



Replace switches

Cluster and storage switches

NetApp
April 05, 2024

Table of Contents

- Replace switches 1
 - Replacement requirements 1
 - Replace a Broadcom-supported BES-53248 cluster switch 1
 - Replace Broadcom BES-53248 cluster switches with switchless connections 13

Replace switches

Replacement requirements

Before replacing the switch, make sure the following conditions are met in the current environment and on the replacement switch.

Existing cluster and network infrastructure

Make sure that:

- The existing cluster is verified as completely functional, with at least one fully connected cluster switch.
- All cluster ports are **up**.
- All cluster logical interfaces (LIFs) are administratively and operationally **up** and on their home ports.
- The ONTAP `cluster ping-cluster -node node1` command must indicate that the settings, basic connectivity and larger than PMTU communication, are successful on all paths.

BES-53248 replacement cluster switch

Make sure that:

- Management network connectivity on the replacement switch is functional.
- Console access to the replacement switch is in place.
- The node connections are ports 0/1 through 0/16 with default licensing.
- All Inter-Switch Link (ISL) ports are disabled on ports 0/55 and 0/56.
- The desired reference configuration file (RCF) and EFOS operating system switch image are loaded onto the switch.
- Initial customization of the switch is complete, as detailed in [Configure the BES-53248 cluster switch](#).

Any previous site customizations, such as STP, SNMP, and SSH, are copied to the new switch.

For more information

- [NetApp Support Site](#)
- [NetApp Hardware Universe](#)

Replace a Broadcom-supported BES-53248 cluster switch

Follow these steps to replace a defective Broadcom-supported BES-53248 cluster switch in a cluster network. This is a nondisruptive procedure (NDU).

About the examples

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

- The names of the existing BES-53248 switches are `cs1` and `cs2`.

- The name of the new BES-53248 switch is `newcs2`.
- The node names are `node1` and `node2`.
- The cluster ports on each node are named `e0a` and `e0b`.
- The cluster LIF names are `node1_clus1` and `node1_clus2` for `node1`, and `node2_clus1` and `node2_clus2` for `node2`.
- The prompt for changes to all cluster nodes is `cluster1::>`

About the topology

This procedure is based on the following cluster network topology:

Show example topology

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

Node: node1

Ignore

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

e0a	Cluster	Cluster		up	9000	auto/10000	healthy
false							
e0b	Cluster	Cluster		up	9000	auto/10000	healthy
false							

Node: node2

Ignore

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

e0a	Cluster	Cluster		up	9000	auto/10000	healthy
false							
e0b	Cluster	Cluster		up	9000	auto/10000	healthy
false							

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

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

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

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

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

Node/ Protocol	Local Port	Discovered Device (LLDP: ChassisID)	Interface	Platform
node2	/cdp			
	e0a	cs1	0/2	BES-
53248				
	e0b	cs2	0/2	BES-
53248				
node1	/cdp			
	e0a	cs1	0/1	BES-
53248				
	e0b	cs2	0/1	BES-
53248				

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

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge,

S - Switch, H - Host, I - IGMP, r - Repeater

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

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

Capability Codes: R - Router, T - Trans Bridge, B - Source Route Bridge,

S - Switch, H - Host, I - IGMP, r - Repeater

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

Steps

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

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

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



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

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

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

- a. You can download the applicable Broadcom EFOS software for your cluster switches from the [Broadcom Ethernet Switch Support](#) site. Follow the steps on the Download page to download the EFOS file for the version of ONTAP software you are installing.
 - b. The appropriate RCF is available from the [Broadcom Cluster Switches](#) page. Follow the steps on the Download page to download the correct RCF for the version of ONTAP software you are installing.
4. On the new switch, log in as `admin` and shut down all of the ports that will be connected to the node cluster interfaces (ports 1 to 16).



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

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



No password is required to enter `enable` mode.

