GbE) adapters (QSA), insert them in non-breakout 40GbE/100GbE switch ports (swp3-swp14). Do not insert the QSA module in a port that is configured for breakout.

When a QSA module 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 link speed to 10000 and set auto-negotiation to off.
- For 25GbE, manually set the link speed to 25000 and set auto-negotiation to off.

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 or QSA module, 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)

See the [Hardware Universe](#) and the Knowledge Base article [What additional information do I need to install my equipment that is not in HWU](#) for more information.

What's next?

After you've reviewed your cabling and configuration requirements, you can [cable the 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?

After you've cabled your shelves, you can [configure the switch](#).

Configure the software

Software install workflow for NVIDIA SN2100 switches

To install and configure software for a NVIDIA SN2100 switch, follow these steps:

1

Configure the switch

Configure the NVIDIA SN2100 switch.

2

Install Cumulus Linux in Cumulus mode

You can install the Cumulus Linux (CL) OS when the switch is running Cumulus Linux.

3

Install Cumulus Linux in ONIE mode

Alternatively, you can install the Cumulus Linux (CL) OS when the switch is running Cumulus Linux in ONIE mode.

4

Upgrade your Cumulus Linux version, as required

You can upgrade your Cumulus Linux (CL) OS as required.

5

Install or upgrade the Reference Configuration File (RCF) script

There are two RCF scripts available for Clustering and Storage applications. The procedure for each is the same.

6

Install the CSHM file

You can install the applicable configuration file for Ethernet switch health monitoring of NVIDIA cluster switches.

7

Reset the switch to factory defaults

Erase the SN2100 cluster switch settings.

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?

After you've configured your switch, you can [install Cumulus Linux in Cumulus mode](#) or [install Cumulus Linux in ONIE mode](#).

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](#)).

Before you begin

Make sure you have the following:

- 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
10.233.204.71/24
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 mgmt listen eth0
cumulus@sw1:~$ nv set service ntp mgmt server <server> iburst on
cumulus@sw1:~$ nv config apply
cumulus@sw1:~$ nv config save
```

See the Knowledge Base article [NTP server configuration is not working with NVIDIA SN2100 switches](#) for further details.

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@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@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 mgmt listen eth0
cumulus@sw1:~$ nv set service ntp mgmt server <server> iburst on
cumulus@sw1:~$ nv config apply
cumulus@sw1:~$ nv config save
```

See the Knowledge Base article [NTP server configuration is not working with NVIDIA SN2100 switches](#) for further details.

- 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?

After you've installed Cumulus Linux in Cumulus mode, you [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.


```

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?

After you've installed Cumulus Linux in ONIE mode, you can [install the Reference Configuration File \(RCF\) script](#).

Upgrade Cumulus Linux versions

Complete the following procedure to upgrade your Cumulus Linux version as required.

Before you begin

Make sure you have the following:

- 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

Cumulus Linux 4.4.x to Cumulus Linux 5.4.0

1. Connect the cluster switch to the management network.
2. Use the ping command to verify connectivity to the server hosting the Cumulus Linux and the RCF.
3. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

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
```

- b. Verify that all the cluster interfaces (LIFs) are on the home port:

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

- c. Verify that the cluster displays information for both cluster switches:

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

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. Check the current Cumulus Linux version and connected ports:

```

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

```

```

cumulus@cumulus: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)					
.					
.					

7. Download the Cumulux Linux 5.4.0 image:

```
cumulus@cumulus: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.
```

8. Reboot the switch:

```
cumulus@cumulus:mgmt:~$ sudo reboot
```

9. 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'
```

10. 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
```

11. 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
.
.
```

12. 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:~$
```

13. 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
```

14. 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:~$

```

15. 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.x to Cumulus Linux 5.4.0

1. Connect the cluster switch to the management network.
2. Use the ping command to verify connectivity to the server hosting the Cumulus Linux and the RCF.
3. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

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
```

- b. Verify that all the cluster interfaces (LIFs) are on the home port:

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

- c. Verify that the cluster displays information for both cluster switches:

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

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. Check the current Cumulus Linux version and connected ports:

```

cumulus@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

cumulus@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

```

7. Download the Cumulux Linux 5.4.0 image:

```
cumulus@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.
```

8. Reboot the switch:

```
cumulus@sw1:mgmt:~$ sudo reboot
```

9. 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'
```

10. 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
```

