

---

# Replacing an I/O expansion module or PCIe card in an I/O expansion module

Replacing the I/O expansion module involves shutting down the system in which the I/O expansion module is installed and moving all cards and adapters from the old I/O expansion module to the new I/O expansion module.

## Before you begin

All other components in the system must be functioning properly; if not, you must contact technical support.

## About this task

- You can only replace the I/O expansion module in a system that already has an I/O expansion module installed.
- You must shut down the system that contains the I/O expansion module before you replace it.
- You can use this procedure with all versions of Data ONTAP supported by your system.

## Steps

1. [Shutting down the target controller](#) on page 1
2. [Replacing the I/O expansion module](#) on page 6
3. [Replacing a PCIe card in the I/O expansion module](#) on page 9
4. [Completing the replacement process](#) on page 15

## Shutting down the target controller

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

### Choices

- [Shutting down a node operating in 7-Mode](#) on page 1
- [Shutting down a controller module in a two-node MetroCluster configuration](#) on page 2
- [Shutting down a node operating in clustered Data ONTAP](#) on page 4

## Shutting down a node operating in 7-Mode

When performing maintenance on a system operating in 7-Mode, you must shut down the node. Depending on your system's configuration, you might also need to turn off the power supplies.

### About this task

Your system's configuration determines whether you turn off the power supplies after shutting down the node:

- If you have one controller module in the chassis that is either part of an HA pair or in a stand-alone configuration, you must turn off the power supplies in the impaired node chassis.

## Shutting down a node in an HA pair

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

### Steps

1. Check the HA status of the impaired node by entering the following command from either node in the HA pair that is displaying the Data ONTAP prompt:

```
cf status
```

2. Take the appropriate action based on the takeover status of the node.

If the impaired node...	Then...
Has been taken over by the healthy node and is halted	Go to the next step.
Has not been taken over by the healthy node and is running	Take over the impaired node from the prompt of the healthy node: <b>cf takeover</b>

3. Wait at least two minutes after takeover of the impaired node to ensure that the takeover was completed successfully.
4. With the impaired node showing the `waiting for giveback` message or halted, shut it down, depending on your configuration:

If the Service Processor (SP)...	Then...
Is configured	Log in to the SP, and then enter the following command: <b>system power off</b>
Is not configured, and the system is in a dual-chassis HA pair in which each controller is in a separate chassis	Manually shut down the power supplies on the impaired node.

5. If the nodes are in a dual-chassis HA pair, unplug the impaired node power cords from the power source.

## Shutting down a node in a stand-alone configuration

For a node that is not configured with a high-availability (HA) partner, you must perform a clean shutdown (ensuring that all data has been written to disk) and disconnect the power supplies.

### Steps

1. Shut down the node if it is not already shut down:

```
halt -t 0
```

2. Shut down the power supplies, and then unplug both power cords from the source.

The system is ready for maintenance.

## 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 clustered Data 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 true for both health and eligibility.	Proceed to Step 3.
The impaired node shows false for health.	Proceed to the next step.
Any nodes show false for eligibility.	Resolve any cluster issues as needed before continuing with this procedure.
Any nodes other than the impaired node show false 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 node operating in clustered Data ONTAP

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

### Steps

1. If the system is running clustered Data ONTAP, check the status of the nodes in the cluster:

- a. Change to the advance privilege level:

```
set -privilege advanced
```

- b. Enter the following command at the system console of either node:

```
cluster show -epsilon *
```

The command produces output similar to the following:

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>a. Exit advanced mode: <code>set -privilege admin</code></li> <li>b. Proceed to Step 3.</li> </ol>

<b>If...</b>	<b>Then...</b>
All nodes show <code>true</code> for both health and eligibility and Epsilon is assigned to the impaired node	Complete the following steps to move Epsilon: <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.	Complete the following steps: <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>
The impaired node shows <code>false</code> for health and is not the Epsilon node.	<ol style="list-style-type: none"> <li>Exit advanced mode: <code>set -privilege admin</code></li> <li>Proceed to the next step.</li> </ol>
Any nodes show <code>false</code> for eligibility.	<ol style="list-style-type: none"> <li>Resolve any cluster issues as needed before continuing with this procedure.</li> <li>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>Correct the problems that cause the health issues on the nodes before continuing with this procedure.</li> <li>Exit advanced mode: <code>set -privilege admin</code></li> </ol>
<ol style="list-style-type: none"> <li>If the impaired node is part of an HA pair, disable <code>auto-giveback</code> from the console of the healthy node: <code>storage failover modify -node local -auto-giveback false</code></li> <li>Bring the impaired node to the LOADER prompt:</li> </ol>	
<b>If the impaired node is in...</b>	<b>Then...</b>
A stand-alone configuration and is running	Halt the impaired node: <code>system -node halt <i>impaired_node_name</i></code>

