



Upgrade the switch

Install and maintain

NetApp
February 20, 2026

Table of Contents

- Upgrade the switch 1
 - Upgrade workflow for BES-53248 cluster switches 1
 - Upgrade the EFOS software 1
 - Prepare for upgrade 1
 - Upgrade the software 3
- Upgrade the Reference Configuration File (RCF) 12
 - Step 1: Prepare for upgrade 13
 - Step 2: Configure ports 19
 - Step 3: Verify the configuration 36
- Verify the ONTAP cluster network after an EFOS software or RCF upgrade of the BES-53248 cluster switches 42

Upgrade the switch

Upgrade workflow for BES-53248 cluster switches

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

1

Upgrade your EFOS version

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

2

Upgrade your RCF version

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

3

Verify the ONTAP cluster network after upgrade

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

Upgrade the EFOS software

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

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

Prepare for upgrade

Before you begin

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

Note the following:

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

Procedure	Current EFOS version	New EFOS version	High level steps
Steps to upgrade EFOS between two (non) FIPS compliant versions	3.4.x.x	3.4.x.x	Upgrade the new EFOS image using Method 1: Upgrade EFOS . The configuration and license information is retained.
	3.4.4.6 (or later 3.4.x.x)	3.7.x.x or later non-FIPS compliant	Upgrade EFOS using Method 1: Upgrade EFOS . Reset the switch to factory defaults and apply the RCF file for EFOS 3.7.x.x or later.
	3.7.x.x or later non-FIPS compliant	3.4.4.6 (or later 3.4.x.x)	Downgrade EFOS using Method 1: Upgrade EFOS . Reset the switch to factory defaults and apply the RCF file for EFOS 3.4.x.x
		3.7.x.x or later non-FIPS compliant	Upgrade the new EFOS image using Method 1: Upgrade EFOS . The configuration and license information is retained.
	3.7.x.x or later FIPS compliant	3.7.x.x or later FIPS compliant	Upgrade the new EFOS image using Method 1: Upgrade EFOS . The configuration and license information is retained.

Steps to upgrade to/from a FIPS compliant EFOS version	Non-FIPS compliant	FIPS compliant	Upgrade of the EFOS image using Method 2: Upgrade EFOS using the ONIE OS installation . The switch configuration and license information will be lost.
	FIPS compliant	Non-FIPS compliant	

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

- On switch IP_switch_a1 (FIPS compliant EFOS):

```
IP_switch_a1 # show fips status
System running in FIPS mode
```

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

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

Upgrade the software

Use one of the following methods:

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



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

Method 1: Upgrade EFOS

Perform the following steps to upgrade the EFOS software.



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

Steps

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

This example verifies that the switch is connected to the server at IP address 172.19.2.1:

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

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

3. Disable auto-revert on the cluster LIFs.

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

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

```
show bootvar
```

Show example

```
(cs2)# show bootvar

Image Descriptions

active :
backup :

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

5. Download the image file to the switch.

Copying the image file to the backup image means that when you reboot, that image establishes the running EFOS version, completing the update.

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

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

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

File transfer operation completed successfully.
```

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

```
show bootvar
```

Show example

```
(cs2)# show bootvar

Image Descriptions

active :
backup :

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

7. Boot the system from the backup configuration:

```
boot system backup
```

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

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

```
show bootvar
```

Show example

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

```
Image Descriptions
```

```
active :
```

```
backup :
```

```
Images currently available on Flash
```

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

9. Save the running configuration to the startup configuration:

```
write memory
```

Show example

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

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

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

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

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

```
Configuration Saved!
```

10. Reboot the switch:

```
reload
```

Show example

```
(cs2)# reload
```

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

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

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

```
Configuration Saved!
```

```
System will now restart!
```

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

```
show version
```

Show example

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

```
Switch: 1
```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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:

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

e. [Configure SNMPv3 for monitoring](#)

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

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

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

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

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

Upgrade the Reference Configuration File (RCF)

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

Before you begin

Make sure you have the following:

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

Suggested documentation

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

About the examples

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

- The names of the two BES-53248 switches are cs1 and cs2.
- The node names are cluster1-01, cluster1-02, cluster1-03, and cluster1-04.
- The cluster LIF names are cluster1-01_clus1, cluster1-01_clus2, cluster1-02_clus1, cluster1-02_clus2, cluster1-03_clus1, cluster1-03_clus2, cluster1-04_clus1, and cluster1-04_clus2.

- The `cluster1::*>` prompt indicates the name of the cluster.
- The examples in this procedure use four nodes. These nodes use two 10GbE cluster interconnect ports `e0a` and `e0b`. See the [Hardware Universe](#) to verify the correct cluster ports on your platforms.



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

About this task

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

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



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

Step 1: Prepare for upgrade

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

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

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



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

The following command suppresses automatic case creation for two hours:

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

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

```
set -privilege advanced
```

The advanced prompt (`*>`) appears.

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

```
network device-discovery show
```

Show example

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

4. Check the administrative and operational status of each cluster port.
 - a. Verify that all the cluster ports are up with a healthy status:

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

Show example