11. 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
.
.
```

12. 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:~$
```

13. 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
```

14. 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:~$

```

15. 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 5.4.0 to Cumulus Linux 5.11.0

1. Connect the cluster switch to the management network.
2. Use the ping command to verify connectivity to the server hosting the Cumulus Linux and the RCF.
3. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

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
```

- b. Verify that all the cluster interfaces (LIFs) are on the home port:

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

- c. Verify that the cluster displays information for both cluster switches:

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

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. Check the current Cumulus Linux version and connected ports:

```

cumulus@sw1:mgmt:~$ nv show system
operational          applied
-----
hostname             cumulus             cumulus
build                Cumulus Linux 5.4.0
uptime               6 days, 8:37:36
timezone             Etc/UTC

cumulus@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

```

7. Download the Cumulux Linux 5.11.0 image:

```
cumulus@sw1:mgmt:~$ sudo onie-install -a -i http://<ip-to-  
webserver>/path/to/cumulus-linux-5.11.0-mlx-amd64.bin  
[sudo] password for cumulus:  
Fetching installer: http://<ip-to-webserver>/path/to/cumulus-linux-  
5.11.0-mlx-amd64.bin  
Downloading URL: http://<ip-to-webserver>/path/to/cumulus-linux-  
5.11.0-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.
```

8. Reboot the switch:

```
cumulus@sw1:mgmt:~$ sudo reboot
```

9. 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 5.11.0-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'
```

10. Check the Cumulus Linux version: `nv show system`

```
cumulus@cumulus:mgmt:~$ nv show system
operational      applied
-----
hostname         cumulus cumulus
build            Cumulus Linux 5.11.0
uptime           14:07:08
timezone         Etc/UTC
```

11. 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/syncd/syncd.conf
.
.
```

12. Log out 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 5.11.0-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:~$
```

13. 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
```

What's next?

After you've upgraded your Cumulus Linux version, you can [install or upgrade the RCF script](#).

Install or upgrade the Reference Configuration File (RCF) script

Follow this procedure to install or upgrade the RCF script.

Before you begin

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.

Customer configurations

The following reference configuration categories are available:

Cluster	On ports configured for 4x10GbE breakout, one port is configured for 4x25GbE breakout and the other ports are configured for 40/100GbE. Supports shared cluster/HA traffic on ports for nodes that use shared cluster/HA ports. See the platform table in the Knowledge Base article What AFF, ASA, and FAS platforms use shared Cluster and HA Ethernet ports? . All ports can be used as dedicated cluster ports too.
Storage	All ports configured for 100GbE NVMe storage connections.

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. Connect the cluster switch to the management network.
2. Use the `ping` command to verify connectivity to the server hosting the Cumulus Linux and the RCF.
3. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

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
```

- b. Verify that all the cluster interfaces (LIFs) are on the home port:

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

- c. Verify that the cluster displays information for both cluster switches:

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

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
```

Note the following:



- If you are upgrading your RCF, you must disable auto-revert for this step.
- If you have just upgraded your Cumulus Linux version, you don't need to disable auto-revert for this step because it's already disabled.

6. 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		

7. 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, for example SFTP, HTTPS, or FTP.

8. 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.

9. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.

10. 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%

11. 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
```

12. 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

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:

```
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

```

14. Verify that the cluster is healthy:

```
cluster show
```

15. Repeat steps 1 to 14 on the second switch.

16. Enable auto-revert on the cluster LIFs.

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

Cumulus Linux 5.4.0

1. Connect the cluster switch to the management network.
2. Use the `ping` command to verify connectivity to the server hosting the Cumulus Linux and the RCF.
3. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

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
```

- b. Verify that all the cluster interfaces (LIFs) are on the home port:

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

- c. Verify that the cluster displays information for both cluster switches:

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

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
```

Note the following:



- If you are upgrading your RCF, you must disable auto-revert for this step.
- If you have just upgraded your Cumulus Linux version, you don't need to disable auto-revert for this step because it's already disabled.