<b>If the impaired node is in...</b>	<b>Then...</b>
A stand-alone configuration and is not running and not at the LOADER prompt	Resolve any issues that caused the node to quit running, power cycle it, and then halt the boot process by entering <b>Ctrl-C</b> and responding <b>Y</b> to take the node to the LOADER prompt.
An HA pair	<p>If the impaired node is at the LOADER prompt, it is ready for service. Otherwise, take the applicable action:</p> <ul style="list-style-type: none"> <li>If the impaired node is showing the Data ONTAP prompt, take over the impaired node over from the healthy node and be prepared to interrupt the reboot: <pre>storage failover takeover -ofnode <i>impaired_node_name</i></pre> When prompted to interrupt the reboot, press <b>Ctrl-C</b> to go to the LOADER prompt. <p><b>Note:</b> In a two-node cluster, if the impaired node holds Epsilon, you must move it to the healthy node before halting the impaired node.</p> </li> <li>If the impaired node display is showing <code>Waiting for giveback</code>, press <b>Ctrl-C</b> and respond <b>Y</b> to take the node to the LOADER prompt.</li> <li>If the impaired node is not showing <code>Waiting for giveback</code> or is not showing a Data ONTAP prompt, power cycle the node. Contact technical support if the node does not respond to the power cycle.</li> </ul>

#### 4. Shut down the impaired node.

**Note:** If the system is in an HA pair the node should be at the LOADER prompt.

The method you use to shut down the node depends on whether remote management through a Service Processor (SP) is used, and whether the system is in a dual-chassis or single-chassis configuration.

<b>If the SP is...</b>	<b>Then...</b>
Configured	Log in to the impaired node SP and turn off the power using the following command: <pre>system power off</pre>
Not configured, and the system is in a single-chassis HA pair in which both controllers are in the same chassis and share power supplies	At the impaired node prompt, press <b>Ctrl-C</b> and respond <b>Y</b> to halt the node.

#### 5. If the system is in a dual-chassis HA pair or stand-alone configuration, turn off the power supplies, and then unplug the impaired node power cords from the power source.

## Replacing the I/O expansion module

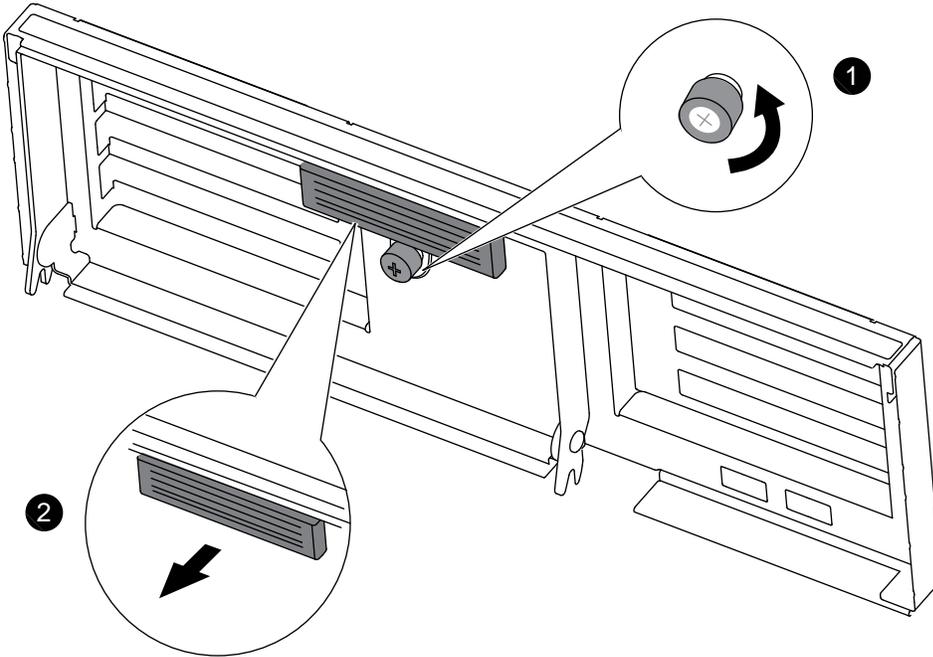
You must shut down the target controller module, remove the cabling from the I/O expansion module (IOXM), remove the module from the chassis, move the PCIe cards to the replacement IOXM, install the replacement IOXM into the chassis, recable it, and reboot the controller.

