



Maintain

Install and maintain

NetApp

February 06, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap-systems/fas8300/maintain-overview.html> on February 06, 2026. Always check docs.netapp.com for the latest.

Table of Contents

Maintain	1
Maintain FAS8300 and FAS8700 hardware	1
System components	1
Boot media - automated recovery	2
Boot media automated recovery workflow - FAS8300 and FAS8700	2
Requirements for automated boot media recovery - FAS8300 and FAS8700	3
Shut down the controller for automated boot media recovery - FAS8300 and FAS8700	3
Replace the boot media for automated boot recovery - FAS8300 and FAS8700	4
Automated boot media recovery from the partner node - FAS8300 and FAS8700	8
Return the failed boot media to NetApp - FAS8300 and FAS8700	14
Boot media - manual recovery	15
Boot media manual recovery workflow - FAS8300 and FAS8700	15
Requirements for manual boot media recovery - FAS8300 and FAS8700	15
Check encryption key support and status - FAS8300 and FAS8700	16
Shut down the controller for manual boot media recovery - FAS8300 and FAS8700	19
Replace the boot media and prepare for manual boot recovery - FAS8300 and FAS8700	23
Manual boot media recovery from a USB drive - FAS8300 and FAS8700	28
Restore encryption - FAS8300 and FAS8700	31
Return the failed boot media to NetApp - FAS8300 and FAS8700	41
Replace the caching module - FAS8300 and FAS8700	41
Step 1: Shut down the impaired controller	42
Step 2: Remove the controller module	45
Step 3: Replace a caching module	46
Step 4: Install the controller module	48
Step 5: Restore the controller module to operation	49
Step 7: Switch back aggregates in a two-node MetroCluster configuration	49
Step 8: Complete the replacement process	51
Chassis	51
Overview of chassis replacement - FAS8300 and FAS8700	51
Shut down the controllers - FAS8300 and FAS8700	51
Move and replace hardware - FAS8300 and FAS8700	54
Complete the restoration and replacement process - FAS8300 and FAS8700	56
Controller	58
Overview of controller module replacement - FAS8300 and FAS8700	58
Shut down the impaired controller - FAS8300 and FAS8700	59
Replace the controller module hardware - FAS8300 and FAS8700	63
Restore and verify the system configuration - FAS8300 and FAS8700	74
Recable the system and reassign disks - FAS8300 and FAS8700	75
Complete system restoration - FAS8300 and FAS8700	81
Replace a DIMM - FAS8300 and FAS8700	83
Step 1: Shut down the impaired controller	83
Step 2: Remove the controller module	88
Step 3: Replace system DIMMs	89

Step 4: Install the controller module	90
Step 5: Restore the controller module to operation	92
Step 6: Switch back aggregates in a two-node MetroCluster configuration.	92
Step 7: Return the failed part to NetApp	93
Hot-swap a fan module - FAS8300 and FAS8700	93
Replace an NVDIMM - FAS8300 and FAS8700	95
Step 1: Shut down the impaired controller	95
Step 2: Remove the controller module	99
Step 3: Replace the NVDIMM	100
Step 4: Install the controller module	102
Step 5: Restore the controller module to operation	103
Step 6: Switch back aggregates in a two-node MetroCluster configuration.	103
Step 7: Return the failed part to NetApp	105
Replace the NVDIMM battery - FAS8300 and FAS8700	105
Step 1: Shut down the impaired controller	105
Step 2: Remove the controller module	109
Step 3: Replace the NVDIMM battery	110
Step 4: Install the controller module	111
Step 5: Restore the controller module to operation	112
Step 6: Switch back aggregates in a two-node MetroCluster configuration.	112
Step 7: Return the failed part to NetApp	114
Replace a PCIe or mezzanine card - FAS8300 and FAS8700	114
Step 1: Shut down the impaired controller	114
Step 2: Remove the controller module	118
Step 3: Replace a PCIe card	119
Step 4: Replace the mezzanine card	120
Step 5: Install the controller module	122
Step 6: Switch back aggregates in a two-node MetroCluster configuration.	123
Step 7: Restore the controller module to operation	124
Step 8: Return the failed part to NetApp	124
Hot-swap a power supply - FAS8300 and FAS8700	125
Replace the real-time clock battery - FAS8300 and FAS8700	126
Step 1: Shut down the impaired controller	126
Step 2: Remove the controller module	130
Step 3: Replace the RTC battery	131
Step 4: Reinstall the controller module and sett time/date after RTC battery replacement	132
Step 5: Switch back aggregates in a two-node MetroCluster configuration.	133
Step 6: Return the failed part to NetApp	135

