
Replacing a DIMM

You must replace a DIMM in the controller module when your system registers an increasing number of correctable error correction codes (ECC); failure to do so causes a system panic.

Before you begin

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

You must replace the failed component with a replacement FRU component you received from your provider.

About this task

- You can use this procedure with all versions of Data ONTAP supported by your system.
- This procedure refers to system DIMMs; if you need to replace the NVRAM9 DIMM, see the procedure *Replacing the NVRAM9 battery and/or DIMM in a 80xx system*.

Steps

1. [Shutting down the target controller](#) on page 1
2. [Opening the system](#) on page 6
3. [Removing the DIMMs](#) on page 8
4. [Installing the DIMMs](#) on page 10
5. [Reinstalling the controller module and booting the system](#) on page 10
6. [Running diagnostics on the DIMM](#) on page 12
7. [Completing the replacement process](#) on page 14

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 controller module in an HA pair](#) on page 1
- [Shutting down a controller module in a two-node MetroCluster configuration](#) on page 4
- [Shutting down a stand-alone controller module](#) on page 6

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 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.	<p>a. Exit advanced mode:</p> <pre>set -privilege admin</pre> <p>b. Proceed to Step 3.</p>
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> <p>a. Remove Epsilon from the node:</p> <pre>cluster modify -node node1 -epsilon false</pre> <p>b. Assign Epsilon to a node in the cluster:</p> <pre>cluster modify -node node4 -epsilon true</pre> <p>c. Exit advanced mode:</p> <pre>set -privilege admin</pre> <p>d. Go to Step 3.</p>

If...	Then...
The impaired node shows <code>false</code> for health and is the Epsilon node.	Complete the following steps: <ol style="list-style-type: none"> a. Change to the advance privilege level: <code>set -privilege advanced</code> b. Remove Epsilon from the node: <code>cluster modify -node node1 -epsilon false</code> c. Assign Epsilon to a node in the cluster: <code>cluster modify -node node4 -epsilon true</code> d. Exit advanced mode: <code>set -privilege admin</code> e. Proceed to the next step.
The impaired node shows <code>false</code> for health and is not the Epsilon node.	<ol style="list-style-type: none"> a. Exit advanced mode: <code>set -privilege admin</code> b. Proceed to the next step.
Any nodes show <code>false</code> for eligibility.	<ol style="list-style-type: none"> a. Resolve any cluster issues as needed before continuing with this procedure. b. Exit advanced mode: <code>set -privilege admin</code>
Any nodes other than the impaired node show <code>false</code> for health.	<ol style="list-style-type: none"> a. Correct the problems that cause the health issues on the nodes before continuing with this procedure. b. Exit advanced mode: <code>set -privilege admin</code>
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>
Clustered Data ONTAP	<code>storage failover show</code>
3. Take one of the following actions, depending on the result of the <code>cf status</code> or <code>storage failover show</code> 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.

If...	Then...
The impaired node took over the healthy node	<ol style="list-style-type: none"> a. Correct the problem that caused the takeover. b. Enter the cf giveback command (7-Mode) or storage failover giveback <i>impaired_node_name</i> command (clustered Data ONTAP) from the impaired node console. c. Go back to the Step 1.

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	cf takeover
Clustered Data ONTAP	<ul style="list-style-type: none"> • In Data ONTAP 8.2 or later: storage failover takeover -ofnode <i>impaired_node_name</i>