### 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 cables and SFPs (if required) from the I/O expansion module, and keep track of where the cables were connected.

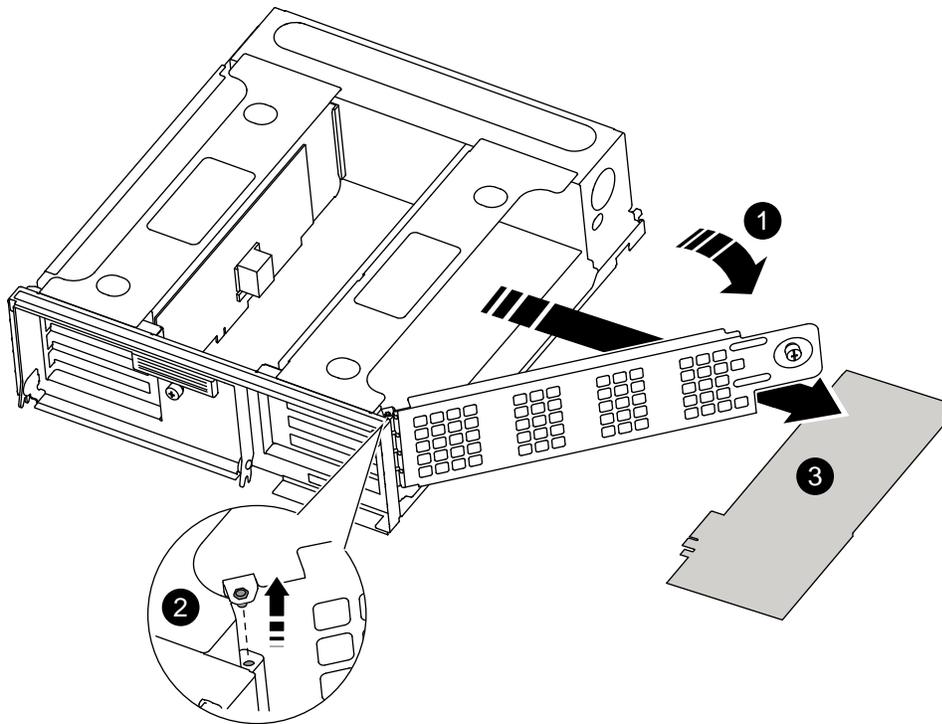
Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

3. Remove the cable management device and set it aside.
4. Loosen the thumbscrew on the cam handle.



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

5. Pull the cam handle down and slide the I/O expansion module out of the chassis until it catches, press the release tab on the side of the module and slide it out of the system, and then place it on an anti-static surface.
6. Loosen the thumb screws on the side panels, swing the side panels away from the module, and remove them from the I/O expansion module.



<b>1</b>	<p>Side panel</p> <p>There are side panels on the left and right sides of the I/O expansion module.</p>
<b>2</b>	<p>Side panel hinge point</p>
<b>3</b>	<p>PCIe card</p> <p>There can be up to eight PCIe cards in the I/O expansion module; four per side. The ports are labeled slot 5 through slot 12. There are also two HA ports directly under slot 12</p>

Repeat this step on the replacement I/O expansion module.

7. Remove the PCIe cards from the old I/O expansion module one at a time, and install them in the same location in the replacement I/O expansion module.
8. Reinstall the side panels on the both I/O expansion modules and tighten the thumbscrews.
9. Align the replacement I/O expansion module with the bottom opening of the bottom bay of the target chassis (labeled B on the bay), and gently slide the I/O expansion module all the way into the chassis using the cam handle.  
The I/O expansion module cam handle begins to engage prior to fully seating into the chassis. Use the cam handle to fully seat the module into the chassis.
10. Push the cam handle all the way to the closed position, and then tighten the thumbscrew on the cam handle.
11. Reinstall the cable management device, and recable the I/O expansion module.
12. Go to [Booting the system](#) on page 9.

## Booting the system

After you have installed the replacement I/O expansion module, you must boot the system.

### Before you begin

The replacement I/O expansion module must be installed and cabled, and the cable management device must be installed.

### Steps

1. Reconnect the system to the power sources, and then turn on the power.

The system begins to reboot.