Show example

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

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

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


Show example topology

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

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

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

Show example topology

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

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

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

Show example

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

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

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

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

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



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

Show example

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

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

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

Show example

The output should be similar to the following:

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

Node: node1

Ignore

						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
```

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.
```

+

```
cs1# show cdp neighbors
```

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge

S - Switch, H - Host, I - IGMP, r - Repeater,
V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute

Device-ID Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0a	Eth1/1	144	H	FAS2980
node2 e0a	Eth1/2	145	H	FAS2980
newcs2 (FDO296348FU) Eth1/65	Eth1/65	176	R S I s	N9K-C92300YC
newcs2 (FDO296348FU) Eth1/66	Eth1/66	176	R S I s	N9K-C92300YC

```
cs2# show cdp neighbors
```

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge

S - Switch, H - Host, I - IGMP, r - Repeater,
V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute

Device-ID Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0b	Eth1/1	139	H	FAS2980
node2 e0b	Eth1/2	124	H	FAS2980
cs1 (FDO220329KU) Eth1/65	Eth1/65	178	R S I s	N9K-C92300YC
cs1 (FDO220329KU) Eth1/66	Eth1/66	178	R S I s	N9K-C92300YC

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?

See [Enable the log collection feature](#) for the steps required to enable cluster health switch log collection used for collecting switch-related log files.

Replace Broadcom BES-53248 cluster switches with switchless connections

You can migrate from a cluster with a switched cluster network to one where two nodes are directly connected for ONTAP 9.3 and later.

Review requirements

Guidelines

Review the following guidelines:

- Migrating to a two-node switchless cluster configuration is a nondisruptive operation. Most systems have two dedicated cluster interconnect ports on each node, but you can also use this procedure for systems with a larger number of dedicated cluster interconnect ports on each node, such as four, six or eight.
- You cannot use the switchless cluster interconnect feature with more than two nodes.
- If you have an existing two-node cluster that uses cluster interconnect switches and is running ONTAP 9.3 or later, you can replace the switches with direct, back-to-back connections between the nodes.


What you'll need

- A healthy cluster that consists of two nodes connected by cluster switches. The nodes must be running the same ONTAP release.
- Each node with the required number of dedicated cluster ports, which provide redundant cluster interconnect connections to support your system configuration. For example, there are two redundant ports for a system with two dedicated cluster interconnect ports on each node.

Migrate the switches

About this task

The following procedure removes the cluster switches in a two-node cluster and replaces each connection to the switch with a direct connection to the partner node.



About the examples

The examples in the following procedure show nodes that are using "e0a" and "e0b" as cluster ports. Your nodes might be using different cluster ports as they vary by system.

Step 1: Prepare for migration

1. Change the privilege level to advanced, entering `y` when prompted to continue:

```
set -privilege advanced
```


The advanced prompt `*>` appears.

2. ONTAP 9.3 and later supports automatic detection of switchless clusters, which is enabled by default.

You can verify that detection of switchless clusters is enabled by running the advanced privilege command:

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

Show example

The following example output shows if the option is enabled.

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

If "Enable Switchless Cluster Detection" is `false`, contact NetApp support.

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

```
system node autosupport invoke -node * -type all -message
MAINT=<number_of_hours>h
```

where `h` is the duration of the maintenance window in hours. The message notifies technical support of this maintenance task so that they can suppress automatic case creation during the maintenance window.

In the following example, the command suppresses automatic case creation for two hours:

Show example


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

Step 2: Configure ports and cabling

1. Organize the cluster ports on each switch into groups so that the cluster ports in group1 go to cluster switch1 and the cluster ports in group2 go to cluster switch2. These groups are required later in the procedure.
2. Identify the cluster ports and verify link status and health:

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

In the following example for nodes with cluster ports "e0a" and "e0b", one group is identified as "node1:e0a" and "node2:e0a" and the other group as "node1:e0b" and "node2:e0b". Your nodes might be using different cluster ports because they vary by system.



