

---

## Replacing a PCIe card

To replace a PCIe card in the system, you must perform a specific sequence of tasks.

### About this task

- You can use this procedure with all versions of ONTAP supported by your system.
- All other components in the system must be functioning properly; if not, you must contact technical support.
- PCIe cards in IOXM expansion modules need to follow separate instructions  
[Replacing an I/O expansion module or PCIe card in an I/O expansion module](#)

### Steps

1. [Shutting down the target controller](#) on page 1
2. [Opening the system](#) on page 6
3. [Removing a PCIe card](#) on page 7
4. [Installing a PCIe card](#) on page 9
5. [Configuring and cabling CNA ports \(80xx systems only\)](#) on page 9
6. [Reinstalling the controller module and booting the system](#) on page 10
7. [Completing the replacement process](#) on page 11

## Shutting down the target controller

You shut down or take over the target controller using different procedures, depending on the storage system hardware configuration.

### Shutting down a controller module in an HA pair

To shut down a controller module, you must determine the status of the impaired node and, if necessary, take over the impaired node so that the healthy partner continues to serve data from the impaired node's storage.

#### About this task

As part of this procedure, you either leave the power supplies on or turn them off, depending on your configuration:

- If you have two controller modules in the same chassis, you must leave the power supplies turned on to provide power to the healthy node.
- If you have one controller module in the chassis, you should turn off the power supplies in the chassis.

#### Steps

1. If the system is running ONTAP, check the status of the nodes in the cluster:
  - a. Change to the advance privilege level:  
`set -privilege advanced`
  - b. Check the health and eligibility of the node:

```
cluster show -epsilon *
```

The command produces output similar to :

```
Node          Health Eligibility  Epsilon
-----
node1         true   true         true
node2         true   true         false
node3         true   true         false
node4         true   true         false

4 entries were displayed.
```

**Note:** Epsilon must not be on a node to be replaced.

**Note:** In a cluster with a single HA pair, Epsilon will not be assigned to either node.

c. Take one of the following actions, depending on the result of the command:

If...	Then...
All nodes show <code>true</code> for both health and eligibility and Epsilon is not assigned to the impaired node.	<ol style="list-style-type: none"><li>Exit advanced mode: <code>set -privilege admin</code></li><li>Proceed to Step 3.</li></ol>
All nodes show <code>true</code> for both health and eligibility and Epsilon is assigned to the impaired node.	<p>Complete the following steps to move Epsilon:</p> <ol style="list-style-type: none"><li>Remove Epsilon from the node: <code>cluster modify -node node1 -epsilon false</code></li><li>Assign Epsilon to a node in the cluster: <code>cluster modify -node node4 -epsilon true</code></li><li>Exit advanced mode: <code>set -privilege admin</code></li><li>Go to Step 3.</li></ol>
The impaired node shows <code>false</code> for health and is the Epsilon node.	<p>Complete the following steps:</p> <ol style="list-style-type: none"><li>Change to the advance privilege level: <code>set -privilege advanced</code></li><li>Remove Epsilon from the node: <code>cluster modify -node node1 -epsilon false</code></li><li>Assign Epsilon to a node in the cluster: <code>cluster modify -node node4 -epsilon true</code></li><li>Exit advanced mode: <code>set -privilege admin</code></li><li>Proceed to the next step.</li></ol>

If...	Then...
The impaired node shows <code>false</code> for health and is not the Epsilon node.	<ol style="list-style-type: none"> <li>a. Exit advanced mode: <code>set -privilege admin</code></li> <li>b. Proceed to the next step.</li> </ol>
Any nodes show <code>false</code> for eligibility.	<ol style="list-style-type: none"> <li>a. Resolve any cluster issues as needed before continuing with this procedure.</li> <li>b. Exit advanced mode: <code>set -privilege admin</code></li> </ol>
Any nodes other than the impaired node show <code>false</code> for health.	<ol style="list-style-type: none"> <li>a. Correct the problems that cause the health issues on the nodes before continuing with this procedure.</li> <li>b. Exit advanced mode: <code>set -privilege admin</code></li> </ol>

2. Check the status of the impaired node (the node you want to perform maintenance on) by entering the following command at the system console of either node:

For..	Issue the command..
7-Mode	<code>cf status</code>
ONTAP	<code>storage failover show</code>

3. Take one of the following actions, depending on the result of the `cf status` or `storage failover show` command:

If..	Then..
Neither node is in takeover mode	Go to the next step in this procedure.
The healthy node took over the impaired node	The impaired node is in a state where you can begin removing it from the system chassis.
The impaired node took over the healthy node	<ol style="list-style-type: none"> <li>a. Correct the problem that caused the takeover.</li> <li>b. Enter the <code>cf giveback</code> command (7-Mode) or <code>storage failover giveback <i>impaired_node_name</i></code> command (ONTAP) from the impaired node console.</li> <li>c. Go back to the Step 1.</li> </ol>

4. Take over and power down the impaired node by taking the following steps:

- a. Enter one of the following commands from the healthy node's console and wait for the takeover to complete:

For systems operating in...	Issue the command..
7-Mode	<code>cf takeover</code>
ONTAP	<ul style="list-style-type: none"> <li>• In ONTAP 8.2 or later: <code>storage failover takeover -ofnode <i>impaired_node_name</i></code></li> </ul>

The impaired node is taken over and then automatically reboots and displays the `waiting for giveback...` message.

- b. Wait for two minutes after takeover of the impaired node to verify that the takeover was completed successfully.
- c. With the impaired node showing the `Waiting for giveback...` message, shut it down.

The method you use to shut down the node depends on whether remote management via a Service Processor (SP) is used and whether or not the system is in a dual-chassis or single-chassis configuration.

Is the SP configured?	Then...
Yes	Log in to the impaired node's SP and issue the following command: <b>system power off</b>
No, and the system is in a dual-chassis HA pair in which each controller module is in a separate chassis.	Proceed to Step 5.
No, and the system is in a single-chassis HA pair in which both controller modules are in the same chassis and share power supplies.	At the impaired node's prompt, press <b>Ctrl-C</b> and respond <b>Y</b> to halt the node.

The impaired node is now in a state where you can proceed to the next task.

- 5. If you are not already grounded, properly ground yourself.
- 6. If the system is in a dual-chassis HA pair, turn off the power supplies for the impaired node.

## Shutting down a controller module in a two-node MetroCluster configuration

To shut down a controller module, you must determine the status of the impaired node and, if necessary, perform a MetroCluster switchover operation from the healthy controller module so that the healthy node continues to serve data from the impaired node's storage.

### Steps

- 1. If the system is running ONTAP, check the status of the nodes in the cluster:
  - a. Enter the following command at the system console of either node:

**cluster show**

The command produces output similar to the following:

```

Node           Health Eligibility
-----
node1          true   true
node2          true   true
node3          true   true
node4          true   true

4 entries were displayed.
```

- b. Take one of the following actions, depending on the result of the command:

If...	Then...
All nodes show <code>true</code> for both health and eligibility.	Proceed to Step 3.
The impaired node shows <code>false</code> for health.	Proceed to the next step.
Any nodes show <code>false</code> for eligibility.	Resolve any cluster issues as needed before continuing with this procedure.

If...	Then...
Any nodes other than the impaired node show <code>false</code> for health.	Correct the problems that cause the health issues on the nodes before continuing with this procedure.

2. Use the `metrocluster check run`, `metrocluster check show` and `metrocluster check config-replication show` commands to make sure no configuration updates are in progress or pending.

3. If the impaired node has not switched over, perform the switchover operation from the healthy node:

```
metrocluster switchover
```

4. Monitor the completion of the switchover:

```
metrocluster operation show
```

### Example

```
mcclA::*> metrocluster operation show
Operation: Switchover
Start time: 10/4/2012 19:04:13
State: in-progress
End time: -
Errors:

mcclA::*> metrocluster operation show
Operation: Switchover
Start time: 10/4/2012 19:04:13
State: successful
End time: 10/4/2012 19:04:22
Errors: -
```

5. Shut down the impaired node.

The method you use to shut down the node depends on whether remote management using a Service Processor (SP) is used:

Is the SP configured?	Then...
Yes	Log in to the impaired node's SP and issue the following command: <code>system power off</code>
No	At the impaired node's prompt, press <b>Ctrl-C</b> and respond <b>Y</b> to halt the node.

