



Replace switches

Cluster and storage switches

NetApp
April 05, 2024

Table of Contents

- Replace switches 1
 - Replace a NVIDIA SN2100 cluster switch 1
 - Replace NVIDIA SN2100 cluster switches with switchless connections 15

Replace switches

Replace a NVIDIA SN2100 cluster switch

Follow this procedure to replace a defective NVIDIA SN2100 switch in a cluster network. This is a nondisruptive procedure (NDU).

Review requirements

Existing cluster and network infrastructure

Ensure that:

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

NVIDIA SN2100 replacement switch

Ensure that:

- Management network connectivity on the replacement switch are functional.
- Console access to the replacement switch are in place.
- The node connections are ports swp1 through swp14.
- All Inter-Switch Link (ISL) ports are disabled on ports swp15 and swp16.
- The desired reference configuration file (RCF) and Cumulus operating system image switch are loaded onto the switch.
- Initial customization of the switch is complete.

Also make sure that any previous site customizations, such as STP, SNMP, and SSH, are copied to the new switch.



You must execute the command for migrating a cluster LIF from the node where the cluster LIF is hosted.

Replace the switch

About the examples

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

- The names of the existing NVIDIA SN2100 switches are *sw1* and *sw2*.
- The name of the new NVIDIA SN2100 switch is *nsw2*.
- The node names are *node1* and *node2*.
- The cluster ports on each node are named *e3a* and *e3b*.
- The cluster LIF names are *node1_clus1* and *node1_clus2* for node1, and *node2_clus1* and *node2_clus2*

for node2.

- The prompt for changes to all cluster nodes is `cluster1:*>`
- Breakout ports take the format: `swp[port]s[breakout port 0-3]`. For example, four breakout ports on `swp1` are *swp1s0*, *swp1s1*, *swp1s2*, and *swp1s3*.

About the cluster network topology

This procedure is based on the following cluster network topology:

Show example topology

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

Node: node1

Ignore

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

e3a	Cluster	Cluster		up	9000	auto/100000	healthy
false							
e3b	Cluster	Cluster		up	9000	auto/100000	healthy
false							

Node: node2

Ignore

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

e3a	Cluster	Cluster		up	9000	auto/100000	healthy
false							
e3b	Cluster	Cluster		up	9000	auto/100000	healthy
false							

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

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

Cluster					
	node1_clus1	up/up	169.254.209.69/16	node1	e3a
true					
	node1_clus2	up/up	169.254.49.125/16	node1	e3b
true					

```

node2_clus1 up/up 169.254.47.194/16 node2 e3a
true
node2_clus2 up/up 169.254.19.183/16 node2 e3b
true

```

```
cluster1::*> network device-discovery show -protocol lldp
```

Node/	Local	Discovered			
Protocol	Port	Device (LLDP: ChassisID)	Interface	Platform	
node1	/lldp				
	e3a	sw1 (b8:ce:f6:19:1a:7e)	swp3	-	
	e3b	sw2 (b8:ce:f6:19:1b:96)	swp3	-	
node2	/lldp				
	e3a	sw1 (b8:ce:f6:19:1a:7e)	swp4	-	
	e3b	sw2 (b8:ce:f6:19:1b:96)	swp4	-	

+

```
cumulus@sw1:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
swp3	100G	Trunk/L2	sw2	e3a
swp4	100G	Trunk/L2	sw2	e3a
swp15	100G	BondMember	sw2	swp15
swp16	100G	BondMember	sw2	swp16

```
cumulus@sw2:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
swp3	100G	Trunk/L2	sw1	e3b
swp4	100G	Trunk/L2	sw1	e3b
swp15	100G	BondMember	sw1	swp15
swp16	100G	BondMember	sw1	swp16

Step 1: Prepare for replacement

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

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

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

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

```
set -privilege advanced
```

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

3. Install the appropriate RCF and image on the switch, nsw2, and make any necessary site preparations.

If necessary, verify, download, and install the appropriate versions of the RCF and Cumulus software for the new switch.