6. 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

```

7. 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, for example SFTP, HTTPS, or FTP.

8. 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.net` 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.net
*****
*****
*
* 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.

9. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
10. 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.

11. 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
```

12. 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

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:

```
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

```

14. Verify that the cluster is healthy:

```
cluster show
```

15. Repeat steps 1 to 14 on the second switch.

16. Enable auto-revert on the cluster LIFs.

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

Cumulus Linux 5.11.0

1. Connect the cluster switch to the management network.
2. Use the `ping` command to verify connectivity to the server hosting the Cumulus Linux and the RCF.
3. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

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
```

- b. Verify that all the cluster interfaces (LIFs) are on the home port:

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

- c. Verify that the cluster displays information for both cluster switches:

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

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
```

Note the following:



- If you are upgrading your RCF, you must disable auto-revert for this step.
- If you have just upgraded your Cumulus Linux version, you don't need to disable auto-revert for this step because it's already disabled.

6. 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

```

7. 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, for example SFTP, HTTPS, or FTP.

8. 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.net` 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.net
*****
*****
*
* 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.

9. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
10. 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: fd20:8b1e: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

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.

11. Verify information for the transceiver in the interface:

```
admin@sw1:mgmt:~$ nv show platform transceiver
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
```

12. 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

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:

```
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

```

14. Verify that the cluster is healthy:

```
cluster show
```

15. Repeat steps 1 to 14 on the second switch.

16. Enable auto-revert on the cluster LIFs.

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

What's next?

After you've installed the RCF, you can [install the CSHM file](#).

Install the Ethernet Switch Health Monitor configuration file

To configure Ethernet switch health monitoring on NVIDIA Ethernet switches, follow this procedure.

These instructions apply if NVIDIA X190006-PE and X190006-PI switches are not detected properly, which can be confirmed by running `system switch ethernet show` and checking if **OTHER** is shown for your model. To identify your NVIDIA switch model, find the part-number with the command `nv show platform hardware` for NVIDIA CL 5.8 and earlier or `nv show platform` for later versions.



These steps are also recommended if you want health monitoring and log collection to work as intended when using NVIDIA CL 5.11.x with the following ONTAP releases. While health monitoring and log collection might still function without these steps, following them ensures everything operates correctly.

- 9.10.1P20, 9.11.1P18, 9.12.1P16, 9.13.1P8, 9.14.1, 9.15.1 and later patch releases

Before you begin

- Make sure that the ONTAP cluster is up and running.
- Enable SSH on the switch 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 switch ethernet configure-health-monitor
```

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 configure-health-monitor 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 configure-health-monitor show
```



If your model is still showing **OTHER** after applying the configuration file, contact NetApp support.

See the [system switch ethernet configure-health-monitor](#) command for further details.

What's next?

After you've installed the CSHM file, you can [configure switch health monitoring](#).

Reset the SN2100 cluster switch to factory defaults

To reset the SN2100 cluster switch to factory defaults:

- For Cumulus Linux 5.10 and earlier, you apply the Cumulus image.
- For Cumulus Linux 5.11 and later, you use the `nv action reset system factory-default` command.

About this task

- You must be connected to the switch using the serial console.
- You must have the root password for sudo access to the commands.



For more information about installing Cumulus Linux, see [Software install workflow for NVIDIA SN2100 switches](#).

Example 5. Steps

Cumulus Linux 5.10 and earlier

1. From the Cumulus console, download and queue the switch software installation with the command `onie-install -a -i` followed by the file path to the switch software, for example:

```
cumulus@sw1:mgmt:~$ sudo onie-install -a -i http://<web-server>/<path>/cumulus-linux-5.10.0-mlx-amd64.bin
```

2. The installer starts the download. Type **y** when prompted to confirm the installation when the image is downloaded and verified.
3. Reboot the switch to install the new software.

```
sudo reboot
```

```
cumulus@sw1:mgmt:~$ sudo reboot
```



The switch reboots and enters the switch software installation which takes some time. When the installation is complete, the switch reboots and remains at the `log-in` prompt.

Cumulus Linux 5.11 and later

1. To reset the switch to the factory defaults and remove all configuration, system files, and log files, run:

```
nv action reset system factory-default
```

For example:

```
cumulus@switch:~$ nv action reset system factory-default
```

```
This operation will reset the system configuration, delete the log files and reboot the switch.
```