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

- b. Wait at least two minutes after takeover of the impaired node to ensure 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: system power off
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 Ctrl-C and respond Y 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 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

```

mcc1A::*> metrocluster operation show
Operation: Switchover
Start time: 10/4/2012 19:04:13
State: in-progress
End time: -
Errors:

mcc1A::*> 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 Ctrl-C and respond Y to halt the node.

- 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

- 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.

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

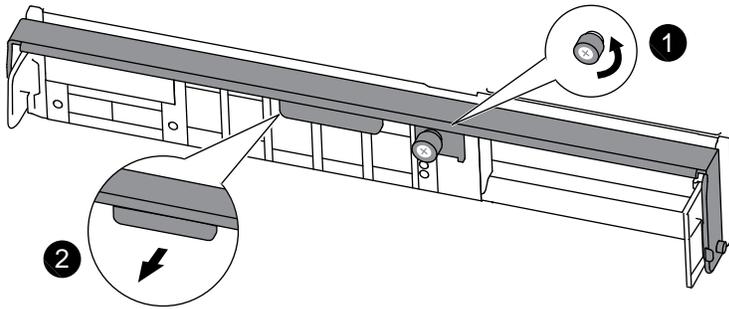
- If you are not already grounded, properly ground yourself.
- Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, and keep track of where the cables were connected.

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

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

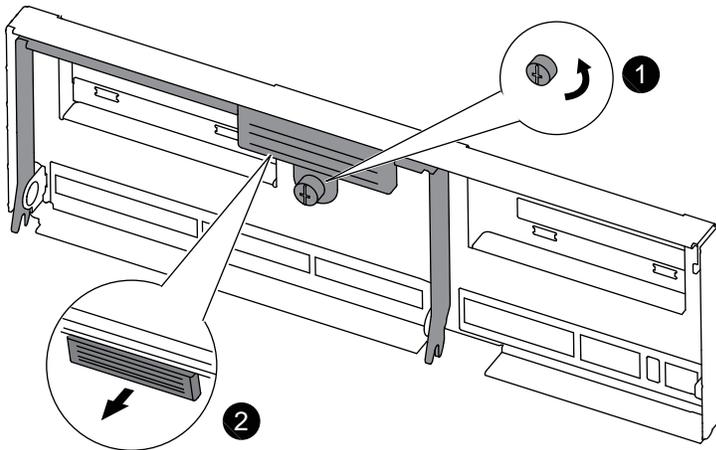
Note: The 8020 controller module uses cable management arms, while the rest of the 80xx family uses a cable management tray.

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

1	Thumbscrew
2	Cam handle

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

If you have...	Then...
A 8020 system	Slide the controller module completely out of the system. Note: Make sure that you support the bottom of the controller module with your free hand and set it aside.

If you have...	Then...
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. Note: Make sure that you support the bottom of the controller module with your free hand.

Removing the DIMMs

To remove a DIMM, you must perform a specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Open the CPU cover.
3. If you have a 80xx system other than an 8020, you might have to remove the PCIe cards from the controller module to locate and remove the DIMM you are replacing.

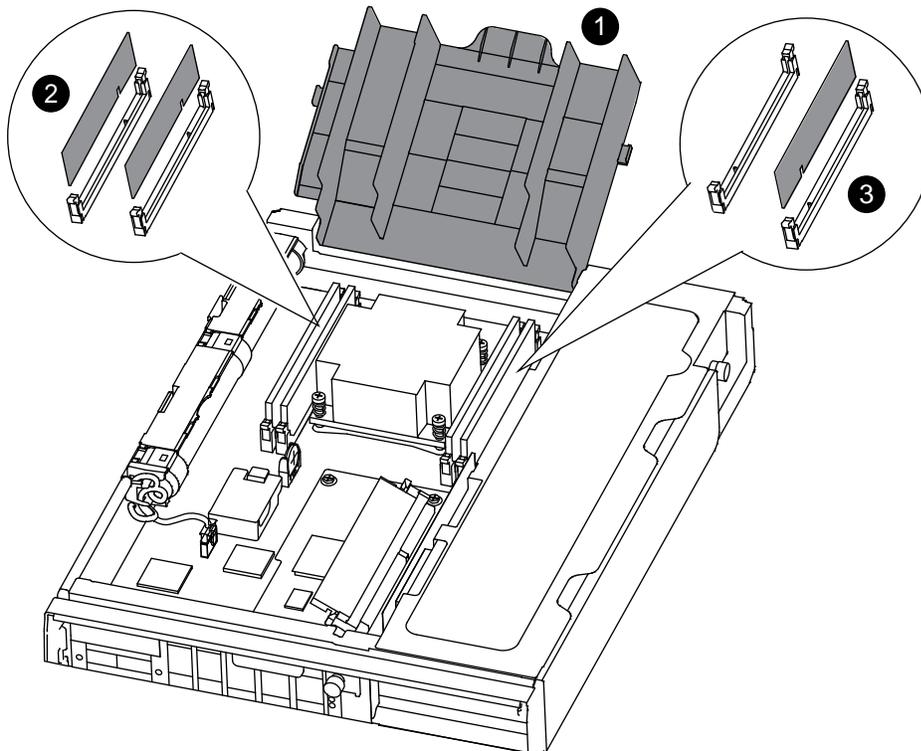
See *Replacing PCI cards in a 80xx system* for instructions for removing PCIe cards.

4. Open the CPU cover.
5. Locate the DIMMs.

Note: Each system memory DIMM has an LED located on the board next to each DIMM slot. The LED is illuminated if the DIMM needs to be replaced.

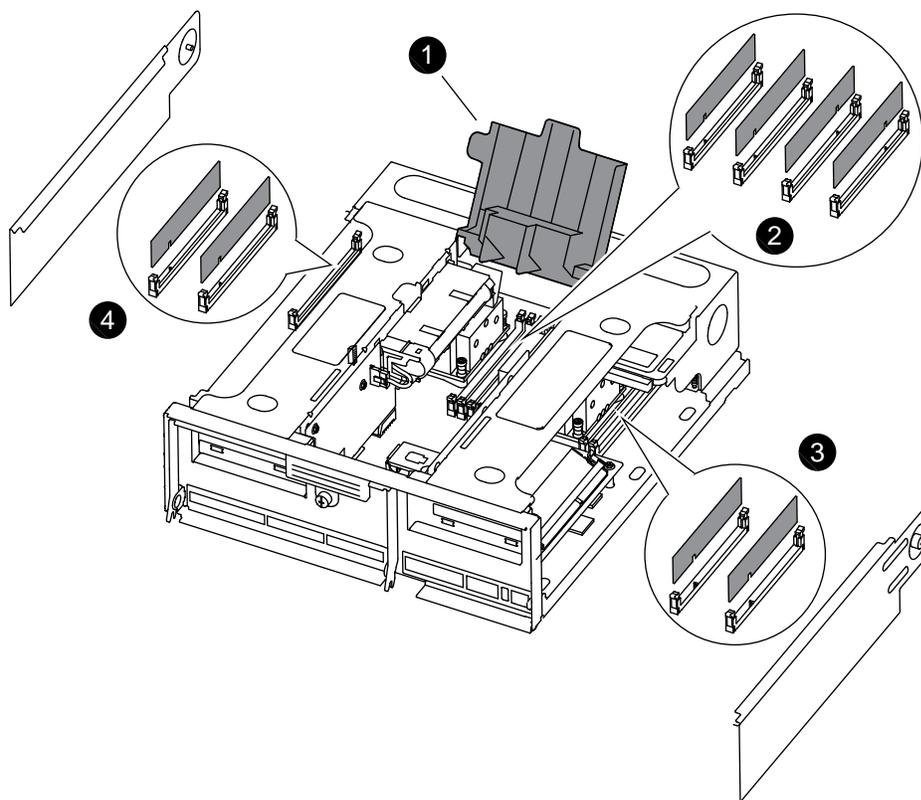
The number and placement of 80xx system DIMMs depends on the model of your system. 8020 have three DIMMs, and other models in the 80xx have up to eight DIMMs.

The following illustrations show the location of system DIMMs in 80xx systems.



8020 system DIMMs

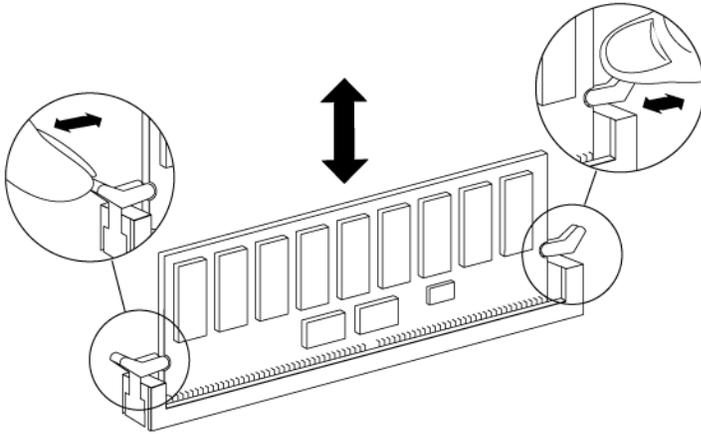
1	CPU cover
2	DIMM 1 and DIMM 2
3	DIMM 3 and DIMM 4 (The DIMM 4 slot is not populated.)



Other 80xx system DIMMs

1	CPU cover
2	DIMM 6, DIMM 5, DIMM 1, DIMM 2 (left to right)
3	Side panel, DIMM 4, and DIMM 3 (left to right)
4	Side panel, DIMM 7, and DIMM 8 (left to right)

6. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