2. Take one of the following actions, depending on your configuration:

If your system is in...	Then...
A stand-alone configuration	Go to the next step.
An HA configuration	Enter the <code>cf giveback</code> command (7-Mode) or <code>storage failover giveback impaired_node_name</code> command (clustered Data ONTAP) from the healthy node console when you see the <code>Waiting for giveback...</code> message.

3. Verify that all of the PCIe cards are listed for the I/O expansion module:

```
sysconfig -v
```

If any PCIe cards do not appear in the output, verify that they are properly seated, and then rerun the command.

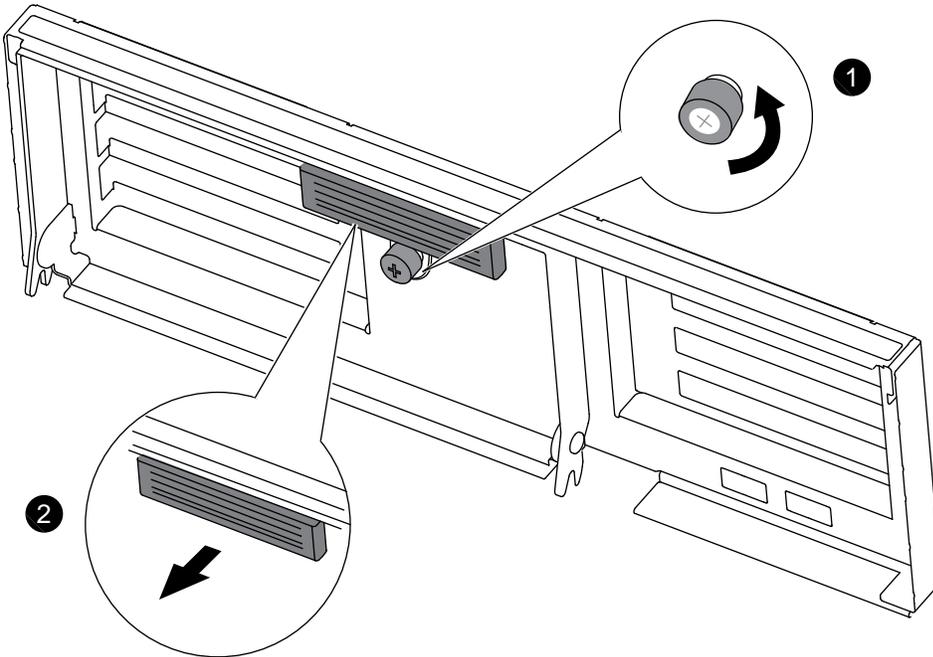
4. Go to [Completing the replacement process](#) on page 15.

## Replacing a PCIe card in the I/O expansion module

You must shut down the node with the faulty PCIe card, remove the I/O expansion module, remove the side panels, remove the faulty card, install the replacement card, reassemble the I/O expansion module, reinstall it in the chassis, and then reboot the node.

