
Replacing the NVRAM battery and/or DIMM

To replace an NVRAM battery or NVRAM DIMM in the system, you must remove the controller module from the system, open it, replace the battery, and close and replace the controller module.

About this task

- You can use this procedure with all versions of Data ONTAP supported by your system.
- All other components in the system must be functioning properly; if not, you must contact technical support.

Steps

1. [Shutting down the target controller](#) on page 1
2. [Opening the system](#) on page 7
3. [Removing the NVRAM battery](#) on page 8
4. [Removing the NVRAM DIMM](#) on page 10
5. [Installing the NVRAM DIMMs](#) on page 13
6. [Installing the NVRAM battery](#) on page 13
7. [Reinstalling the controller module and booting the system](#) on page 13
8. [Checking the status of the battery and running diagnostics](#) on page 15
9. [Completing the replacement process](#) on page 17

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 node operating in clustered Data ONTAP](#) on page 3
- [Shutting down a controller module in a two-node MetroCluster configuration](#) on page 5

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: <pre>cf takeover</pre>

3. Wait at least two minutes after takeover of the impaired node to ensure that the takeover was completed successfully.
4. If the impaired node is not at the LOADER prompt and is showing `waiting for giveback`, press **Ctrl-C** and respond **Y** to halt the node.
5. 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: <pre>system power off</pre>
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.

6. 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 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.	<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>
<p>2. 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></p> <p>3. Bring the impaired node to the LOADER prompt:</p>	
If the impaired node is in...	Then...
A stand-alone configuration and is running	Halt the impaired node: <code>system -node halt <i>impaired_node_name</i></code>
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 Ctrl-C and responding Y to take the node to the LOADER prompt.

If the impaired node is in...	Then...
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"> 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 Ctrl-C to go to the LOADER prompt. <p>Note: 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> If the impaired node display is showing <code>Waiting for giveback</code>, press Ctrl-C and respond Y to take the node to the LOADER prompt. 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.

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.

If the SP is...	Then...
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 Ctrl-C and respond Y 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.

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:

- 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.
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 Ctrl-C and respond y to halt the node.

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

Opening the system

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

Steps

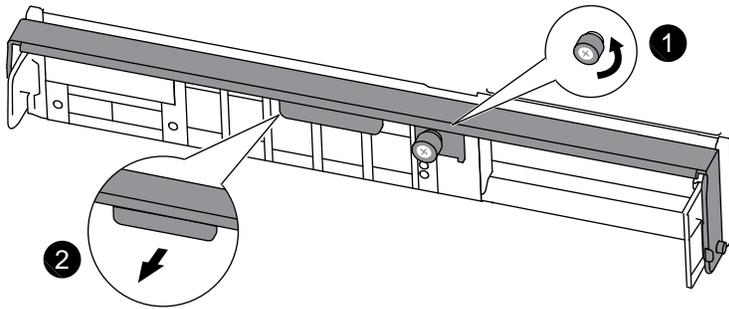
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, and keep track of where the cables were connected.

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.

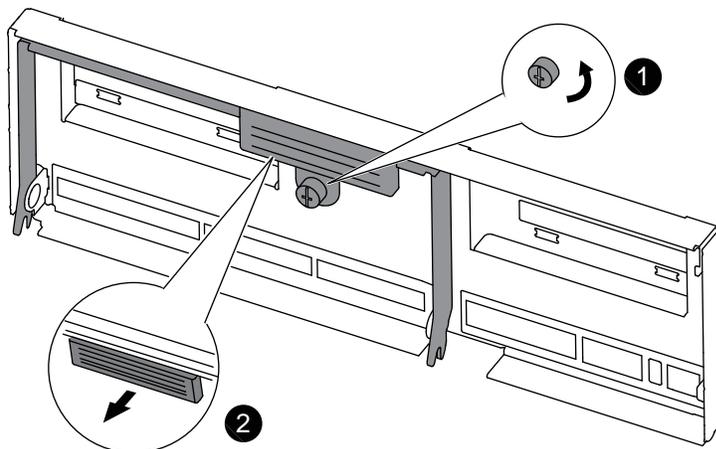
Note: The 8020 controller module uses cable management arms, while the rest of the 80xx family uses a cable management tray.

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

1	Thumbscrew
2	Cam handle

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

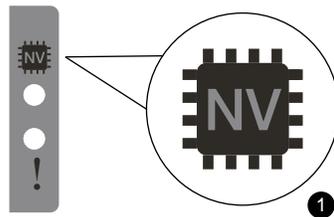
If you have...	Then...
A 8020 system	Slide the controller module completely out of the system. Note: Make sure that you support the bottom of the controller module with your free hand and set it aside.
Any other 80xx model	Slide the controller module out of the system until it catches, press the release 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 NVRAM battery

The process of removing the NVRAM battery from a controller module in a 80xx system is similar between the models. The difference between the models is the physical appearance of the battery packs and where it connects to the system.