```
Type [y] continue.
```

```
Type [n] to abort.
```

```
Do you want to continue? [y/n] y
```

See the NVIDIA [Factory Reset](#) documentation for further details.

What's next

After you've reset your switches, you can [reconfigure](#) them as needed.

Migrate the switches

Migrate CN1610 cluster switches to NVIDIA SN2100 cluster switches

You can migrate NetApp CN1610 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 NetApp CN1610 cluster switches with NVIDIA SN2100 cluster switches. See [Overview of installation and configuration for NVIDIA SN2100 switches](#).

Supported switches

The following cluster switches are supported:

- NetApp CN1610
- NVIDIA SN2100

For details of supported ports and their configurations, see the [Hardware Universe](#).

Before you begin

Verify that you meet the following requirements for your configuration:

- The existing cluster is correctly 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 correct version of Cumulus Linux installed with the reference configuration file (RCF) applied.
- The existing cluster network configuration has the following:
 - A redundant and fully functional NetApp cluster using CN1610 switches.
 - Management connectivity and console access to both the CN1610 switches and the new switches.
 - All cluster LIFs in the up state with the cluster LIFs on their home ports.
 - ISL ports enabled and cabled between the CN1610 switches and between the new switches.
- Some of the ports are configured on NVIDIA SN2100 switches to run at 40GbE or 100GbE.
- You have planned, migrated, and documented 40GbE and 100GbE connectivity from nodes to NVIDIA SN2100 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 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.
 - The cabling between the nodes and `c2` is 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.
 - The cabling between the nodes and `c1` is then disconnected from `c1` and reconnected to `sw1`.



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.

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
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	node1_clus2	up/up	169.254.49.125/16	node1
e3a	node2_clus1	up/up	169.254.47.194/16	node2
e3b	node2_clus2	up/up	169.254.19.183/16	node2

- 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	0/1	124	H	AFF-A400
node2 e3a	0/2	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	0/1	124	H	AFF-A400
node2 e3b	0/2	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.

```

(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
(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:

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

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)	0/1	-
	e3b	sw2 (b8:ce:f6:19:1a:7e)	swp3	-
node2	/lldp			
	e3a	c1 (6a:ad:4f:98:3b:3f)	0/2	-
	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
(c1) (Config) # interface 0/1-0/12
(c1) (Interface 0/1-0/12) # shutdown
(c1) (Interface 0/1-0/12) # 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	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 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. On switch sw2, shut down and restart all cluster ports to trigger an auto-revert of all cluster LIFs that are not on their home ports.

Cumulus 4.4.3

```
cumulus@sw2:mgmt:~$ net add interface swp1-14 link down
cumulus@sw2:mgmt:~$ net pending
cumulus@sw2:mgmt:~$ net commit
```

(Wait for 5-10 seconds before re-enabling the ports)

```
cumulus@sw2:mgmt:~$ net add interface swp1-14 link up
cumulus@sw2:mgmt:~$ net pending
cumulus@sw2:mgmt:~$ net commit
```

(After executing the link state up command, the nodes detect the change and begin to auto-revert the cluster LIFs to their home ports)

Cumulus 5.x

```
cumulus@sw2:mgmt:~$ nv set interface swp1-14 link state down
cumulus@sw2:mgmt:~$ nv config apply
cumulus@sw2:mgmt:~$ nv show interface
```

(Wait for 5-10 seconds before re-enabling the ports)

```
cumulus@sw2:mgmt:~$ nv set interface swp1-14 link state up
cumulus@sw2:mgmt:~$ nv config apply
cumulus@sw2:mgmt:~$ nv show interface
```

(After executing the link state up command, the nodes detect the change and begin to auto-revert the cluster LIFs to their home ports)

3. Verify that the cluster LIFs have reverted to their home ports (this might take a minute):

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

If any of the cluster LIFs have not reverted to their home port, manually revert them. You must connect to each node-mgmt LIF or SP/BMC system console of the local node that owns the LIF:

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

4. Change the privilege level back to admin:

```
set -privilege admin
```

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
```

What's next?

After you've migrated your switches, you can [configure switch health monitoring](#).

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      Speed (Mbps)
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      Speed (Mbps)
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 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
```

						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
```

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. On switch sw2, shut down and restart all cluster ports to trigger an auto-revert of all cluster LIFs that are not on their home ports.

Cumulus 4.4.3

```
cumulus@sw2:mgmt:~$ net add interface swp1-14 link down
cumulus@sw2:mgmt:~$ net pending
cumulus@sw2:mgmt:~$ net commit
```

(Wait for 5-10 seconds before re-enabling the ports)