Maintain

Maintain FAS8300 and FAS8700 hardware

Maintain the hardware of your FAS8300 and FAS8700 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the FAS8300 and FAS8700 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the FAS8300 and FAS8700 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the storage system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure .
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure .
Caching module	You must replace the controller's caching module when your system registers a single AutoSupport (ASUP) message that the module has gone offline.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
Fan	The fan cools the controller.
NVDIMM	The NVDIMM (non-volatile dual in-line memory module) manages the data transfer from the volatile memory to the non-volatile storage, and maintains data integrity in the event of a power loss or system shutdown.

NVDIMM battery	A NVDIMM battery is responsible for maintaining power to the NVDIMM module.
PCIe card and risers	A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard or into risers plugged into the motherboard.
Power supply	A power supply provides a redundant power source in a controller shelf.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - FAS8300 and FAS8700

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on the partner node to reinstall ONTAP on the replacement boot media in your FAS8300 or FAS8700 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the [manual boot recovery procedure](#).

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1

[Review the boot media requirements](#)

Review the requirements for boot media replacement.

2

[Shut down the controller](#)

Shut down the controller in your storage system when you need to replace the boot media.

3

[Replace the boot media](#)

Remove the failed boot media from the controller module and install the replacement boot media.

4

[Restore the image on the boot media](#)

Restore the ONTAP image from the partner controller.

5

[Return the failed part to NetApp](#)

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - FAS8300 and FAS8700

Before replacing the boot media in your FAS8300 or FAS8700, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the [manual boot recovery procedure](#).

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
 - /cfcards/kmip/servers.cfg file.
 - /cfcards/kmip/certs/client.crt file.
 - /cfcards/kmip/certs/client.key file.
 - /cfcards/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
 - The *impaired controller* is the controller on which you are performing maintenance.
 - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you [shut down the controller](#).

Shut down the controller for automated boot media recovery - FAS8300 and FAS8700

Shut down the impaired controller in your FAS8300 or FAS8700 storage system to prevent data loss and ensure system stability when replacing the boot media.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the [manual boot recovery procedure](#).

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

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, [quorum status](#) of that node, availability status of that node, and

operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

Steps

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

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter *y* when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> The <i>-halt true</i> parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you [replace the boot media](#).

Replace the boot media for automated boot recovery - FAS8300 and FAS8700

The boot media in your FAS8300 or FAS8700 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the [manual boot recovery procedure](#).

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

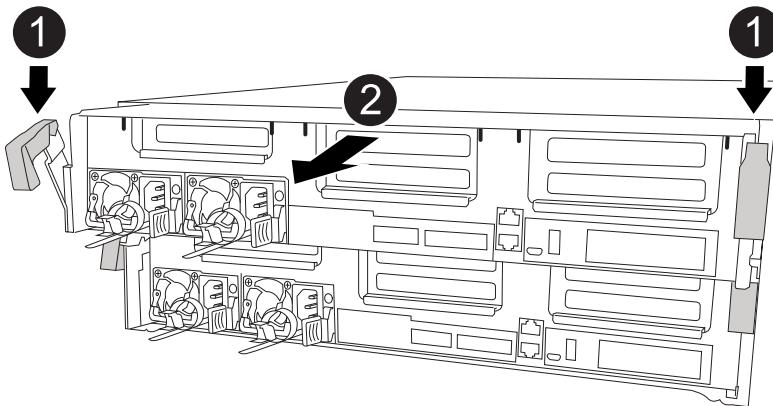
Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. 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, keeping 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.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