Steps

1. Check the NVRAM LED:

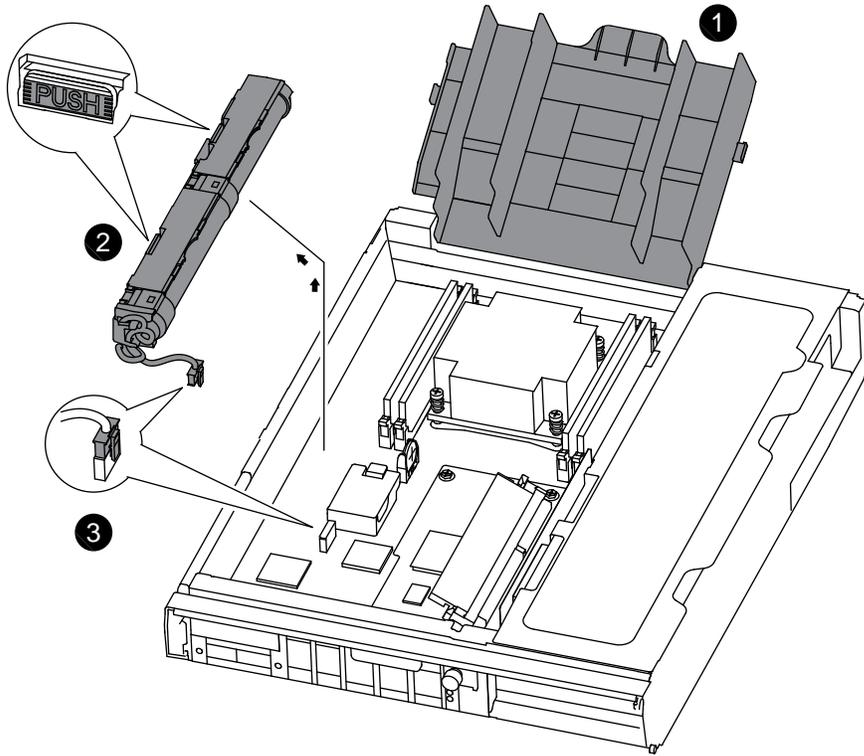


If your system is in...	Then...
HA configuration	Go to the next step.
Syand-alone configuration	Cleanly shut down the controller module, and then check the NVRAM LED.

Attention: The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

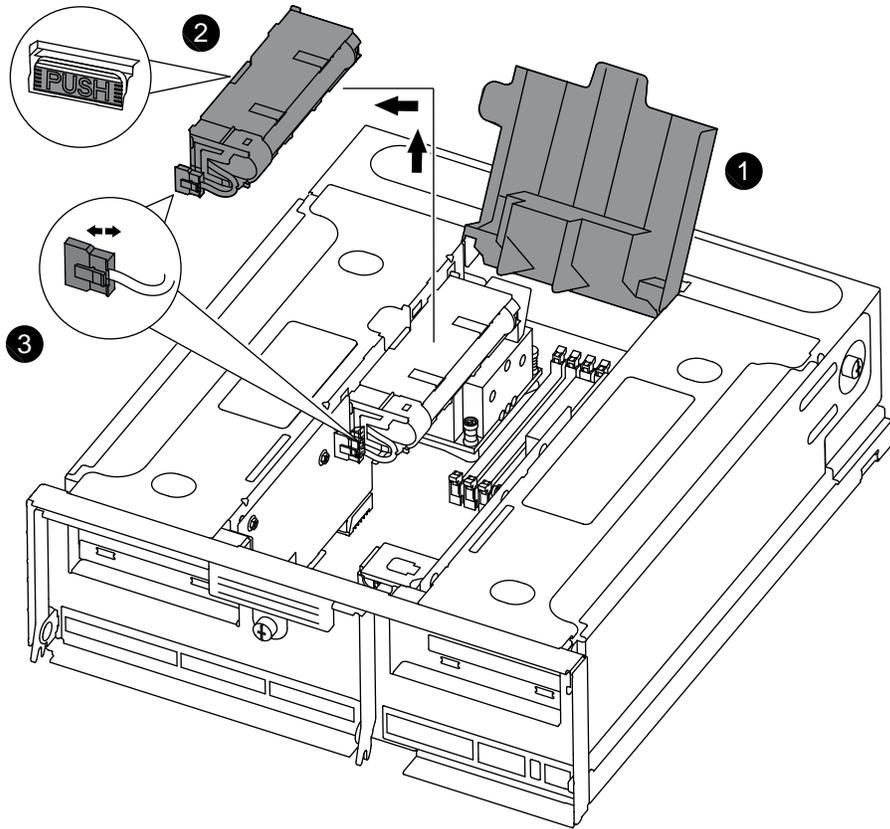
- If power is lost without a clean shutdown, the NVRAM LED flashes until the destage is complete, and then the LED turns off.
- If the LED is on and power is on, unwritten data is stored on NVRAM that has not been destaged. This typically occurs during an uncontrolled shutdown after Data ONTAP successfully booted.