```
cumulus@sw2:mgmt:~$ net add interface swp1-14 link up
cumulus@sw2:mgmt:~$ net pending
cumulus@sw2:mgmt:~$ net commit
```

(After executing the link state up command, the nodes detect the change and begin to auto-revert the cluster LIFs to their home ports)

Cumulus 5.x

```
cumulus@sw2:mgmt:~$ nv set interface swp1-14 link state down
cumulus@sw2:mgmt:~$ nv config apply
cumulus@sw2:mgmt:~$ nv show interface
```

(Wait for 5-10 seconds before re-enabling the ports)

```
cumulus@sw2:mgmt:~$ nv set interface swp1-14 link state up
cumulus@sw2:mgmt:~$ nv config apply
cumulus@sw2:mgmt:~$ nv show interface
```

(After executing the link state up command, the nodes detect the change and begin to auto-revert the cluster LIFs to their home ports)

3. Verify that the cluster LIFs have reverted to their home ports (this might take a minute):

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

If any of the cluster LIFs have not reverted to their home port, manually revert them. You must connect to each node-mgmt LIF or SP/BMC system console of the local node that owns the LIF:

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

4. Change the privilege level back to admin:

```
set -privilege admin
```

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
```

What's next?

After you've migrated your switches, you can [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	oss-g-int-rcf10
swp15					
swp16	up	100G	9216	swp	oss-g-int-rcf10
swp16					

10. Verify that all cluster ports are up:

```
network port show -ip space 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
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
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. On switch sw2, shut down and restart all cluster ports to trigger an auto-revert of all cluster LIFs that are not on their home ports.

Cumulus 4.4.3

```
cumulus@sw2:mgmt:~$ net add interface swp1-14 link down
cumulus@sw2:mgmt:~$ net pending
cumulus@sw2:mgmt:~$ net commit
```

(Wait for 5-10 seconds before re-enabling the ports)

```
cumulus@sw2:mgmt:~$ net add interface swp1-14 link up
cumulus@sw2:mgmt:~$ net pending
cumulus@sw2:mgmt:~$ net commit
```

(After executing the link state up command, the nodes detect the change and begin to auto-revert the cluster LIFs to their home ports)

Cumulus 5.x

```
cumulus@sw2:mgmt:~$ nv set interface swp1-14 link state down
cumulus@sw2:mgmt:~$ nv config apply
cumulus@sw2:mgmt:~$ nv show interface
```

(Wait for 5-10 seconds before re-enabling the ports)

```
cumulus@sw2:mgmt:~$ nv set interface swp1-14 link state up
cumulus@sw2:mgmt:~$ nv config apply
cumulus@sw2:mgmt:~$ nv show interface
```

(After executing the link state up command, the nodes detect the change and begin to auto-revert the cluster LIFs to their home ports)

3. Verify that the cluster LIFs have reverted to their home ports (this might take a minute):

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

If any of the cluster LIFs have not reverted to their home port, manually revert them. You must connect to each node-mgmt LIF or SP/BMC system console of the local node that owns the LIF:

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

4. 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
```

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					

5. Verify that the settings are disabled:

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

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 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)

```

8. Change the privilege level back to admin:

```
set -privilege admin
```

9. 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?

After you've migrated your switches, you can [configure switch health monitoring](#).

Replace the 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
```

						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							

```
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

```

```
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

Cumulus Linux 4.4.3

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.

```
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
```

```
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 disabled:

```
net interface show -vserver Cluster -fields auto-revert
```

4. Shut down the ISL ports swp15 and swp16 on the SN2100 switch sw1.

```
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.

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 -ipospace Cluster
```

The output should be similar to the following:

```
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					

8. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

```

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
```

```

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
```

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.

```
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
```

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
```

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

Cumulus Linux 5.x

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.

```
cumulus@nsw2:~$ nv set interface swp15-16 link state down
cumulus@nsw2:~$ nv config apply
```

2. Disable auto-revert on the cluster LIFs:

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

```
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 disabled:

```
network interface show -vserver Cluster -fields auto-revert
```

4. Shut down the ISL ports swp15 and swp16 on the SN2100 switch sw1.