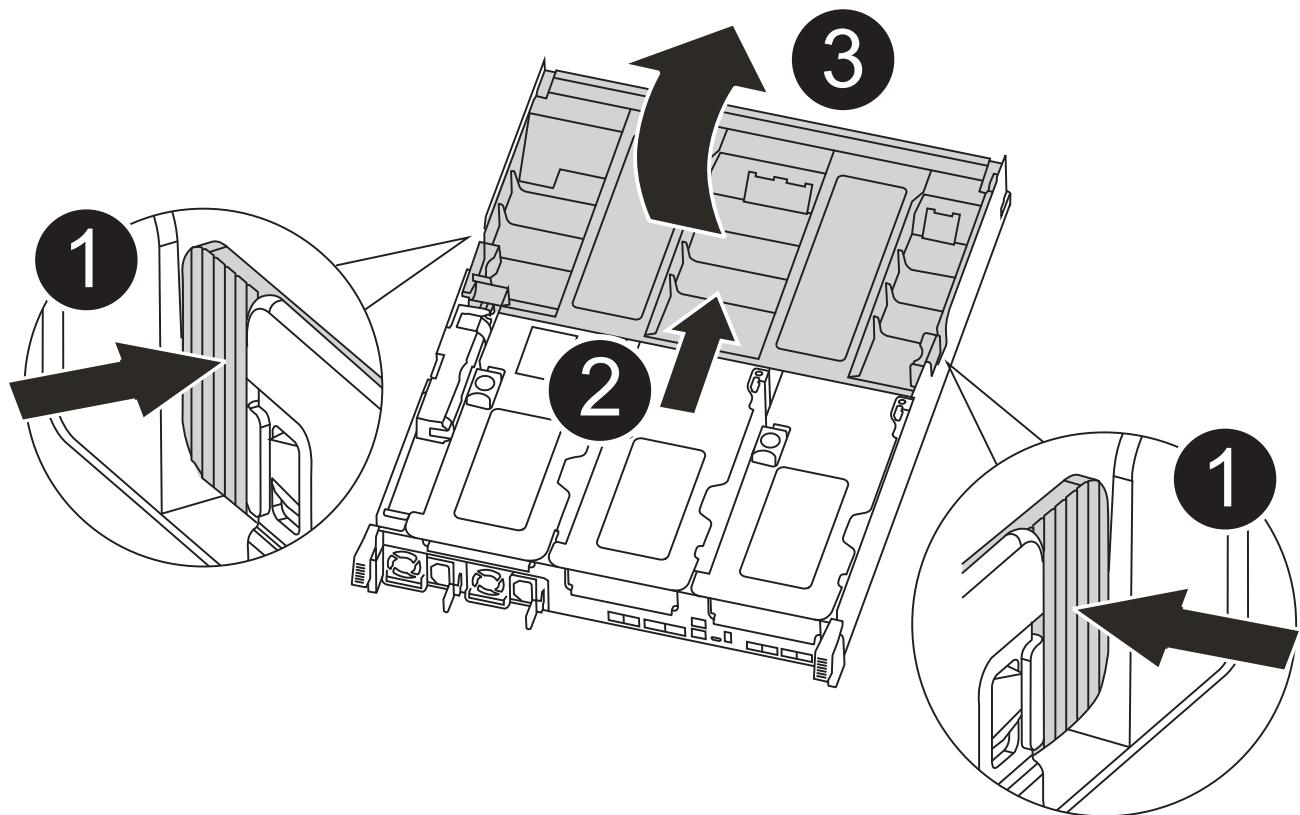


1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