- a. You can download the applicable Cumulus software for your cluster switches from the *NVIDIA Support* site. Follow the steps on the Download page to download the Cumulus Linux for the version of ONTAP software you are installing.
- b. The appropriate RCF is available from the [NVIDIA Cluster and Storage Switches](#) page. Follow the steps on the Download page to download the correct RCF for the version of ONTAP software you are installing.

Step 2: Configure ports and cabling

1. On the new switch nsw2, log in as admin and shut down all of the ports that will be connected to the node cluster interfaces (ports swp1 to swp14).

The LIFs on the cluster nodes should have already failed over to the other cluster port for each node.

Show example

```
cumulus@nsw2:~$ net add interface swp1s0-3, swp2s0-3, swp3-14 link  
down  
cumulus@nsw2:~$ net pending  
cumulus@nsw2:~$ net commit
```

2. Disable auto-revert on the cluster LIFs:

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

Show example

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

```
Warning: Disabling the auto-revert feature of the cluster logical  
interface may effect the availability of your cluster network. Are  
you sure you want to continue? {y|n}: y
```

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

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

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

Show example

```
cumulus@sw1:~$ net add interface swp15-16 link down
cumulus@sw1:~$ net pending
cumulus@sw1:~$ net commit
```

5. Remove all the cables from the SN2100 sw1 switch, and then connect them to the same ports on the SN2100 nsw2 switch.
6. Bring up the ISL ports swp15 and swp16 between the sw1 and nsw2 switches.

Show example

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

```
cumulus@sw1:~$ net del interface swp15-16 link down
cumulus@sw1:~$ net pending
cumulus@sw1:~$ net commit
```

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

```
cumulus@sw1:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp15	100G	9216	BondMember	nsw2 (swp15)	Master: cluster_isl (UP)
UP	swp16	100G	9216	BondMember	nsw2 (swp16)	Master: cluster_isl (UP)

+ The following example shows that the ISL ports are up on switch nsw2:

+

```
cumulus@nsw2:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP	Summary
UP	swp15	100G	9216	BondMember	sw1 (swp15)	Master: cluster_isl (UP)
UP	swp16	100G	9216	BondMember	sw1 (swp16)	Master: cluster_isl (UP)

7. Verify that port e3b is up on all nodes:

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

Show example

The output should be similar to the following:

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

Node: node1

Ignore

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/100000
healthy	false					

Node: node2

Ignore

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/100000
healthy	false					

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

Show example

```
cluster1::~*> network device-discovery show -protocol lldp
```

Node/	Local	Discovered			
Protocol	Port	Device	(LLDP: ChassisID)	Interface	Platform
-----	-----	-----	-----	-----	
node1	/lldp				
	e3a	sw1	(b8:ce:f6:19:1a:7e)	swp3	-
	e3b	nsw2	(b8:ce:f6:19:1b:b6)	swp3	-
node2	/lldp				
	e3a	sw1	(b8:ce:f6:19:1a:7e)	swp4	-
	e3b	nsw2	(b8:ce:f6:19:1b:b6)	swp4	-

9. Verify that all node cluster ports are up:

```
net show interface
```

Show example

```
cumulus@nsw2:~$ net show interface
```

State	Name	Spd	MTU	Mode	LLDP
Summary	-----	-----	-----	-----	-----

...					
...					
UP	swp3	100G	9216	Trunk/L2	
Master: bridge(UP)					
UP	swp4	100G	9216	Trunk/L2	
Master: bridge(UP)					
UP	swp15	100G	9216	BondMember	sw1 (swp15)
Master: cluster_isl(UP)					
UP	swp16	100G	9216	BondMember	sw1 (swp16)
Master: cluster_isl(UP)					

10. Verify that both nodes each have one connection to each switch:

```
net show lldp
```

Show example

The following example shows the appropriate results for both switches:

```
cumulus@sw1:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
-----	-----	-----	-----	-----
swp3	100G	Trunk/L2	node1	e3a
swp4	100G	Trunk/L2	node2	e3a
swp15	100G	BondMember	nsw2	swp15
swp16	100G	BondMember	nsw2	swp16

