



Broadcom-supported BES-53248

Install and maintain

NetApp

February 13, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap-systems-switches/switch-bes-53248/configure-new-switch-overview.html> on February 13, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Broadcom-supported BES-53248	1
Get started	1
Installation and setup workflow for BES-53248 switches	1
Configuration requirements for BES-53248 cluster switches	1
Components and part numbers for BES-53248 cluster switches	2
Documentation requirements for BES-53248 cluster switches	4
Install the hardware	5
Hardware install workflow for BES-53248 switches	5
Install the hardware for the BES-53248 cluster switch	6
Review cabling and configuration considerations	6
Configure the software	8
Software install workflow for BES-53248 switches	8
Configure the BES-53248 cluster switch	9
Install the EFOS software	14
Install the Reference Configuration File (RCF) and license file	23
Install licenses for BES-53248 cluster switches	55
Install the Reference Configuration File (RCF)	68
Enable SSH on BES-53248 cluster switches	101
Reset the BES-53248 cluster switch to factory defaults	105
Upgrade the switch	106
Upgrade workflow for BES-53248 cluster switches	106
Upgrade the EFOS software	106
Upgrade the Reference Configuration File (RCF)	118
Verify the ONTAP cluster network after an EFOS software or RCF upgrade of the BES-53248 cluster switches	148
Migrate the switches	150
Migrate CN1610 cluster switches to BES-53248 cluster switches	150
Migrate to a switched NetApp cluster environment	170
Replace the switches	187
Replacement requirements	187
Replace a Broadcom-supported BES-53248 cluster switch	188
Replace Broadcom BES-53248 cluster switches with switchless connections	199

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 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-53248 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 licensed 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:



[Install the switch hardware](#)

Install and configure the BES-53248 switch hardware.

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			
write	read		write	read		link status
	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			
write	read		write	read		link status
	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:

```
systemshell -node <node> -command sudo sysctl  
dev.device.<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.

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) # sntp 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.

Show example

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

Show example

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

The switch boots into ONIE installation mode.

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

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

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



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

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

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

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

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

```
ping
```

Show example

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

6. Install the new switch software:

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

Show example

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

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

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

```
show bootvar
```

Show example

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

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

- Configure switch
- Install licenses
- Install the RCF file
- 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                                         Speed (Mbps)
Health   Health
Port     IPspace      Broadcast Domain Link MTU Admin/Oper
Status   Status
-----  -----
-----  -----
e0a     Cluster      Cluster          up    9000  auto/100000
healthy false
e0b     Cluster      Cluster          up    9000  auto/100000
healthy false

Node: cluster1-02

Ignore                                         Speed (Mbps)
Health   Health
Port     IPspace      Broadcast Domain Link MTU Admin/Oper
Status   Status
-----  -----
-----  -----
e0a     Cluster      Cluster          up    9000  auto/100000
healthy false
e0b     Cluster      Cluster          up    9000  auto/100000
healthy false
8 entries were displayed.

Node: cluster1-03

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

```

Node: cluster1-04

Ignore

          Speed (Mbps)

Health   Health
Port      IPspace      Broadcast Domain Link MTU Admin/Oper
Status   Status
-----  -----
-----  -----
e0a      Cluster      Cluster          up    9000  auto/10000
healthy  false
e0b      Cluster      Cluster          up    9000  auto/10000
healthy  false
cluster1:>*:>

```

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

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

Show example

```
cluster1::*> network interface show -vserver Cluster
      Logical          Status      Network
  Current      Current  Is
  Vserver      Interface
  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  e0b      true
      cluster1-02_clus1  up/up    169.254.3.8/23
cluster1-02  e0a      true
      cluster1-02_clus2  up/up    169.254.3.9/23
cluster1-02  e0b      true
      cluster1-03_clus1  up/up    169.254.1.3/23
cluster1-03  e0a      true
      cluster1-03_clus2  up/up    169.254.1.1/23
cluster1-03  e0b      true
      cluster1-04_clus1  up/up    169.254.1.6/23
cluster1-04  e0a      true
      cluster1-04_clus2  up/up    169.254.1.7/23
cluster1-04  e0b      true
```

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

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command:

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

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

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

```
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command:

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

```

cluster1::*> system cluster-switch show -is-monitoring-enabled
- operational true
Switch Type Address Model
-----
-----
cs1 cluster-network 10.228.143.200 BES-
53248
    Serial Number: QTWCU22510008
    Is Monitored: true
    Reason: None
    Software Version: 3.10.0.3
    Version Source: CDP/ISDP

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

```

6. Disable auto-revert on the cluster LIFs.

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

Step 2: Configure ports

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

```
show isdp neighbor
```

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

```

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

```

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

Link      Physical      Physical
Media     Flow
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	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	cluster-network	10.228.143.200	BES-
53248			
	Serial Number: QTWCU22510008		
	Is Monitored: true		
	Reason: None		
	Software Version: 3.10.0.3		
	Version Source: CDP/ISDP		
cs2	cluster-network	10.228.143.202	BES-
53248			
	Serial Number: QTWCU22510009		
	Is Monitored: true		
	Reason: None		
	Software Version: 3.10.0.3		
	Version Source: CDP/ISDP		

```
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command:

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

```

cluster1::*> system cluster-switch show -is-monitoring-enabled
- operational true
Switch Type Address Model
-----
-----
cs1 cluster-network 10.228.143.200 BES-
53248
    Serial Number: QTWCU22510008
    Is Monitored: true
    Reason: None
    Software Version: 3.10.0.3
    Version Source: CDP/ISDP

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

```

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_clus1  up/up      169.254.3.4/23
cluster1-01      e0a      false
      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      e0a      false
      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      e0a      false
      cluster1-03_clus2  up/up      169.254.1.1/23
cluster1-03      e0b      true
      cluster1-04_clus1  up/up      169.254.1.6/23
cluster1-04      e0a      false
      cluster1-04_clus2  up/up      169.254.1.7/23
cluster1-04      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

Link      Physical      Physical
Media     Flow
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_clus1  up/up      169.254.3.4/23
cluster1-01      e0a      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      e0a      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      e0a      true
      cluster1-03_clus2  up/up      169.254.1.1/23
cluster1-03      e0b      true
      cluster1-04_clus1  up/up      169.254.1.6/23
cluster1-04      e0a      true
      cluster1-04_clus2  up/up      169.254.1.7/23
cluster1-04      e0b      true
```

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
                                         Source          Destination
Packet
Node    Date          LIF          LIF
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 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	Intf	Type	Admin Mode	Physical Mode	Physical Status	Link Status	Link Trap	LACP 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							
	0/49		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                                         Speed (Mbps)
Health   Health
Port     IPspace      Broadcast Domain Link MTU Admin/Oper
Status   Status
-----  -----
-----  -----
e0a     Cluster      Cluster          up    9000  auto/100000
healthy false
e0b     Cluster      Cluster          up    9000  auto/100000
healthy false

Node: cluster1-02

Ignore                                         Speed (Mbps)
Health   Health
Port     IPspace      Broadcast Domain Link MTU Admin/Oper
Status   Status
-----  -----
-----  -----
e0a     Cluster      Cluster          up    9000  auto/100000
healthy false
e0b     Cluster      Cluster          up    9000  auto/100000
healthy false
8 entries were displayed.

Node: cluster1-03

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

```

Node: cluster1-04

Ignore

          Speed (Mbps)

Health   Health
Port      IPspace      Broadcast Domain Link MTU Admin/Oper
Status   Status
-----  -----
-----  -----
e0a      Cluster      Cluster          up    9000  auto/10000
healthy  false
e0b      Cluster      Cluster          up    9000  auto/10000
healthy  false
cluster1:>*:>

```

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

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

Show example

```
cluster1::*> network interface show -vserver Cluster
      Logical          Status      Network
  Current      Current  Is
  Vserver      Interface
  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  e0b      true
      cluster1-02_clus1  up/up      169.254.3.8/23
cluster1-02  e0a      true
      cluster1-02_clus2  up/up      169.254.3.9/23
cluster1-02  e0b      true
      cluster1-03_clus1  up/up      169.254.1.3/23
cluster1-03  e0a      true
      cluster1-03_clus2  up/up      169.254.1.1/23
cluster1-03  e0b      true
      cluster1-04_clus1  up/up      169.254.1.6/23
cluster1-04  e0a      true
      cluster1-04_clus2  up/up      169.254.1.7/23
cluster1-04  e0b      true
```

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

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command:

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

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

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

```
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command:

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