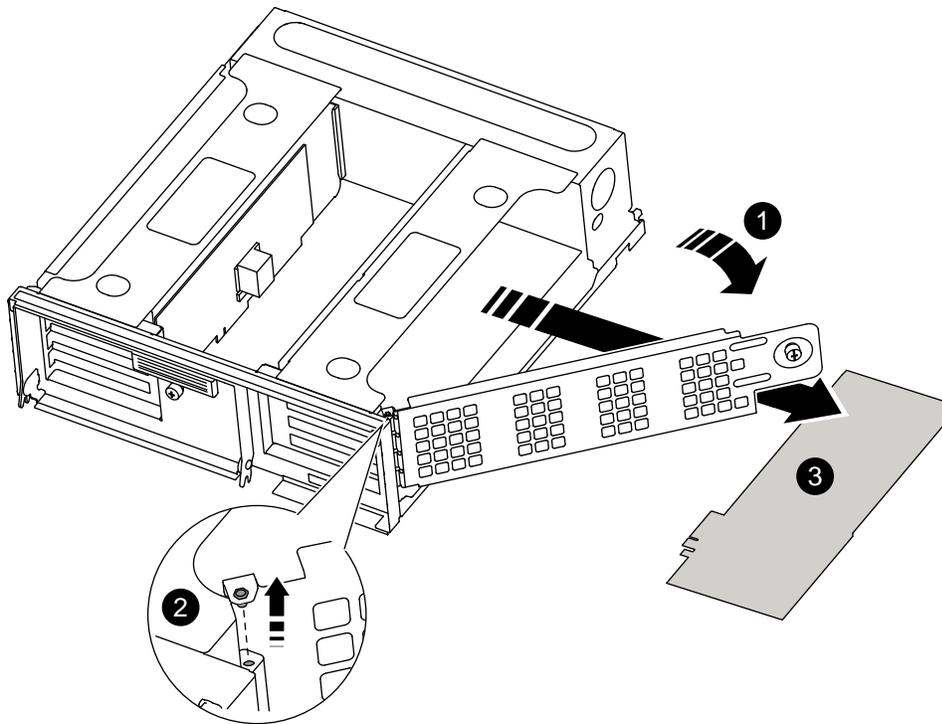
### Steps

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 cables and SFPs (if required) from the I/O expansion 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 and set it aside.
4. Loosen the thumbscrew on the cam handle.



1	Thumbscrew
2	Cam handle

5. Pull the cam handle down and slide the I/O expansion module out of the chassis until it catches, press the release tab on the side of the module and then slide it out of the system, and then place it on an anti-static surface.
6. Loosen the thumb screws on the side panels, swing the side panels away from the module, and remove them from both I/O expansion modules.



<b>1</b>	<p><b>Side panel</b></p> <p>There is a side panel on the left and right sides of the I/O expansion module.</p>
<b>2</b>	<p><b>Side panel hinge point</b></p>
<b>3</b>	<p><b>PCIe card</b></p> <p>There can be up to eight PCIe cards in the I/O expansion module; four per side. The ports are labeled slot 5 through slot 12.</p>

7. Remove the target PCIe card from the I/O expansion module.
8. Install the replacement card by aligning it with the slot in the I/O expansion module and then firmly, but gently, pushing the card down into the slot.
9. Reinstall the side panel and tighten the thumbscrew.
10. Align the replacement I/O expansion module with the bottom opening of the bottom bay of the target chassis (labeled B on the bay), and gently slide the I/O expansion module all the way into the chassis using the cam handle.  
 The I/O expansion module cam handle begins to engage prior to fully seating into the chassis. Use the cam handle to fully seat the module into the chassis.
11. Push the cam handle all the way to the closed position, and then tighten the thumbscrew on the cam handle.
12. Reinstall the cable management device, and recable the I/O expansion module.

## Running diagnostics on a PCIe card

After installing a PCIe card, you should run diagnostics.

### Before you begin

You need to observe the following prerequisites depending on the version of Data ONTAP you are running:

If you are running a version of...	Then...
Data ONTAP 8.2.1 or later	You do not have to recable your PCIe cards prior to beginning diagnostics tests. External loopback tests are disabled by default.
Data ONTAP prior to 8.2.1	<p>You need to cable your PCIe cards according to the following guidelines prior to running diagnostics tests. If you do not have loopback cables, skip running diagnostics and bring the system back online.</p> <ul style="list-style-type: none"><li>NIC interfaces must be cabled in a pair-wise manner, with adjacent ports connected together. Ports e0M and e0P must not be connected together because of the internal switch connection.</li><li>SAS card ports must be connected to storage or connected to the adjacent SAS ports. If the SAS cards or systems with more than two ports are not connected to storage, they must be connected to ports A to B, C to D, and so forth.</li><li>The FC card ports must be connected to storage or the ports must be terminated with loopback plugs.</li></ul>

### Steps

1. If the boot process was interrupted too late and the Boot Menu appeared, perform the following steps:

- a. Select the Maintenance mode option from the displayed menu.
- b. After the system boots to Maintenance mode, enter the following command at the prompt:

```
halt
```

After you issue the command, wait until the system stops at the LOADER prompt.

**Important:** During the boot process, you might see the following prompt:

- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy node remains down.

You can safely respond **y** to the prompt.

2. On the node with the replaced component, enter the following command at the LOADER prompt:

```
boot_diags
```

**Note:** You must enter this command from the LOADER prompt for system-level diagnostics to function properly. The `boot_diags` command starts special drivers designed specifically for system-level diagnostics.

**Important:** During the `boot_diags` process, you might see the following prompt:

- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the partner remains down.

You can safely respond **y** to the prompt.

The Maintenance mode prompt (\*>) appears.

3. Clear the status logs by entering the following command:

```
sldiag device clearstatus
```

4. Verify that the log is cleared by entering the following command:

```
sldiag device status
```

The following default response is displayed:

```
SLDIAG: No log messages are present.
```

5. View the types of devices available for checking by entering the following command:

```
sldiag device show
```

6. Run the test for the type of PCI card you installed by entering the following command at the Maintenance mode prompt:

```
sldiag device run -dev fcal/sas/ata
```

See the *System-Level Diagnostics Guide* for a description of other types of tests you can run.