```
cumulus@nsw2:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
-----	-----	-----	-----	-----
swp3	100G	Trunk/L2	node1	e3b
swp4	100G	Trunk/L2	node2	e3b
swp15	100G	BondMember	sw1	swp15
swp16	100G	BondMember	sw1	swp16

11. Enable auto-revert on the cluster LIFs:

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

12. On switch nsw2, bring up the ports connected to the network ports of the nodes.

Show example

```
cumulus@nsw2:~$ net del interface swp1-14 link down  
cumulus@nsw2:~$ net pending  
cumulus@nsw2:~$ net commit
```

13. Display information about the nodes in a cluster:

```
cluster show
```

Show example

This example shows that the node health for node1 and node2 in this cluster is true:

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

Node	Health	Eligibility
-----	-----	-----
node1	true	true
node2	true	true

14. Verify that all physical cluster ports are up:

```
network port show ipspace Cluster
```

Show example

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

Node node1

Ignore

					Speed (Mbps)	
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

Node: node2

Ignore

					Speed (Mbps)	
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e3b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

Step 3: Complete the procedure

1. Verify that the cluster network is healthy.

Show example

```
cumulus@sw1:~$ net show lldp
```

LocalPort	Speed	Mode	RemoteHost	RemotePort
-----	-----	-----	-----	-----
swp3	100G	Trunk/L2	node1	e3a
swp4	100G	Trunk/L2	node2	e3a
swp15	100G	BondMember	nsw2	swp15
swp16	100G	BondMember	nsw2	swp16

2. Create a password for the Ethernet switch health monitor log collection feature:

```
system switch ethernet log setup-password
```

Show example

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: <return>
```

```
The switch name entered is not recognized.
```

```
Choose from the following list:
```

```
cs1
```

```
cs2
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: cs1
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
Enter the password: <enter switch password>
```

```
Enter the password again: <enter switch password>
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: cs2
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
Enter the password: <enter switch password>
```

```
Enter the password again: <enter switch password>
```

3. Enable the Ethernet switch health monitor log collection feature.

```
system switch ethernet log modify -device <switch-name> -log-request true
```

Show example

```
cluster1::*> system switch ethernet log modify -device cs1 -log  
-request true
```

```
Do you want to modify the cluster switch log collection  
configuration? {y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

```
cluster1::*> system switch ethernet log modify -device cs2 -log  
-request true
```

```
Do you want to modify the cluster switch log collection  
configuration? {y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

Wait for 10 minutes and then check that the log collection completes:

```
system switch ethernet log show
```

Show example

```
cluster1::*> system switch ethernet log show  
Log Collection Enabled: true
```

Index	Switch	Log Timestamp	Status
-----	-----	-----	-----
1	cs1 (b8:ce:f6:19:1b:42)	4/29/2022 03:05:25	complete
2	cs2 (b8:ce:f6:19:1b:96)	4/29/2022 03:07:42	complete



If any of these commands return an error or if the log collection does not complete, contact NetApp support.

4. Change the privilege level back to admin:

```
set -privilege admin
```

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


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

Replace NVIDIA SN2100 cluster switches with switchless connections

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

Review requirements

Guidelines

Review the following guidelines:

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


What you'll need

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

Migrate the switches

About this task

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



About the examples

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

Step 1: Prepare for migration

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

```
set -privilege advanced
```

The advanced prompt `*>` appears.

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

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

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

Show example

The following example output shows if the option is enabled.

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

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

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

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

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

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

Show example


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

Step 2: Configure ports and cabling

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

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

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



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

Show example

```
cluster::> network port show -ipspace Cluster
Node: node1

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

Node: node2

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

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

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

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

Show example

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

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

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

4. Disable auto-revert for the cluster LIFs:

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

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

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

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

Show example

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

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

6. Verify the 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:



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.