```

cluster1::*> system cluster-switch show -is-monitoring-enabled
- operational true
Switch Type Address Model
-----
-----
cs1 cluster-network 10.228.143.200 BES-
53248
    Serial Number: QTWCU22510008
    Is Monitored: true
    Reason: None
    Software Version: 3.10.0.3
    Version Source: CDP/ISDP

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

```

6. Disable auto-revert on the cluster LIFs.

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

Step 2: Configure ports

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

```
show isdp neighbor
```

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

```

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

```

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 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

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 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.
```

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
2 configuration script(s) found.
2042 Kbytes free.
```

10. 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
```

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
```

```
(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

Link      Physical      Physical
Media     Flow
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

          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

```

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

```
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command:

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

```

cluster1::*> system cluster-switch show -is-monitoring-enabled
- operational true
Switch Type Address Model
-----
-----
cs1 cluster-network 10.228.143.200 BES-
53248
    Serial Number: QTWCU22510008
    Is Monitored: true
    Reason: None
    Software Version: 3.10.0.3
    Version Source: CDP/ISDP

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

```

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_clus1  up/up      169.254.3.4/23
cluster1-01      e0a      false
      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      e0a      false
      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      e0a      false
      cluster1-03_clus2  up/up      169.254.1.1/23
cluster1-03      e0b      true
      cluster1-04_clus1  up/up      169.254.1.6/23
cluster1-04      e0a      false
      cluster1-04_clus2  up/up      169.254.1.7/23
cluster1-04      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

Link      Physical      Physical
Media     Flow
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_clus1  up/up      169.254.3.4/23
cluster1-01      e0a      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      e0a      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      e0a      true
      cluster1-03_clus2  up/up      169.254.1.1/23
cluster1-03      e0b      true
      cluster1-04_clus1  up/up      169.254.1.6/23
cluster1-04      e0a      true
      cluster1-04_clus2  up/up      169.254.1.7/23
cluster1-04      e0b      true
```

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
                                         Source          Destination
Packet
Node    Date          LIF          LIF
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
<passphase>
```

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
<passphase>
```

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
<passphase>
```

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	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
```

12. Repeat steps 5 through to 11 on the switch cs1.

13. Enable auto-revert on the cluster LIFs.

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

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

```
network interface show -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

1. Disable auto-revert on the cluster LIFs.

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

2. Boot the switch into ONIE installation mode.

During boot, select ONIE when you see the prompt:

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

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

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

The switch boots into ONIE installation mode.

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

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

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



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

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

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

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

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

```
ping
```

Show example

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

6. Install the new switch software:

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

Show example

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

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

7. Verify that the new switch software is installed:

```
show bootvar
```

Show example

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

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

- [Install licenses](#)
- [Install the RCF](#)
- [Enable SSH](#)
- [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
                                         Speed (Mbps)
Health   Health
Port     IPspace      Broadcast Domain Link MTU Admin/Oper
Status   Status
----- -----
----- 
e0a      Cluster      Cluster          up    9000  auto/100000
healthy  false
e0b      Cluster      Cluster          up    9000  auto/100000
healthy  false

Node: cluster1-02

Ignore
                                         Speed (Mbps)
Health   Health
Port     IPspace      Broadcast Domain Link MTU Admin/Oper
Status   Status
----- -----
----- 
e0a      Cluster      Cluster          up    9000  auto/100000
healthy  false
e0b      Cluster      Cluster          up    9000  auto/100000
healthy  false
8 entries were displayed.

Node: cluster1-03

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

```
Node: cluster1-04
```

```
Ignore
```

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

```
cluster1::*>
```

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

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

Show example

```
cluster1::*> network interface show -vserver Cluster
      Logical          Status      Network
  Current      Current  Is
  Vserver      Interface
  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  e0b      true
      cluster1-02_clus1  up/up      169.254.3.8/23
cluster1-02  e0a      true
      cluster1-02_clus2  up/up      169.254.3.9/23
cluster1-02  e0b      true
      cluster1-03_clus1  up/up      169.254.1.3/23
cluster1-03  e0a      true
      cluster1-03_clus2  up/up      169.254.1.1/23
cluster1-03  e0b      true
      cluster1-04_clus1  up/up      169.254.1.6/23
cluster1-04  e0a      true
      cluster1-04_clus2  up/up      169.254.1.7/23
cluster1-04  e0b      true
```

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

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command:

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

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

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

```
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command:

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