6. If you are not already grounded, properly ground yourself.

## Shutting down a stand-alone controller module

For a stand-alone controller module, you must perform a clean shutdown to ensure that all data has been written to disk. You must also disconnect the power supplies.

### Steps

1. Enter the following command from the system console of the impaired node:

If your system is configured in...	Then issue this command...
7-Mode	<code>halt</code>
Clustered Data ONTAP	<code>halt local</code>

After you issue the command and respond **Y** to halt the node, wait until the system stops at the **LOADER** prompt.

**Attention:** You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the NVRAM.

The NVRAM has a green LED on the rear of the controller module that blinks while destaging data to the flash memory when you halt the system. After the destage is complete, the LED turns off.

- If power is lost without a clean shutdown, the NVRAM LED will flash until the destage is complete, and then the LED will turn off.
- If the LED is on and power is on, unwritten data is stored on NVRAM. This typically occurs during an uncontrolled shutdown after Data ONTAP had successfully booted.

2. If you are not already grounded, properly ground yourself.
3. Turn off the power supplies, and then unplug both power cords from the power source.

## Opening the system

If you want to access components inside the controller module, you must open the system.

### Steps

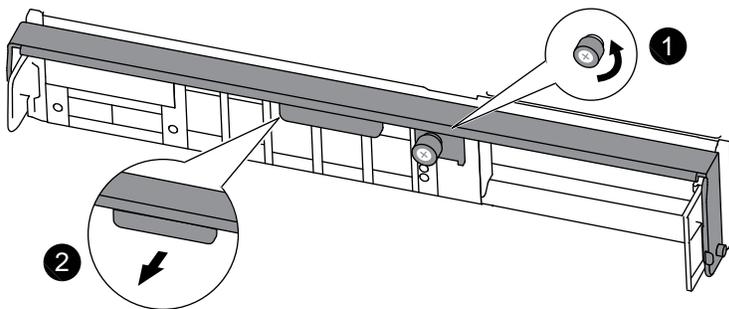
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, and keep track of where the cables were connected.

Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove the cable management device from the controller module and set it aside.

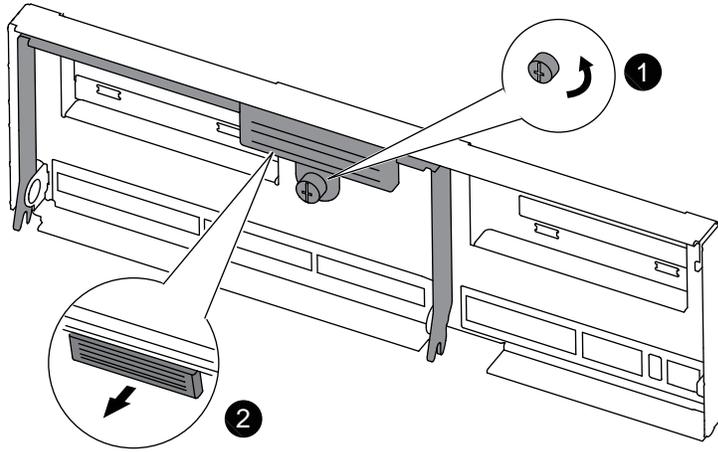
**Note:** The 8020 controller module uses cable management arms, while the rest of the 80xx family uses a cable management tray.

4. Loosen the thumbscrew on the cam handle of the controller module.



### 8020 controller module cam handle

1	Thumbscrew
2	Cam handle



**Any other 80xx controller module cam handle**

<b>1</b>	Thumbscrew
<b>2</b>	Cam handle

5. Pull the cam handle downward and begin to slide the controller module out of the chassis.

<b>If you have...</b>	<b>Then...</b>
An 8020 system	Slide the controller module completely out of the system.  <b>Note:</b> Ensure that you support the bottom of the controller module with your free hand and set it aside.
Any other 80xx model	Slide the controller module out of the system until it catches, press the release latch on the left side of the controller module, and then slide the controller module out of the system and set it aside.  <b>Note:</b> Ensure that you support the bottom of the controller module with your free hand.