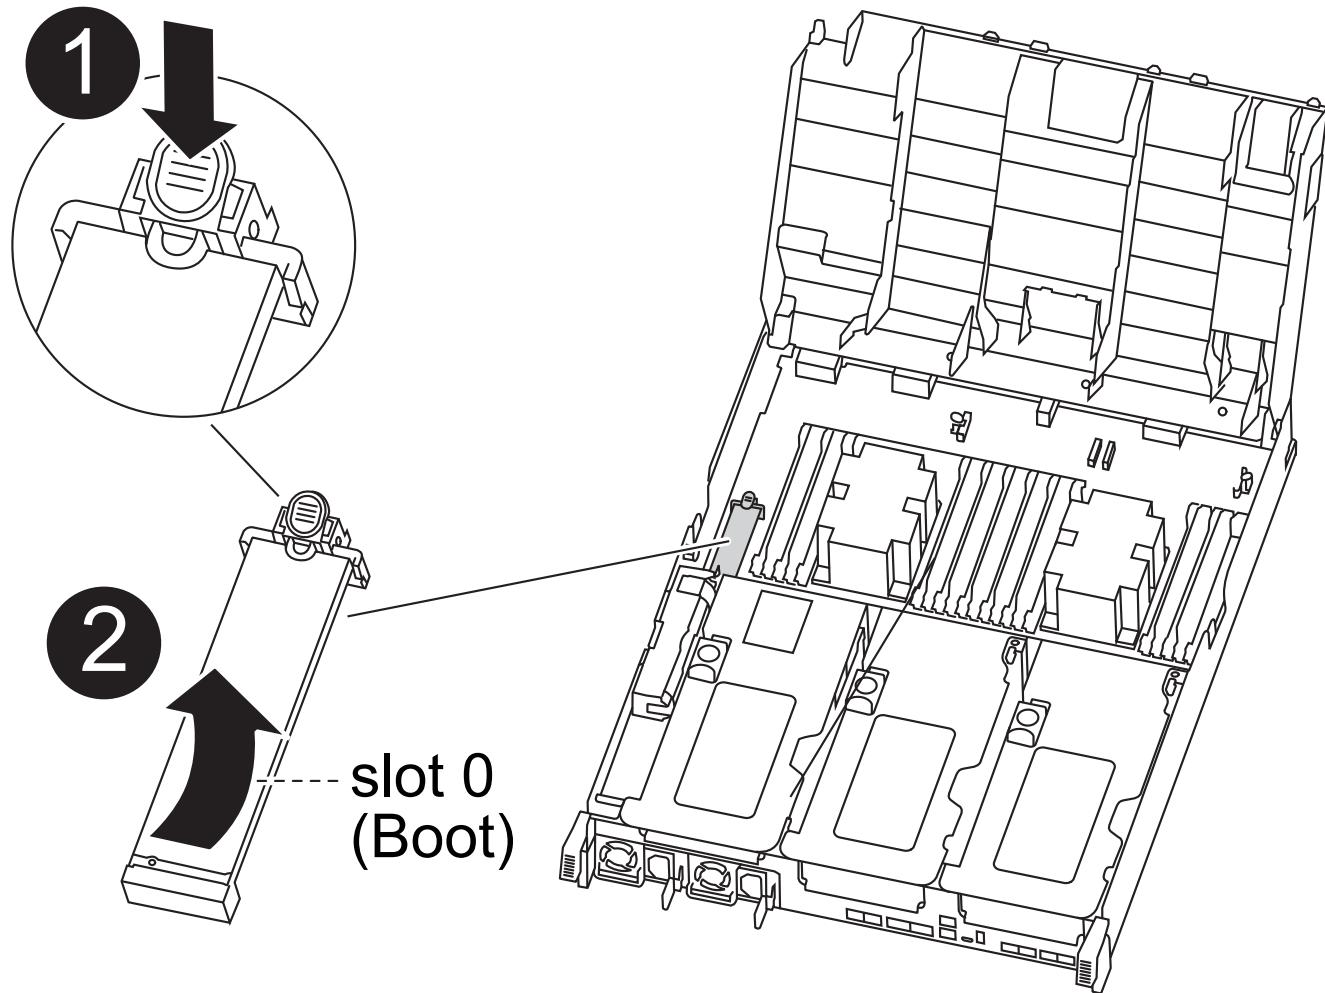
7. Place the controller module on a stable, flat surface.
8. Open the air duct:



1	Locking tabs
2	Slide air duct toward back of controller
3	Rotate air duct up

- a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
- b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

9. Locate and remove the boot media from the controller module:



1	Press blue button
2	Rotate boot media up and remove from socket

- a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
- b. Rotate the boot media up and gently pull the boot media out of the socket.
10. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
11. Check the boot media to make sure that it is seated squarely and completely in the socket.
If necessary, remove the boot media and reseat it into the socket.
12. Lock the boot media in place:
 - a. Rotate the boot media down toward the motherboard.
 - b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
 - c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
13. Close the air duct.

What's next

After physically replacing the impaired boot media, [restore the ONTAP image from the partner node](#).

Automated boot media recovery from the partner node - FAS8300 and FAS8700

After installing the new boot media device in your FAS8300 or FAS8700 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the [manual boot recovery procedure](#).

Before you begin

- Determine your key manager type:
 - Onboard Key Manager (OKM): Requires cluster-wide passphrase and backup data
 - External Key Manager (EKM): Requires the following files from the partner node:
 - /cfcard/kmip/servers.cfg
 - /cfcard/kmip/certs/client.crt
 - /cfcard/kmip/certs/client.key
 - /cfcard/kmip/certs/CA.pem

Steps

1. From the LOADER prompt, start the boot media recovery process:

```
boot_recovery -partner
```

The screen displays the following message:

```
Starting boot media recovery (BMR) process. Press Ctrl-C to abort...
```

2. Monitor the boot media install recovery process.

The process completes and displays the Installation complete message.

3. The system checks for encryption and displays one of the following messages:

If you see this message...	Do this...
key manager is not configured. Exiting.	<p>Encryption is not installed on the system.</p> <ol style="list-style-type: none"> Wait for the login prompt to display. Log into the node and give back the storage: <pre>storage failover giveback -ofnode <i>impaired_node_name</i></pre> <ol style="list-style-type: none"> Go to re-enabling automatic giveback if it was disabled.
key manager is configured.	Encryption is installed. Go to restoring the key manager .



If the system cannot identify the key manager configuration, it displays an error message and prompts you to confirm whether key manager is configured and which type (onboard or external). Answer the prompts to proceed.

4. Restore the key manager using the appropriate procedure for your configuration:

Onboard Key Manager (OKM)

The system displays the following message and begins running BootMenu Option 10:

```
key manager is configured.  
Entering Bootmenu Option 10...
```

```
This option must be used only in disaster recovery procedures. Are  
you sure? (y or n):
```

- a. Enter **y** at the prompt to confirm you want to start the OKM recovery process.
- b. Enter the passphrase for onboard key management when prompted.
- c. Enter the passphrase again when prompted to confirm.
- d. Enter the backup data for onboard key manager when prompted.

Show example of passphrase and backup data prompts

```
Enter the passphrase for onboard key management:  
-----BEGIN PASSPHRASE-----  
<passphrase_value>  
-----END PASSPHRASE-----  
Enter the passphrase again to confirm:  
-----BEGIN PASSPHRASE-----  
<passphrase_value>  
-----END PASSPHRASE-----  
Enter the backup data:  
-----BEGIN BACKUP-----  
<passphrase_value>  
-----END BACKUP-----
```

- e. Monitor the recovery process as it restores the appropriate files from the partner node.

When the recovery process is complete, the node reboots. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....  
Setting recovery material for the onboard key manager  
Recovery secrets set successfully  
Trying to delete any existing km_onboard.keydb file.  
  