2. Open the CPU air duct and locate the NVRAM battery.



8020 controller module

1	CPU air duct
2	NVRAM battery pack
3	NVRAM battery plug



Other 80xx controller modules

1	CPU air duct
2	NVRAM battery pack
3	NVRAM battery plug

3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Press the blue locking tab or tabs on the edge of the battery, and then slide the battery out of the controller module.

Note: There are two locking tabs on the 8020 battery housing, and there is only one locking tab on the battery housing on the rest of the 80xx models.

Removing the NVRAM DIMM

To remove the NVRAM DIMM, you must perform a specific sequence of steps, beginning with disconnecting the NVRAM battery and ending with removing the NVRAM DIMM from the socket.

Before you begin

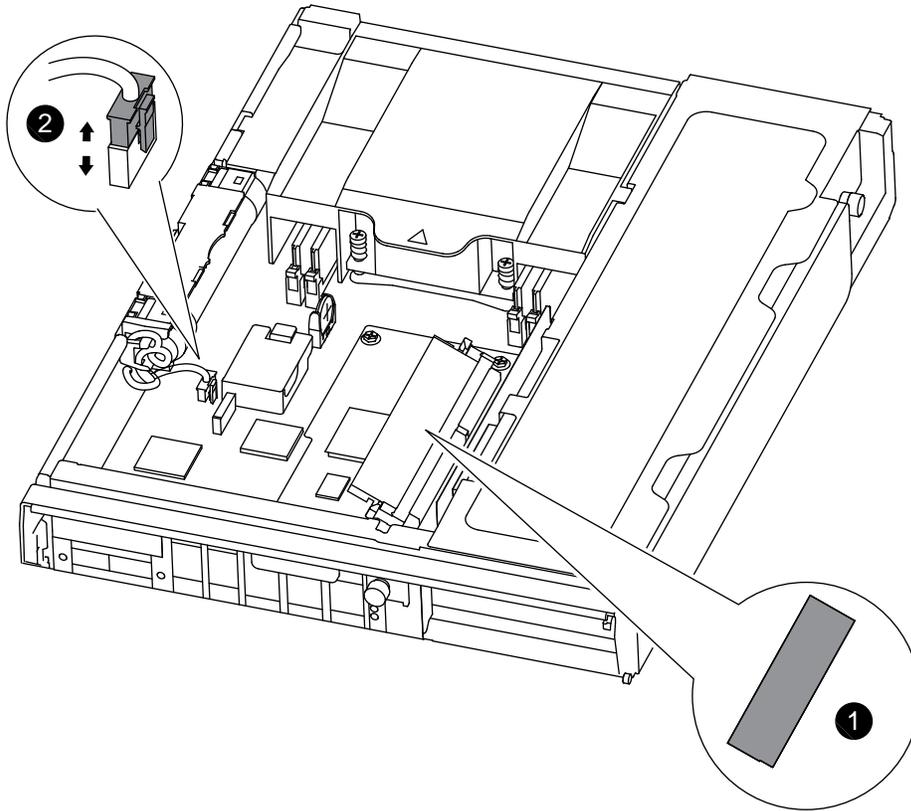
You must unplug the NVRAM battery from the controller module before you remove the NVRAM DIMM.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the NVRAM battery and disconnect it from the controller module.
3. Locate the NVRAM DIMM on the controller module.