## Removing a PCIe card

To remove a PCIe card from the system, you must perform a specific sequence of steps.

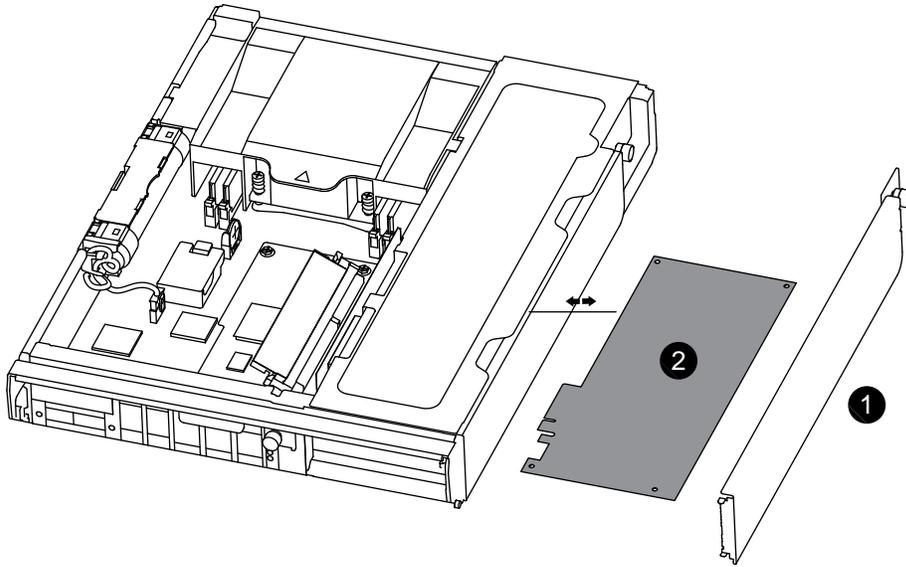
### Steps

1. Loosen the thumbscrew on the controller module side panel.

<b>If you have a...</b>	<b>Then...</b>
A 8020 system	There is a single side panel with two PCIe cards accessible.
Any other 80xx system	There are two side panels with two PCIe cards accessible per side of the controller module.

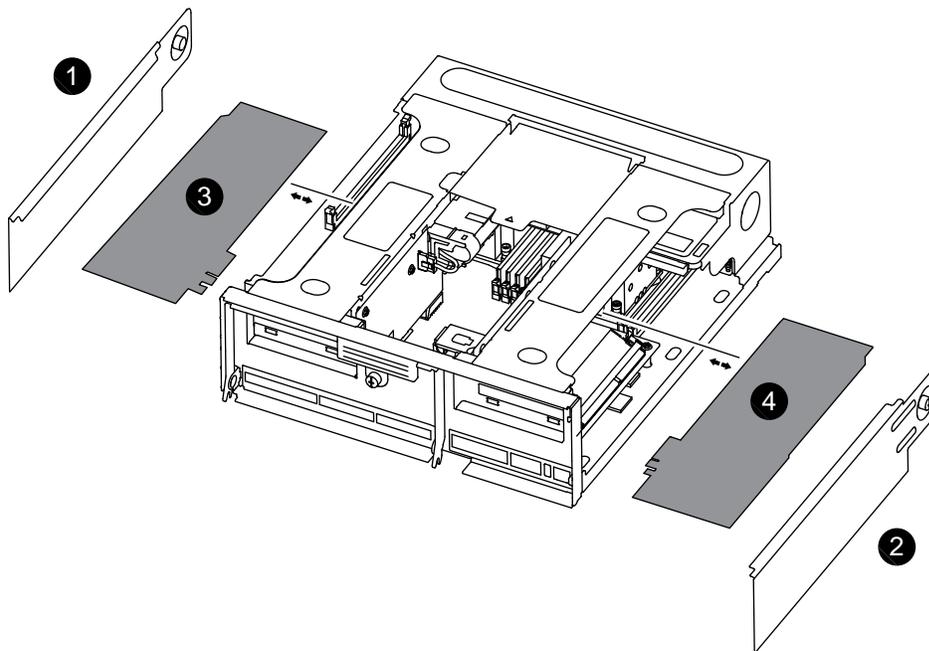
2. Swing the side panel or panels off the controller module.