7. Slowly press down on the two DIMM ejector tabs, one at a time, to eject the DIMM from its slot, and then lift it out of the slot.



Attention: Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

8. Repeat these steps to remove additional DIMMs as needed.

Installing the DIMMs

To install one or more DIMMs, you must perform a specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the slot where you are installing the new DIMM.
8020 systems use three 8 GB DIMMs, 8040 systems use four 8 GB DIMMs, and 8060 systems use eight 8 GB DIMMs.
3. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it over the slot. The notch among the pins on the DIMM should line up with the tab in the socket.
4. Insert the DIMM squarely into the slot.
The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.
Attention: Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.
5. Push carefully, but firmly, on the top edge of the DIMM until the latches snap into place over the notches at the ends of the DIMM.
6. Close the CPU cover.
7. If you have a 80xx system other than the 8020, reinstall the PCIe cards, if necessary, and then close and lock the side panel.

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	<ol style="list-style-type: none">a. Be prepared to interrupt the boot process. The controller module begins to boot as soon as it is fully seated in the chassis.b. 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. Attention: Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors.c. Boot to Maintenance mode by entering halt to go to the LOADER prompt:<ul style="list-style-type: none">• If you are running Data ONTAP 8.2.1 and earlier, enter boot_ontap, and press Ctrl-C when prompted to go to the boot menu, and then select Maintenance mode from the menu.• If you are running Data ONTAP 8.2.2 and later, enter boot_ontap maint at the LOADER prompt.d. 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.e. Bind the cables to the cable management device with the hook and loop strap.
A stand-alone configuration	<ol style="list-style-type: none">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. Attention: Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors.b. Reconnect the power cables to the power supplies and to the power sources, turn on the power to start the boot process, and then press Ctrl-C to interrupt the boot process when you see the message Press Ctrl-C for Boot Menu. Note: If you miss the prompt and the controller module boots to Data ONTAP, enter halt and at the LOADER prompt enter boot_ontap, and press Ctrl-C when prompted, and then repeat this step.c. From the boot menu, select the option for Maintenance mode.d. 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.e. Bind the cables to the cable management device with the hook and loop strap.

Running diagnostics on the DIMM

After installing a new DIMM, you should run diagnostics.

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. Check the status of the system memory by entering the following command at the prompt:

```
sldiag device run -dev mem
```

6. 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>
```

7. 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.

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

If the system-level diagnostics tests...	Then...
--	---------

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:

If your system is in a...	Issue this command from the partner's console...
HA pair running 7-Mode Data ONTAP	cf giveback
HA pair running clustered Data ONTAP	storage failover giveback
Two-node MetroCluster running clustered Data ONTAP	metrocluster switchback
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 1242 308">a. Exit Maintenance mode by entering the following command: halt After you issue the command, wait until the system stops at the LOADER prompt. <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 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 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 data-bbox="609 745 1536 989">d. Boot the controller module you are servicing, interrupting the boot by pressing Ctrl-C when prompted. This takes you to the Boot Menu: <ul style="list-style-type: none"> <li data-bbox="649 823 1518 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 data-bbox="649 934 1502 989">• If you have one controller module in the chassis, connect the power supplies and turn them on. <li data-bbox="609 1012 1128 1045">e. Select Boot to maintenance mode from the menu. <li data-bbox="609 1068 1242 1123">f. Exit Maintenance mode by entering the following command: halt After you issue the command, wait until the system stops at the LOADER prompt. <li data-bbox="609 1188 1388 1222">g. Enter boot_diags at the prompt and rerun the system-level diagnostic test.

Related information

[System-Level Diagnostics Guide](#)

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 © 1994–2015 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. 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