```

cluster1::*> system cluster-switch show -is-monitoring-enabled
- operational true
Switch Type Address Model
-----
-----
cs1 cluster-network 10.228.143.200 BES-
53248
    Serial Number: QTWCU22510008
    Is Monitored: true
    Reason: None
    Software Version: 3.10.0.3
    Version Source: CDP/ISDP

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

```

6. Disable auto-revert on the cluster LIFs.

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

Step 2: Configure ports

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

```
show isdp neighbor
```

2. On 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

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/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.
```

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
2 configuration script(s) found.
2042 Kbytes free.
```

10. 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 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                                         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	Admin/Oper
Status	Status				
<hr/>					
<hr/>					
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	cluster-network	10.228.143.200	BES-
53248			
	Serial Number: QTWCU22510008		
	Is Monitored: true		
	Reason: None		
	Software Version: 3.10.0.3		
	Version Source: CDP/ISDP		
cs2	cluster-network	10.228.143.202	BES-
53248			
	Serial Number: QTWCU22510009		
	Is Monitored: true		
	Reason: None		
	Software Version: 3.10.0.3		
	Version Source: CDP/ISDP		

```
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command:

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

```

cluster1::*> system cluster-switch show -is-monitoring-enabled
- operational true
Switch Type Address Model
-----
-----
cs1 cluster-network 10.228.143.200 BES-
53248
    Serial Number: QTWCU22510008
    Is Monitored: true
    Reason: None
    Software Version: 3.10.0.3
    Version Source: CDP/ISDP

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

```

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

Link      Physical      Physical
Media     Flow
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_clus1  up/up      169.254.3.4/23
cluster1-01      e0a      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      e0a      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      e0a      true
      cluster1-03_clus2  up/up      169.254.1.1/23
cluster1-03      e0b      true
      cluster1-04_clus1  up/up      169.254.1.6/23
cluster1-04      e0a      true
      cluster1-04_clus2  up/up      169.254.1.7/23
cluster1-04      e0b      true
```

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

cluster1-02
      3/5/2022 19:21:18 -06:00  cluster1-02_clus2  cluster1-02_clus1
none
      3/5/2022 19:21:20 -06:00  cluster1-02_clus2  cluster1-02_clus2
none
```

All ONTAP releases

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

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

```

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

```

6. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

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

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

Steps

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

```
network port show -ipspace Cluster
```

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

Show example

The following example shows the output from the command:

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

Node: node1

Ignore

Health
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
```

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

          Logical      Status      Network      Current
Current Is
Vserver      Interface  Admin/Oper  Address/Mask      Node
Port      Home
-----  -----  -----  -----  -----
-----  -----
Cluster
      node1_clus1  up/up    169.254.217.125/16 node1
e0a      true
      node1_clus2  up/up    169.254.205.88/16 node1
e0b      true
      node2_clus1  up/up    169.254.252.125/16 node2
e0a      true
      node2_clus2  up/up    169.254.110.131/16 node2
e0b      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/      Local   Discovered
Protocol   Port    Device (LLDP: ChassisID)  Interface
Platform

-----
-----
node2      /cdp
          e0a    CL1          0/2
CN1610
          e0b    CL2          0/2
CN1610
node1      /cdp
          e0a    CL1          0/1
CN1610
          e0b    CL2          0/1
CN1610
```

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

a. Verify that all the cluster ports are up with a healthy status:

```
network port show -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
```

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
                                         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. 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
          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
e0a      false
      node2_clus1  up/up      169.254.47.194/16  node2
e0a      true
      node2_clus2  up/up      169.254.19.183/16  node2
e0a      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

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
```

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
CN1610
          e0b    cs2
          0/2
          BES-
53248
node1      /cdp
          e0a    CL1
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
          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
```

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

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
```

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
53248
            e0b    cs2
53248
node2      /cdp
            e0a    cs1
53248
            e0b    cs2
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
                                         Source          Destination
Packet
Node    Date          LIF          LIF
Loss
-----
-----
node1
      3/5/2022 19:21:18 -06:00  node1_clus2      node2_clus1
none
      3/5/2022 19:21:20 -06:00  node1_clus2      node2_clus2
none

node2
      3/5/2022 19:21:18 -06:00  node2_clus2      node1_clus1
none
      3/5/2022 19:21:20 -06:00  node2_clus2      node1_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 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
```

```

cluster::*> 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)	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		

5. Verify that all cluster LIFs are up and operational:

```
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
      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

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
```

11. Verify that all cluster LIFs are up and operational:

```
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
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
```

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

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
```

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
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
```

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
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
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
                                         Source          Destination
Packet
Node    Date           LIF           LIF
Loss
-----
-----
node1
      3/5/2022 19:21:18 -06:00  node1_clus2      node2_clus1
none
      3/5/2022 19:21:20 -06:00  node1_clus2      node2_clus2
none

node2
      3/5/2022 19:21:18 -06:00  node2_clus2      node1_clus1
none
      3/5/2022 19:21:20 -06:00  node2_clus2      node1_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 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
  -message 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

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

Node: node2

Ignore

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

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

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	-----
-----	-----	-----	-----	-----	-----
Cluster	node1_clus1	up/up	169.254.209.69/16	node1	e0a
true	node1_clus2	up/up	169.254.49.125/16	node1	e0b
true					

```

true          node2_clus1  up/up      169.254.47.194/16  node2          e0a
true          node2_clus2  up/up      169.254.19.183/16  node2          e0b

cluster1::> network device-discovery show -protocol cdp
Node/          Local   Discovered
Protocol      Port    Device (LLDP: ChassisID)  Interface          Platform
-----  -----  -----  -----  -----  -----
-----  -----
node2          /cdp
53248          e0a    cs1           0/2          BES-
53248          e0b    cs2           0/2          BES-
node1          /cdp
53248          e0a    cs1           0/1          BES-
53248          e0b    cs2           0/1          BES-

```

```
(cs1) # show isdp neighbors
```

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge,
S - Switch, H - Host, I - IGMP, r - Repeater

Device ID Port ID	Intf	Holdtime	Capability	Platform
node1 e0a	0/1	175	H	FAS2750
node2 e0a	0/2	152	H	FAS2750
cs2 0/55	0/55	179	R	BES-53248
cs2 0/56	0/56	179	R	BES-53248

```
(cs2) # show isdp neighbors
```

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge,
S - Switch, H - Host, I - IGMP, r - Repeater

Device ID Port ID	Intf	Holdtime	Capability	Platform
node1 e0b	0/1	129	H	FAS2750
node2 e0b	0/2	165	H	FAS2750
cs1 0/55	0/55	179	R	BES-53248
cs1 0/56	0/56	179	R	BES-53248

Steps

1. Review the [Replacement requirements](#).
2. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

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

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



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

3. Install the appropriate Reference Configuration File (RCF) and image on the switch, newcs2, and make any necessary site preparations.

If necessary, verify, download, and install the appropriate versions of the RCF and EFOS software for the new switch. If you have verified that the new switch is correctly set up and does not need updates to the RCF and EFOS software, continue to step 2.

- a. You can download the applicable Broadcom EFOS software for your cluster switches from the [Broadcom Ethernet Switch Support](#) site. Follow the steps on the Download page to download the EFOS file for the version of ONTAP software you are installing.
- b. The appropriate RCF is available from the [Broadcom Cluster Switches](#) page. Follow the steps on the Download page to download the correct RCF for the version of ONTAP software you are installing.

4. On the new switch, log in as `admin` and shut down all of the ports that will be connected to the node cluster interfaces (ports 1 to 16).



If you purchased additional licenses for additional ports, shut down these ports too.

If the switch that you are replacing is not functional and is powered down, the LIFs on the cluster nodes should have already failed over to the other cluster port for each node.



No password is required to enter `enable` mode.

Show example

```
User: admin
Password:
(newcs2) > enable
(newcs2) # config
(newcs2) (config) # interface 0/1-0/16
(newcs2) (interface 0/1-0/16) # shutdown
(newcs2) (interface 0/1-0/16) # exit
(newcs2) (config) # exit
(newcs2) #
```

5. Verify that all cluster LIFs have auto-revert enabled:

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

Show example topology

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

Logical
Vserver      Interface      Auto-revert
-----  -----
Cluster      node1_clus1    true
Cluster      node1_clus2    true
Cluster      node2_clus1    true
Cluster      node2_clus2    true
```

6. Shut down the ISL ports 0/55 and 0/56 on the BES-53248 switch cs1:

Show example topology

```
(cs1)# config
(cs1)(config)# interface 0/55-0/56
(cs1)(interface 0/55-0/56)# shutdown
```

7. Remove all cables from the BES-53248 cs2 switch, and then connect them to the same ports on the BES-53248 newcs2 switch.
8. Bring up the ISLs ports 0/55 and 0/56 between the cs1 and newcs2 switches, and then verify the port channel operation status.