The following illustrations show the side panels and PCIe cards on 80xx controller modules.



**8020 controller module**

<p>1</p>	<p>Side panel</p>
<p>2</p>	<p>PCIe card</p>



**Other 80xx systems**

<p>1 and 2</p>	Side panels
<p>3 and 4</p>	PCIe cards

3. Remove the PCIe card from the controller module and set it aside.
4. Repeat these steps as needed.

## Installing a PCIe card

To install a PCIe card in the system, you must perform a specific sequence of steps.

### Steps

1. Open the controller module side panel, if necessary, and slide off the PCIe card filler plate, as needed.
2. Install the PCIe card.

Be sure that you properly align the card in the slot and exert even pressure on the card when seating it in the socket. The adapter must be fully and evenly seated in the slot.

3. Repeat these steps as needed for additional PCIe cards.
4. Close the side panel and tighten the thumbscrew.

## Configuring and cabling CNA ports (80xx systems only)

If you are adding a controller module to an 80xx system, you must check the configuration of the CNA ports on the new controller module and, if necessary, change the default settings to match the CNA port configuration of the existing controller module.

### Before you begin

You must have installed the SFP+ modules for the CNA ports.

### Steps

1. Enter Maintenance mode:
  - a. Boot the system: `boot_ontap`.
  - b. Press **Ctrl-C** when you see the message `Press Ctrl-C for Boot Menu`
  - c. Answer **y** when prompted by the system.
  - d. Select the Maintenance mode option from the displayed menu.
2. Check how the ports are currently configured on the new controller module:

`ucadmin show`

The system displays output similar to the following example:

## Example

```
*> ucaadmin show
Adapter      Current Mode  Current Type  Pending Mode  Pending Type  Status
-----
0e           fc      initiator -      -      online
0f           cna     target  -      -      online
0g           cna     target  -      -      online
...
```

3. If the current SFP+ module does not match the desired use, replace the module with the correct SFP+ module.
4. If the current configuration does not match the desired use, change the configuration as required:

If the desired use is for..	Then enter the following command..
FC initiator	<code>ucaadmin modify -t initiator adapter_name</code>
FC target	<code>ucaadmin modify -t target adapter_name</code>
Ethernet	<code>ucaadmin modify -m cna adapter_name</code>

5. If you have changed the default settings, reboot the new controller module to implement the configuration changes.
6. After the new controller module reboots, verify the settings:  
`ucaadmin show`
7. Cable the port.

## Reinstalling the controller module and booting the system

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it to a state where you can run diagnostic tests on the replaced component.

### About this task

For HA pairs with two controller modules in the same chassis, the sequence in which you reinstall the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

### Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.  
**Note:** Do not completely insert the controller module in the chassis until instructed to do so.
2. Recable the system, as needed.  
If you removed the media converters (SFPs), remember to reinstall them if you are using fiber optic cables.
3. Complete the reinstall of the controller module:

If your system is in...	Then perform these steps...						
An HA pair	<p>a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.</p> <p><b>Attention:</b> Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors.</p> <p>b. Enter one of the following commands from the healthy node's console and wait for the giveback to complete:</p> <table border="1" data-bbox="649 493 1518 745"> <thead> <tr> <th data-bbox="649 493 1079 535">For systems operating in...</th> <th data-bbox="1079 493 1518 535">Issue the command...</th> </tr> </thead> <tbody> <tr> <td data-bbox="649 535 1079 588">7-Mode</td> <td data-bbox="1079 535 1518 588"><code>cf giveback</code></td> </tr> <tr> <td data-bbox="649 588 1079 745">Clustered Data ONTAP</td> <td data-bbox="1079 588 1518 745"> <ul style="list-style-type: none"> <li>In Data ONTAP 8.2 or later:  <code>storage failover giveback - ofnode impaired_node_name</code></li> </ul> </td> </tr> </tbody> </table> <p>c. If you have not already done so, reinstall the cable management device, and then tighten the thumbscrew on the cam handle on back of the controller module.</p> <p>d. Bind the cables to the cable management device with the hook and loop strap.</p>	For systems operating in...	Issue the command...	7-Mode	<code>cf giveback</code>	Clustered Data ONTAP	<ul style="list-style-type: none"> <li>In Data ONTAP 8.2 or later:  <code>storage failover giveback - ofnode impaired_node_name</code></li> </ul>
For systems operating in...	Issue the command...						
7-Mode	<code>cf giveback</code>						
Clustered Data ONTAP	<ul style="list-style-type: none"> <li>In Data ONTAP 8.2 or later:  <code>storage failover giveback - ofnode impaired_node_name</code></li> </ul>						
A stand-alone configuration	<p>a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.</p> <p><b>Attention:</b> Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors.</p> <p>b. Reconnect the power cables to the power supplies and to the power sources, turn on the power to start the boot process.</p> <p>c. If you have not already done so, reinstall the cable management device, and then tighten the thumbscrew on the cam handle on back of the controller module.</p> <p>d. Bind the cables to the cable management device with the hook and loop strap.</p>						