```
cumulus@sw1:~$ nv set interface swp15-16 link state down
cumulus@sw1:~$ nv config apply
```

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.

The following commands enable ISL ports swp15 and swp16 on switch sw1:

```
cumulus@sw1:~$ nv set interface swp15-16 link state down
cumulus@sw1:~$ nv config apply
```

The following example shows that the ISL ports are up on switch sw1:

```
cumulus@sw1:~$ nv 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:~$ nv 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
```

The output should be similar to the following:

```
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					

8. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

```

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:

```
nv show interface
```

```

cumulus@nsw2:~$ nv 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:

```
nv show interface lldp
```

The following example shows the appropriate results for both switches:

```
cumulus@sw1:~$ nv show interface 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:~$ nv show interface 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.

```
cumulus@nsw2:~$ nv set interface swp1-14 link state up  
cumulus@nsw2:~$ nv config apply
```

13. Display information about the nodes in a cluster:

```
cluster show
```

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
```

```
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

Cumulus Linux 4.4.3

1. Verify that the cluster network is healthy.

```
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 Linux 5.x

1. Verify that the cluster network is healthy.

```
cumulus@sw1:~$ nv show interface 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?

After you've replaced your switches, you can [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.

Before you begin

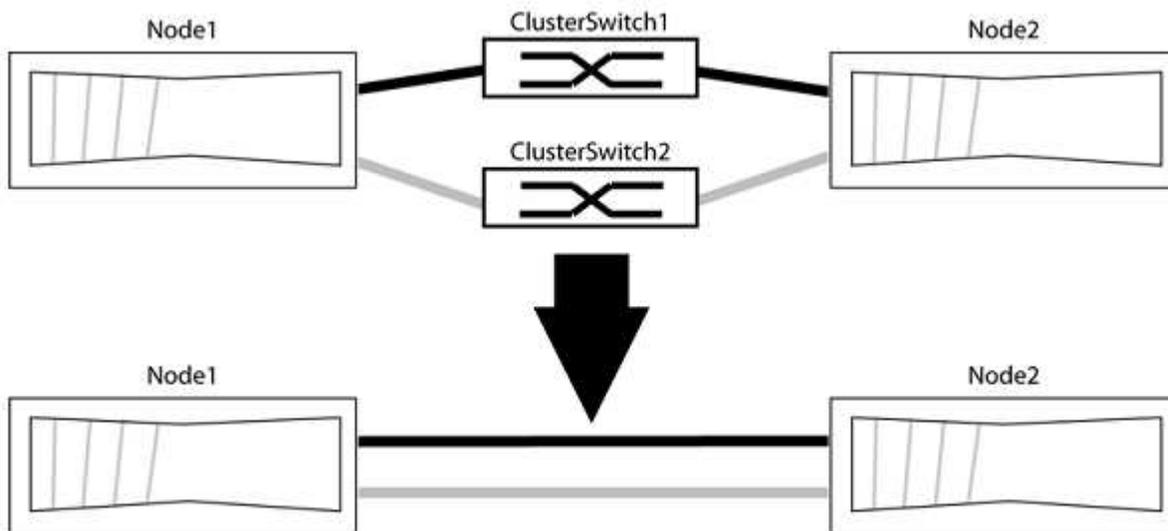
Make sure you have the following:

- 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

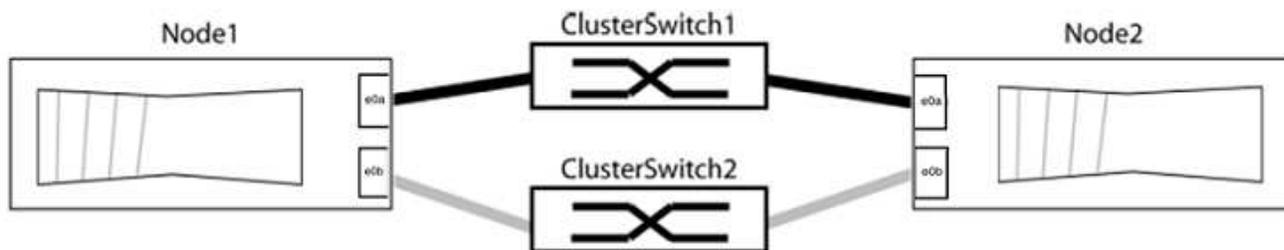
```