The Link State for port-channel 1/1 should be **up** and all member ports should be True under the Port Active heading.

Show example

This example enables ISL ports 0/55 and 0/56 and displays the Link State for port-channel 1/1 on switch cs1:

```
(cs1) # config
(cs1) (config) # interface 0/55-0/56
(cs1) (interface 0/55-0/56) # no shutdown
(cs1) (interface 0/55-0/56) # exit
(cs1) # show port-channel 1/1

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

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

9. On the new switch newcs2, re-enable all of the ports that are connected to the node cluster interfaces (ports 1 to 16).



If you purchased additional licenses for additional ports, shut down these ports too.

Show example

```
User:admin
Password:
(newcs2) > enable
(newcs2) # config
(newcs2) (config) # interface 0/1-0/16
(newcs2) (interface 0/1-0/16) # no shutdown
(newcs2) (interface 0/1-0/16) # exit
(newcs2) (config) # exit
```

10. Verify that port e0b is **up**:

```
network port show -ipspace Cluster
```

Show example

The output should be similar to the following:

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

Node: node1

Ignore                                         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/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

          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
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
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
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

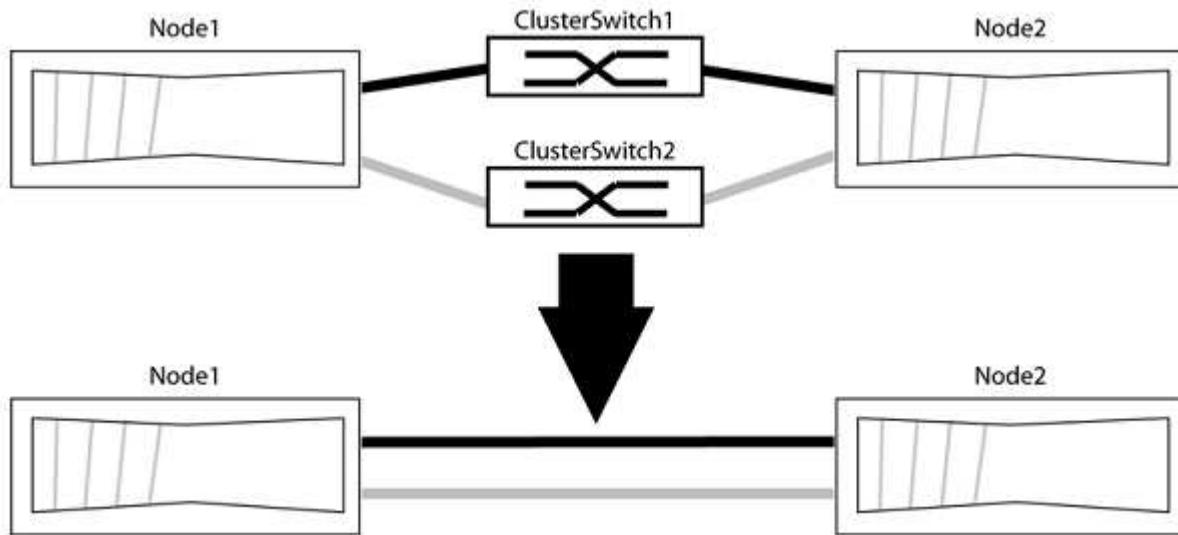
Make sure you have the following:

- 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

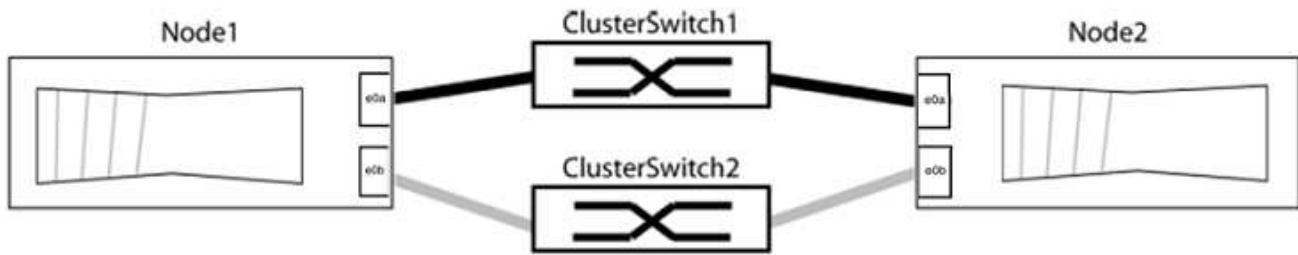
```
cluster::*> system node autosupport invoke -node * -type all
-message MAINT=2h
```

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
                                                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
      e0b    cs2
node2/cdp
      e0a    cs1
      e0b    cs2
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
                                         Source          Destination
Packet
Node    Date          LIF          LIF
Loss
-----
-----
node1
  3/5/2022 19:21:18 -06:00  node1_clus2      node2-clus1
none
  3/5/2022 19:21:20 -06:00  node1_clus2      node2_clus2
none
node2
  3/5/2022 19:21:18 -06:00  node2_clus2      node1_clus1
none
  3/5/2022 19:21:20 -06:00  node2_clus2      node1_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 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.

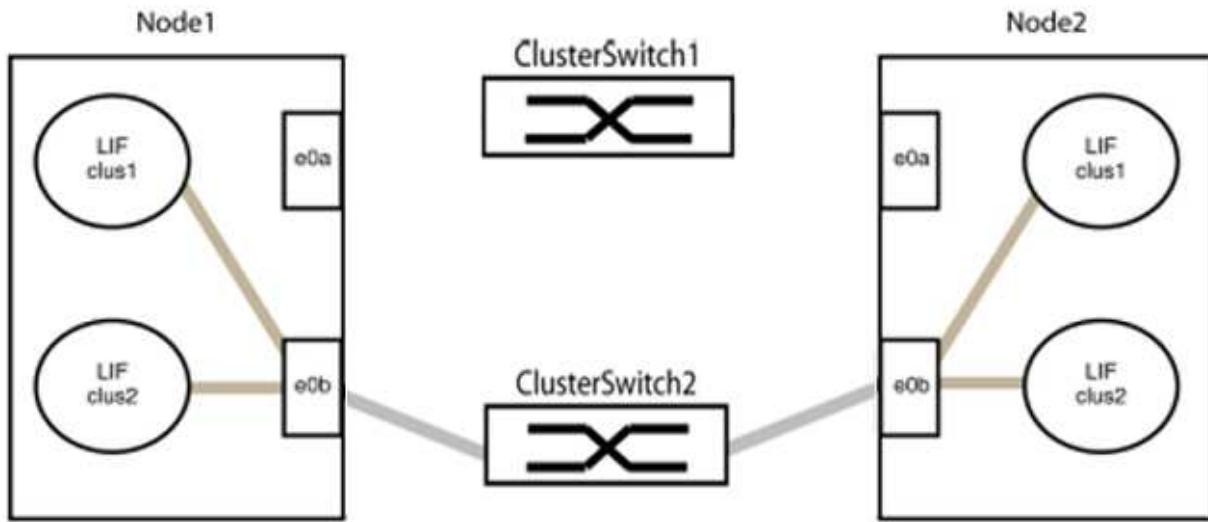
8. Set up the switchless configuration for the ports in group 1.



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**.

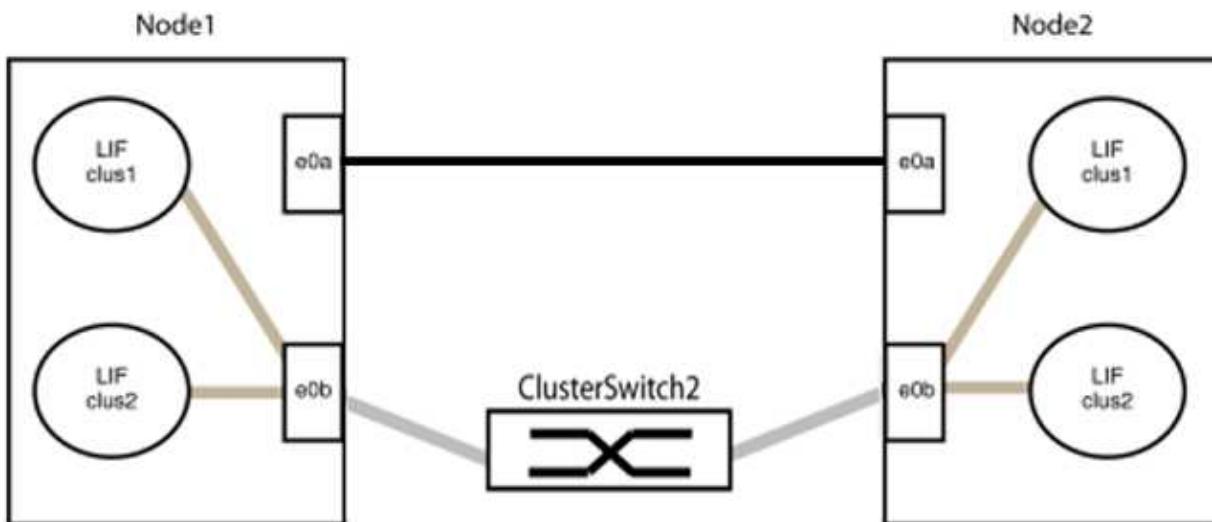
a. Disconnect all the cables from the ports in group1 at the same time.