7. View the status of the test by entering the following command:

```
sldiag device status
```

Your storage system provides the following output while the tests are still running:

```
There are still test(s) being processed.
```

After all the tests are complete, the following response appears by default:

```
*> <SLDIAG: _ALL_TESTS_COMPLETED>
```

8. Verify that no hardware problems resulted from the addition or replacement of hardware components on your system by entering the following command:

```
sldiag device status [-dev devtype] [-name device] -long -state failed
```

System-level diagnostics returns you to the prompt if there are no test failures, or lists the full status of failures resulting from testing the component.

9. Proceed based on the result of the preceding step:

---

<b>If the system-level diagnostics tests...</b>	<b>Then...</b>
---	----------------

---

Were completed without any failures

- a. Clear the status logs by entering the following command:  
**sldiag device clearstatus**
- b. Verify that the log is cleared by entering the following command:  
**sldiag device status**  
The following default response is displayed:  
SLDIAG: No log messages are present.
- c. Exit Maintenance mode by entering the following command:  
**halt**
- d. Enter the following command at the LOADER prompt to boot the storage system:  
**boot\_ontap**
- e. Return the replacement node to normal operation:

<b>If your system is in a...</b>	<b>Issue this command from the partner's console...</b>
HA pair running 7-Mode Data ONTAP	<code>cf giveback</code>
HA pair running clustered Data ONTAP	<code>storage failover giveback</code>
Two-node MetroCluster running clustered Data ONTAP	<code>metrocluster switchback</code>
Standalone configuration	None required

You have completed system-level diagnostics.

---

If the system-level diagnostics tests...	Then...
Resulted in some test failures	<p data-bbox="609 220 966 254">Determine the cause of the problem:</p> <ol style="list-style-type: none"> <li data-bbox="609 275 1234 308">a. Exit Maintenance mode by entering the following command: <b>halt</b> After you issue the command, wait until the system stops at the LOADER prompt.</li> <li data-bbox="609 401 1536 455">b. Turn off or leave on the power supplies, depending on how many controller modules are in the chassis: <ul style="list-style-type: none"> <li data-bbox="649 478 1502 533">• If you have two controller modules in the chassis, leave the power supplies turned on to provide power to the other controller module.</li> <li data-bbox="649 556 1453 611">• If you have one controller module in the chassis, turn off the power supplies and unplug them from the power sources.</li> </ul> </li> <li data-bbox="609 634 1485 722">c. Check the controller module you are servicing and verify that you have observed all the considerations identified for running system-level diagnostics, that cables are securely connected, and that hardware components are properly installed in the storage system.</li> <li data-bbox="609 745 1536 989">d. Boot the controller module you are servicing, interrupting the boot by pressing <b>Ctrl-C</b> when prompted. This takes you to the Boot Menu: <ul style="list-style-type: none"> <li data-bbox="649 823 1510 911">• If you have two controller modules in the chassis, fully seat the controller module you are servicing in the chassis. The controller module boots up when fully seated.</li> <li data-bbox="649 934 1502 989">• If you have one controller module in the chassis, connect the power supplies and turn them on.</li> </ul> </li> <li data-bbox="609 1012 1120 1045">e. Select Boot to maintenance mode from the menu.</li> <li data-bbox="609 1068 1234 1123">f. Exit Maintenance mode by entering the following command: <b>halt</b> After you issue the command, wait until the system stops at the LOADER prompt.</li> <li data-bbox="609 1188 1388 1222">g. Enter <b>boot_diags</b> at the prompt and rerun the system-level diagnostic test.</li> </ol>

#### Related information

[System-Level Diagnostics Guide - now.netapp.com/NOW/knowledge/docs/hardware/NetApp/diag/sldiag.pdf](http://now.netapp.com/NOW/knowledge/docs/hardware/NetApp/diag/sldiag.pdf)

## 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](http://support.netapp.com), 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.

## Disposing of batteries

Dispose of batteries according to local regulations regarding battery recycling or disposal. If you cannot properly dispose of the battery, return it to NetApp, as described in the RMA instructions shipped with the kit.

#### Related information

[Safety Information and Regulatory Notices at support.netapp.com](http://support.netapp.com)

## Copyright information

Copyright © 1994–2015 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

NetApp, the NetApp logo, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, 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 at <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 to [docomments@netapp.com](mailto:docomments@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