cluster::*> system node autosupport invoke -node * -type all
-message MAINT=2h
```

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.

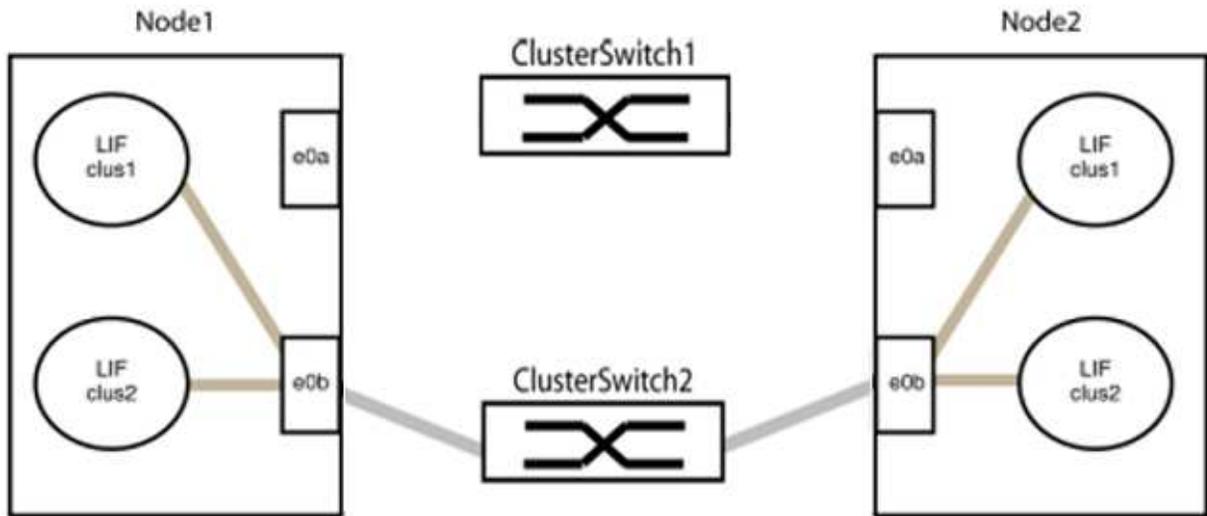
8. Set up the switchless configuration for the ports in group 1.



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**.

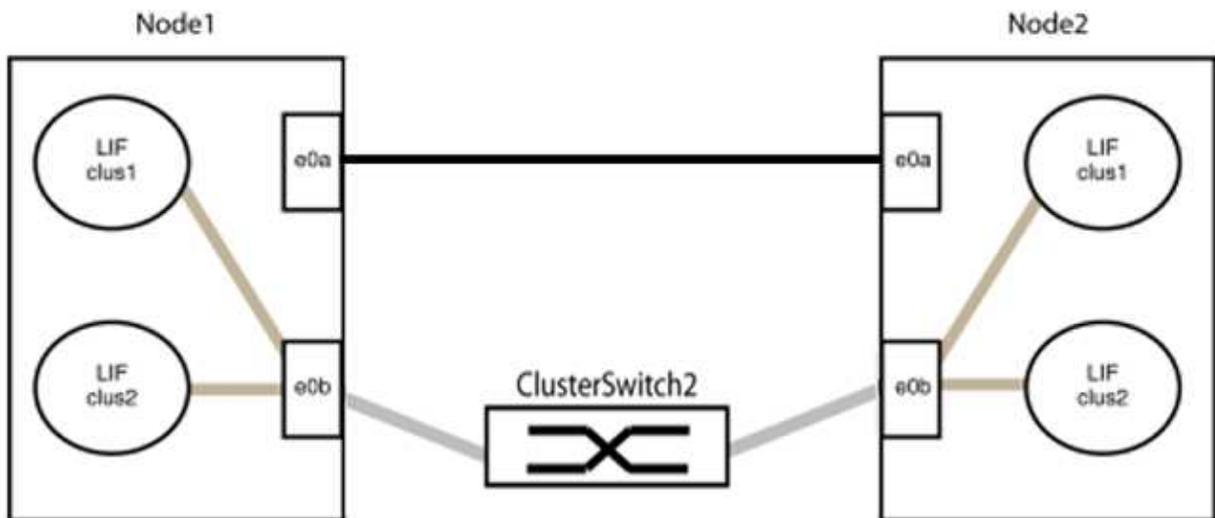
a. Disconnect all the cables from the ports in group1 at the same time.

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.

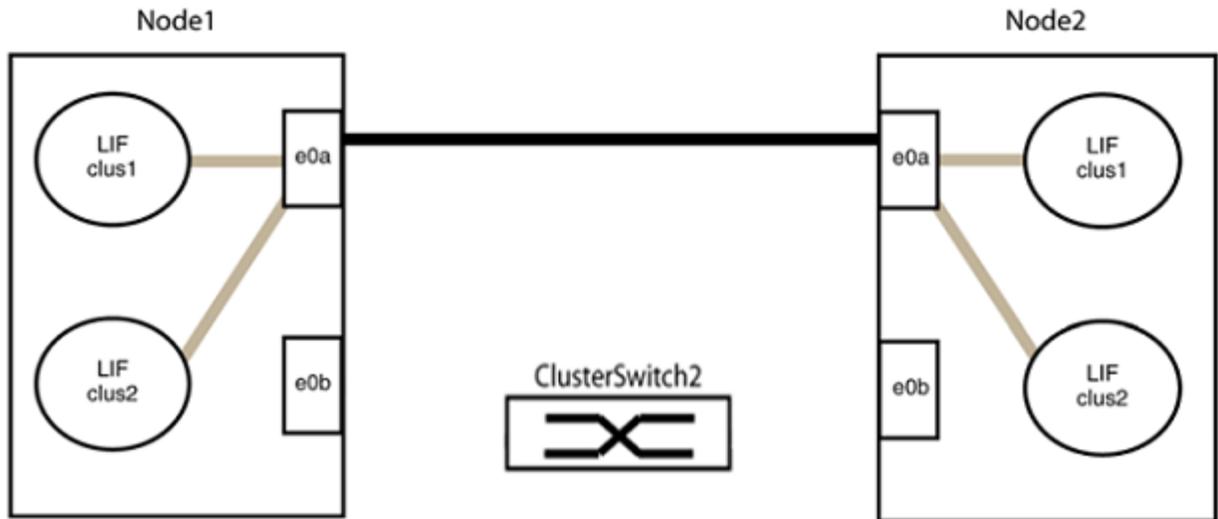
11. Set up the switchless configuration for the ports in group 2.



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**.

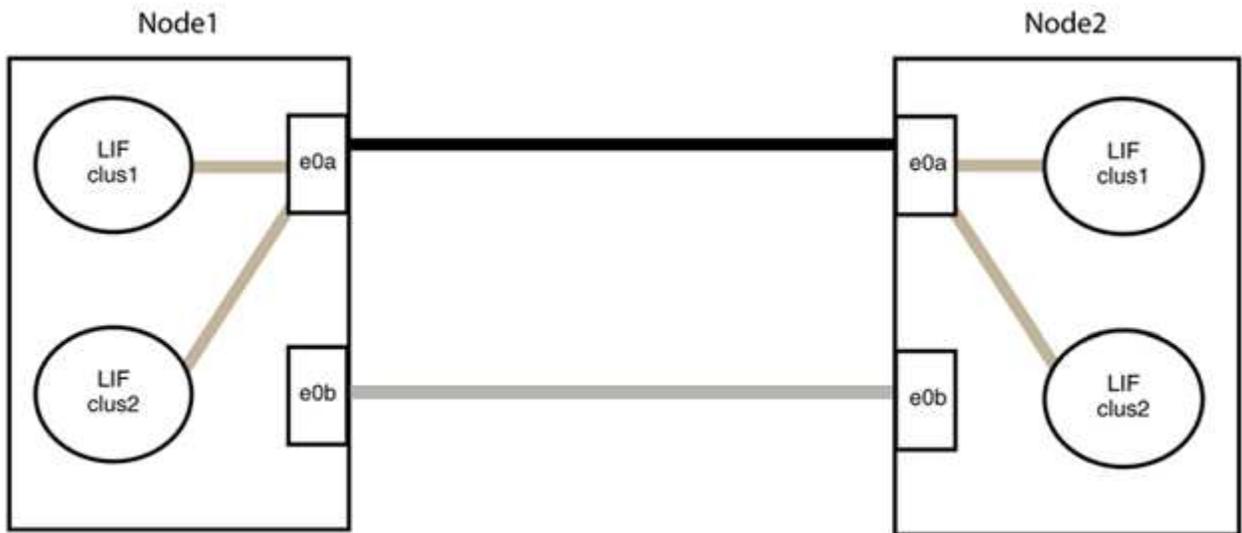
- a. Disconnect all the cables from the ports in group2 at the same time.

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
```

What's next?

After you've replaced your switches, you can [configure switch health monitoring](#).

Copyright information

Copyright © 2026 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.