```
cluster1::*> network port show -ipspace Cluster
Node: cluster1-01

Ignore

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

Node: cluster1-02

Ignore

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

Node: cluster1-03

Ignore

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

```
Node: cluster1-04
```

```
Ignore
```

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

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

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

Show example

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

Current Vserver Port	Logical Current Home	Is Interface	Status Admin/Oper	Network Address/Mask	Node
Cluster					
cluster1-01	cluster1-01_e0a	clus1 true	up/up	169.254.3.4/23	
cluster1-01	cluster1-01_e0b	clus2 true	up/up	169.254.3.5/23	
cluster1-02	cluster1-02_e0a	clus1 true	up/up	169.254.3.8/23	
cluster1-02	cluster1-02_e0b	clus2 true	up/up	169.254.3.9/23	
cluster1-03	cluster1-03_e0a	clus1 true	up/up	169.254.1.3/23	
cluster1-03	cluster1-03_e0b	clus2 true	up/up	169.254.1.1/23	
cluster1-04	cluster1-04_e0a	clus1 true	up/up	169.254.1.6/23	
cluster1-04	cluster1-04_e0b	clus2 true	up/up	169.254.1.7/23	

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

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command:

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

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

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

```
cluster1::*>
```

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, use the command:

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

```

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

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

```

6. Disable auto-revert on the cluster LIFs.

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

Step 2: Configure ports

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

```
show isdp neighbor
```

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

```

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

```

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

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

Show example

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

4. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
cluster1::*> cluster show
Node                Health  Eligibility  Epsilon
-----
cluster1-01        true    true         false
cluster1-02        true    true         false
cluster1-03        true    true         true
cluster1-04        true    true         false
```

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

```
show running-config
```

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



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



Clearing the configuration does not delete licenses.

- a. SSH into the switch.

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

- b. Enter privilege mode:

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

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

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

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

d. Save the running configuration to the startup configuration:

write memory

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

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

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

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

e. Perform a reboot of the switch:

reload

```
(cs2)# reload
```

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

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

7. Note the following:

- a. If additional port licenses have been installed on the switch, you must modify the RCF to configure the additional licensed ports. See [Activate newly licensed ports](#) for details. However, when you upgrade to RCF 1.12 or later, the modifications are no longer needed because all interfaces are now pre-configured.
- b. Record any customizations that were made in the previous RCF and apply these to the new RCF. For example, setting port speeds or hard-coding FEC mode.

EFOS version 3.12.x and later

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

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

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

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

```
script list
```

```
(cs2)# script list

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

- Apply the script to the switch:

```
script apply
```

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

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

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

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

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

```
Configuration Saved!
```

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

All other EFOS versions

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

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

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

```
Remote Password:**
```

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

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

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

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

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

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

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

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

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

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

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

```
script list
```

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

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

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

10. Apply the script to the switch:

```
script apply
```

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

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

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

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

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

```
Configuration Saved!
```

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

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

```
show clibanner
```

Show example

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

```
Banner Message configured :
```

```
=====
```

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

```
Switch    : BES-53248
```

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

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

```
Version   : v1.9
```

```
Port Usage:
```

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

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

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

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

```
NOTE:
```

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

```
speed:
```

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

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

```
group
```

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

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

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

```
startup-config'
```

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

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

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

Show example

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

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

```

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

```

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

```
show running-config
```

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

```
write memory
```

Show example

```

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

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

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

```

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

```
reload
```

```

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

```

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

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

17. Save the running configuration to the startup configuration:

```
write memory
```

Show example

```
(cs2)# write memory

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

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

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

18. Verify the ports on switch cs2:

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

Show example

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

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

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

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

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

Show example

```
cluster1::*> network port show -ipspace Cluster
Node: cluster1-01

Ignore

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

Node: cluster1-02

Ignore

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

Node: cluster1-03

Ignore

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

```
Node: cluster1-04
```

```
Ignore
```

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

b. Verify the switch health from the cluster:

```
network device-discovery show
```

Show example

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

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

ONTAP 9.8 and later

Beginning with ONTAP 9.8, use the command:

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

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

Switch	Type	Address	Model
cs1	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
```

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

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

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

Show example

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

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

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

Show example

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

4. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
cluster1::*> cluster show
Node          Health  Eligibility  Epsilon
-----
cluster1-01   true    true         false
cluster1-02   true    true         false
cluster1-03   true    true         true
cluster1-04   true    true         false
```

5. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

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

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

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

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

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

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

All ONTAP releases

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

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

```

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

```

6. Change the privilege level back to admin:

```
set -privilege admin
```

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

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

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

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

Steps

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

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

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

Show example

The following example shows the output from the command:

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

Node: node1

Ignore

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

Node: node2

Ignore

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

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

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

Show example

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

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

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

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

```
cluster show
```

Show example

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

Node	Health	Eligibility	Epsilon

node1	true	true	false
node2	true	true	false

What's next?

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

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.