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
                                         Source          Destination
Packet
Node    Date          LIF          LIF
Loss
-----
-----
node1
    3/5/2022 19:21:18 -06:00  node1_clus2      node2-clus1
none
    3/5/2022 19:21:20 -06:00  node1_clus2      node2_clus2
none
node2
    3/5/2022 19:21:18 -06:00  node2_clus2      node1_clus1
none
    3/5/2022 19:21:20 -06:00  node2_clus2      node1_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 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.

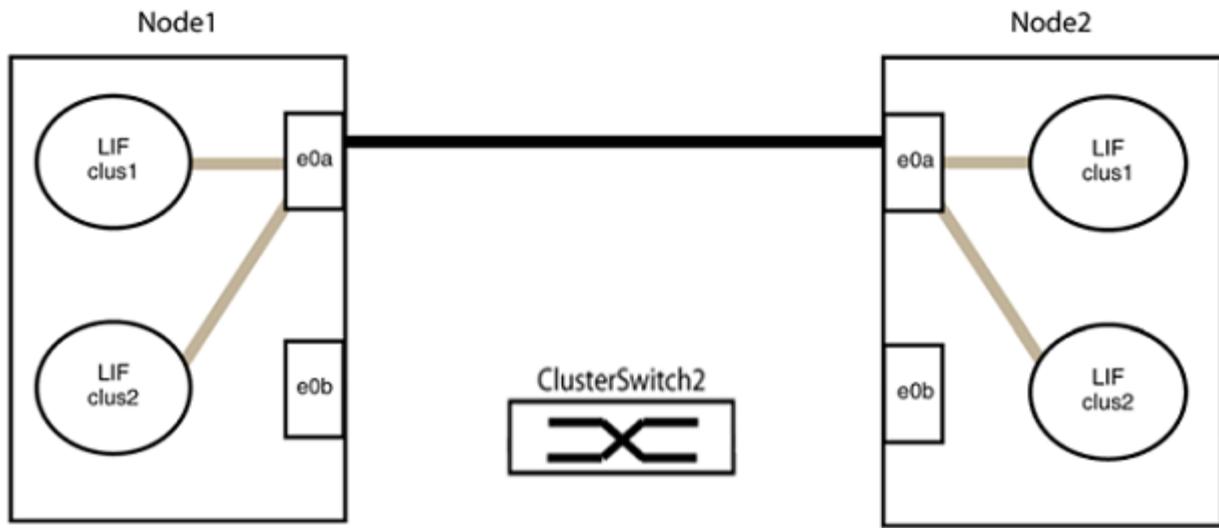
11. Set up the switchless configuration for the ports in group 2.



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
      e0b    node2
node1/lldp
      e0a    node2 (00:a0:98:da:16:44) e0a
      e0b    node2 (00:a0:98:da:16:44) e0b
node2/cdp
      e0a    node1
      e0b    node1
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  
Source Destination  
Packet  
Node Date LIF LIF  
Loss  
-----  
-----  
node1  
3/5/2022 19:21:18 -06:00 node1_clus2 node2-clus1  
none  
3/5/2022 19:21:20 -06:00 node1_clus2 node2_clus2  
none  
node2  
3/5/2022 19:21:18 -06:00 node2_clus2 node1_clus1  
none  
3/5/2022 19:21:20 -06:00 node2_clus2 node1_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 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—with 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.