Verify that the ports have a value of up for the “Link” column and a value of healthy for the “Health Status” column.

Show example

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

```
Node: node1
```

```
Ignore
```

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

```
Node: node2
```

```
Ignore
```

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

```
4 entries were displayed.
```

3. Confirm that all the cluster LIFs are on their home ports.

Verify that the “is-home” column is true for each of the cluster LIFs:

```
network interface show -vserver Cluster -fields is-home
```

Show example

```
cluster::*> net int show -vserver Cluster -fields is-home
(network interface show)
vserver  lif          is-home
-----  -
Cluster  node1_clus1  true
Cluster  node1_clus2  true
Cluster  node2_clus1  true
Cluster  node2_clus2  true
4 entries were displayed.
```

If there are cluster LIFs that are not on their home ports, revert those LIFs to their home ports:

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

4. Disable auto-revert for the cluster LIFs:

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

5. Verify that all ports listed in the previous step are connected to a network switch:

```
network device-discovery show -port cluster_port
```

The “Discovered Device” column should be the name of the cluster switch that the port is connected to.

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to cluster switches "cs1" and "cs2".

```
cluster::> network device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1/cdp
          e0a    cs1                      0/11       BES-53248
          e0b    cs2                      0/12       BES-53248
node2/cdp
          e0a    cs1                      0/9        BES-53248
          e0b    cs2                      0/9        BES-53248
4 entries were displayed.
```

6. Verify the cluster connectivity:

```
cluster ping-cluster -node local
```

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 that the cluster network is not disrupted:

```
cluster ping-cluster -node local
```



Before proceeding to the next step, you must wait at least two minutes to confirm a working back-to-back connection on group 1.


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



To avoid potential networking issues, you must disconnect the ports from group2 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.

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

In the following example, the cables are disconnected from port "e0b" on each node, and cluster traffic continues through the direct connection between the "e0a" ports:



- b. Cable the ports in group2 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2 and "e0b" on node1 is connected to "e0b" on node2:



Step 3: Verify the configuration

1. Verify that the ports on both nodes are correctly connected:

```
network device-discovery show -port cluster_port
```

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to the corresponding port on the cluster partner:

```
cluster::> net device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1/cdp
          e0a    node2                      e0a        AFF-A300
          e0b    node2                      e0b        AFF-A300
node1/lldp
          e0a    node2 (00:a0:98:da:16:44) e0a        -
          e0b    node2 (00:a0:98:da:16:44) e0b        -
node2/cdp
          e0a    node1                      e0a        AFF-A300
          e0b    node1                      e0b        AFF-A300
node2/lldp
          e0a    node1 (00:a0:98:da:87:49) e0a        -
          e0b    node1 (00:a0:98:da:87:49) e0b        -
8 entries were displayed.
```

2. Re-enable auto-revert for the cluster LIFs:

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

3. Verify that all LIFs are home. This might take a few seconds.

```
network interface show -vserver Cluster -lif lif_name
```

Show example

The LIFs have been reverted if the “Is Home” column is `true`, as shown for `node1_clus2` and `node2_clus2` in the following example:

```
cluster::> network interface show -vserver Cluster -fields curr-  
port,is-home  
vserver  lif                curr-port is-home  
-----  
Cluster  node1_clus1         e0a      true  
Cluster  node1_clus2         e0b      true  
Cluster  node2_clus1         e0a      true  
Cluster  node2_clus2         e0b      true  
4 entries were displayed.
```

If any cluster LIFS have not returned to their home ports, revert them manually from the local node:

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

4. Check the cluster status of the nodes from the system console of either node:

```
cluster show
```

Show example

The following example shows `epsilon` on both nodes to be `false`:

```
Node  Health  Eligibility Epsilon  
-----  
node1 true    true       false  
node2 true    true       false  
2 entries were displayed.
```

5. Confirm connectivity between the cluster ports:

```
cluster ping-cluster local
```

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
```

Copyright information

Copyright © 2024 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.