4. Run Config Advisor, and check the disk paths.

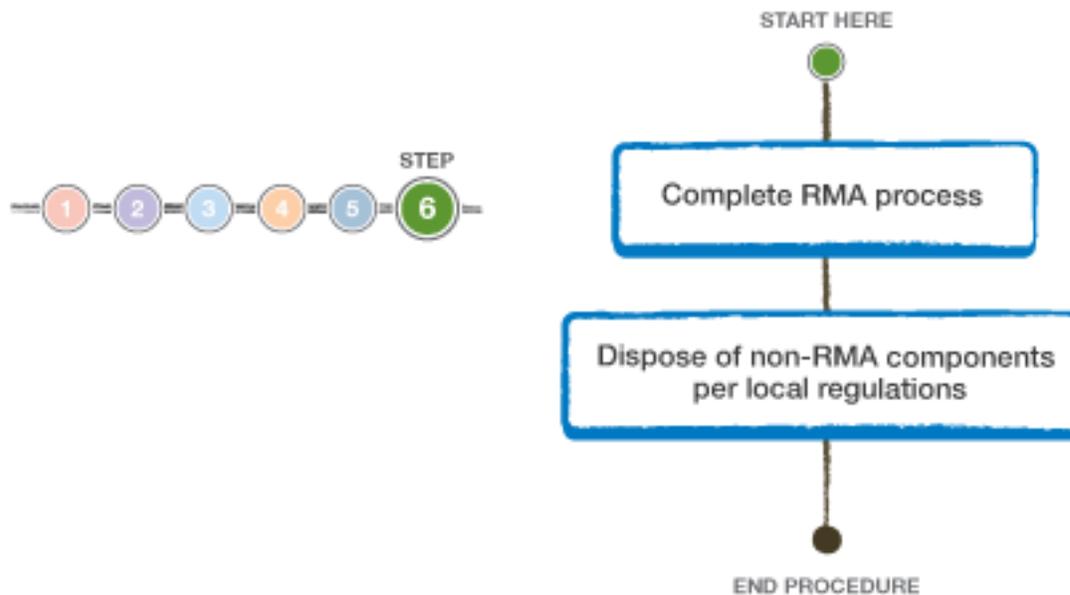
Config Advisor is a configuration validation and health check utility for NetApp systems that runs a series of commands on the hardware and then verifies cabling, configuration, availability, and conformance with NetApp best practice settings.

**Note:** You can check disk paths when the PCIe card you are replacing is a SAS or FC card.

## Completing the replacement process

After you replace the part, you can return the failed part to NetApp, as described in the RMA instructions shipped with the kit. Contact technical support at [NetApp Support](#), 888-463-8277 (North America), 00-800-44-638277 (Europe), or

+800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.



## Copyright information

Copyright © 1994–2016 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.

**RESTRICTED RIGHTS LEGEND:** Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

## Trademark information

Active IQ, AltaVault, Arch Design, ASUP, AutoSupport, Campaign Express, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlexArray, FlexCache, FlexClone, FlexGroup, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, Fueled by SolidFire, GetSuccessful, Helix Design, LockVault, Manage ONTAP, MetroCluster, MultiStore, NetApp, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANscreen, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, SolidFire Helix, StorageGRID, SyncMirror, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

## How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[doccomments@netapp.com](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277