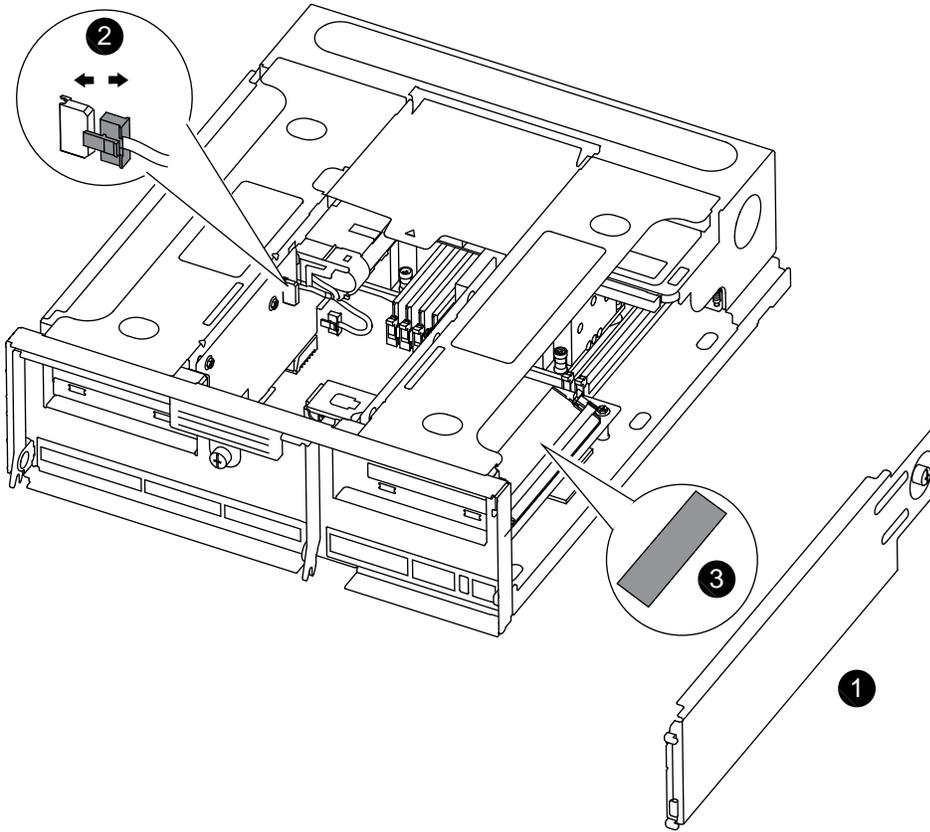
The NVRAM DIMM holder has black ejector tabs, while the system DIMMS have white ejector tabs.

Note: For all models except the 8020, the NVRAM DIMM is only accessible by removing the PCIe side panel, as shown in the illustration. Remove any PCIe cards if you cannot easily access the NVRAM DIMM.



FAS8020 controller module

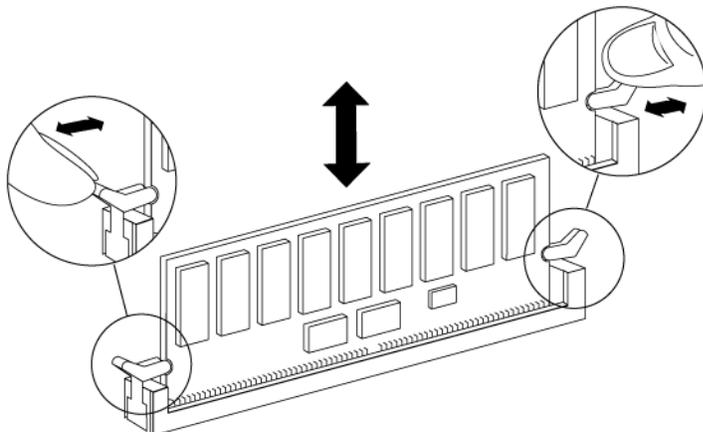
1	NVRAM DIMM
2	NVRAM battery plug



Other FAS80xx controller modules

1	PCIe side panel
2	NVRAM battery plug
3	NVRAM DIMM

4. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
5. 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.

Installing the NVRAM DIMMs

To install the NVRAM DIMM, you must perform a specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the NVRAM DIMM slot where you are installing the replacement NVRAM DIMM.
The NVRAM DIMM slot is located behind the PCIe cover for all models except the 8020, which is located in the middle of the motherboard.
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. If you have a 80xx system other than the 8020, reinstall the PCIe cards, if necessary, and then close and lock the side panel.
7. Plug the NVRAM battery into the controller module.
Make sure that the plug locks down to the socket.

Installing the NVRAM battery

To install the NVRAM battery in the controller module, you must perform a specific sequence of steps.

Steps

1. Align the tab or tabs on the battery holder with the notches in the controller module side, and gently push down on the battery housing until the battery housing clicks into place.
2. Plug the battery plug back into the controller module.
3. Close the CPU air duct.
Make sure that the plug locks down to the socket.

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.

Checking the status of the battery and running diagnostics

After installing a new battery, you should run diagnostics and check the status of the battery.

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. Enter the following command at the prompt:

```
sldiag device run -dev nvram
```

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>Determine the cause of the problem:</p> <ol style="list-style-type: none"> a. Exit Maintenance mode by entering the following command: <code>halt</code> After you issue the command, wait until the system stops at the LOADER prompt. 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"> • If you have two controller modules in the chassis, leave the power supplies turned on to provide power to the other controller module. • If you have one controller module in the chassis, turn off the power supplies and unplug them from the power sources. 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. 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"> • 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. • If you have one controller module in the chassis, connect the power supplies and turn them on. e. Select Boot to maintenance mode from the menu. f. Exit Maintenance mode by entering the following command: <code>halt</code> After you issue the command, wait until the system stops at the LOADER prompt. g. Enter <code>boot_diags</code> 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.

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](#)

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