Successfully recovered keymanager secrets.
```

- f. After the node reboots, verify that the system is back online and operational.

g. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

h. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster:

```
security key-manager onboard sync
```

Go to [re-enabling automatic giveback](#) if it was disabled.

External Key Manager (EKM)

The system displays the following message and begins running BootMenu Option 11:

```
key manager is configured.  
Entering Bootmenu Option 11...
```

a. Enter the EKM configuration settings when prompted:

i. Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file:

Show example of client certificate contents

```
-----BEGIN CERTIFICATE-----  
<certificate_value>  
-----END CERTIFICATE-----
```

ii. Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file:

Show example of client key file contents

```
-----BEGIN RSA PRIVATE KEY-----  
<key_value>  
-----END RSA PRIVATE KEY-----
```

iii. Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file:

Show example of KMIP server file contents

```
-----BEGIN CERTIFICATE-----  
<KMIP_certificate_CA_value>  
-----END CERTIFICATE-----
```

iv. Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file:

Show example of server configuration file contents

```
xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx
xxx.xxx.xxx.xxx:5696.port=5696
xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem
xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4
1xxx.xxx.xxx.xxx:5696.timeout=25
xxx.xxx.xxx.xxx:5696.nbio=1
xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.crt
xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key
xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:
!RC2:!RC4:!SEED:!eNULL:!aNULL"
xxx.xxx.xxx.xxx:5696.verify=true
xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>
```

v. If prompted, enter the ONTAP Cluster UUID from the partner node. You can check the cluster UUID from the partner node using the `cluster identify show` command.

Show example of ONTAP Cluster UUID prompt

```
Notice: bootarg.mgwd.cluster_uuid is not set or is empty.
Do you know the ONTAP Cluster UUID? {y/n} y
Enter the ONTAP Cluster UUID: <cluster_uuid_value>
```

```
System is ready to utilize external key manager(s).
```

vi. If prompted, enter the temporary network interface and settings for the node:

- The IP address for the port
- The netmask for the port
- The IP address of the default gateway

Show example of temporary network setting prompts

In order to recover key information, a temporary network interface needs to be configured.

```
Select the network port you want to use (for example,  
'e0a')  
e0M
```

```
Enter the IP address for port : xxx.xxx.xxx.xxx  
Enter the netmask for port : xxx.xxx.xxx.xxx  
Enter IP address of default gateway: xxx.xxx.xxx.xxx  
Trying to recover keys from key servers....  
[discover_versions]  
[status=SUCCESS reason= message=]
```

b. Verify the key restoration status:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored. The process restores the appropriate files from the partner node and reboots the node. Proceed to the next step.
- If the key is not successfully restored, the system halts and displays error and warning messages. Rerun the recovery process from the **LOADER** prompt: `boot_recovery -partner`

Show example of key recovery error and warning messages

```
ERROR: kmip_init: halting this system with encrypted
mroot...
WARNING: kmip_init: authentication keys might not be
available.
*****
*          A T T E N T I O N
*
*      System cannot connect to key managers.
*
*****
ERROR: kmip_init: halting this system with encrypted
mroot...
.
Terminated

Uptime: 11m32s
System halting...

LOADER-B>
```

- c. After the node reboots, verify that the system is back online and operational.
- d. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

Go to [re-enabling automatic giveback](#) if it was disabled.

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

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

What's next

After you've restored the ONTAP image and the node is up and serving data, you [return the failed part to NetApp](#).

Return the failed boot media to NetApp - FAS8300 and FAS8700

If a component in your FAS8300 or FAS8700 system fails, return the failed part to NetApp. See the [Part Return and Replacements](#) page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - FAS8300 and FAS8700

Get started with replacing the boot media in your FAS8300 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the [automated boot recovery procedure](#). If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

1

Review the boot media requirements

Review the requirements for replacing the boot media.

2

Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.

3

Shut down the controller

Shut down the controller when you need to replace the boot media.

4

Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.

5

Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

6

Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONTAP boot menu.

7

Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - FAS8300 and FAS8700

Before replacing the boot media in your FAS8300 or FAS8700 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you

have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the [automated boot recovery procedure](#). If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to [check encryption key support and status on the boot media](#).

Check encryption key support and status - FAS8300 and FAS8700

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the [automated boot recovery procedure](#). If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Check NVE support and download the correct ONTAP image

Determine whether your ONTAP version supports NetApp Volume Encryption (NVE) so you can download the correct ONTAP image for the boot media replacement.

Steps

1. Check if your ONTAP version supports encryption:

```
version -v
```

If the output includes `10no-DARE`, NVE is not supported on your cluster version.

2. Download the appropriate ONTAP image based on NVE support:

- If NVE is supported: Download the ONTAP image with NetApp Volume Encryption

- If NVE is not supported: Download the ONTAP image without NetApp Volume Encryption



Download the ONTAP image from the NetApp Support Site to your HTTP or FTP server or a local folder. You will need this image file during the boot media replacement procedure.

Step 2: Verify key manager status and back up configuration

Before shutting down the impaired controller, verify the key manager configuration and back up the necessary information.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	<pre>security key-manager keystore show</pre> <ul style="list-style-type: none"> • If EKM is enabled, EKM is listed in the command output. • If OKM is enabled, OKM is listed in the command output. • If no key manager is enabled, No key manager keystores configured is listed in the command output.
ONTAP 9.13.1 or earlier	<pre>security key-manager show-key-store</pre> <ul style="list-style-type: none"> • If EKM is enabled, external is listed in the command output. • If OKM is enabled, onboard is listed in the command output. • If no key manager is enabled, No key managers configured is listed in the command output.

2. Depending on whether a key manager is configured on your system, do one of the following:

If no key manager is configured:

You can safely shut down the impaired controller and proceed to the shutdown procedure.

If a key manager is configured (EKM or OKM):

- a. Enter the following query command to display the status of the authentication keys in your key manager:

```
security key-manager key query
```

- b. Review the output and check the value in the Restored column. This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Complete the appropriate procedure based on your key manager type:

External Key Manager (EKM)

Complete these steps based on the value in the Restored column.

If all keys show true in the Restored column:

You can safely shut down the impaired controller and proceed to the shutdown procedure.

If any keys show a value other than true in the Restored column:

- Restore the external key management authentication keys to all nodes in the cluster:

```
security key-manager external restore
```

If the command fails, contact NetApp Support.

- Verify that all authentication keys are restored:

```
security key-manager key query
```

Confirm that the Restored column displays true for all authentication keys.

- If all keys are restored, you can safely shut down the impaired controller and proceed to the shutdown procedure.

Onboard Key Manager (OKM)

Complete these steps based on the value in the Restored column.

If all keys show true in the Restored column:

- Back up the OKM information:

- Switch to advanced privilege mode:

```
set -priv advanced
```

Enter y when prompted to continue.

- Display the key management backup information:

```
security key-manager onboard show-backup
```

- Copy the backup information to a separate file or your log file.

You will need this backup information if you need to manually recover OKM during the replacement procedure.

- Return to admin mode:

```
set -priv admin
```

- You can safely shut down the impaired controller and proceed to the shutdown procedure.

If any keys show a value other than true in the Restored column:

- a. Synchronize the onboard key manager:

```
security key-manager onboard sync
```

Enter the 32-character alphanumeric onboard key management passphrase when prompted.



This is the cluster-wide passphrase you created when you initially configured the Onboard Key Manager. If you do not have this passphrase, contact NetApp Support.

- b. Verify all authentication keys are restored:

```
security key-manager key query
```

Confirm that the Restored column displays `true` for all authentication keys and the Key Manager type shows `onboard`.

- c. Back up the OKM information:

- i. Switch to advanced privilege mode:

```
set -priv advanced
```

Enter `y` when prompted to continue.

- ii. Display the key management backup information:

```
security key-manager onboard show-backup
```

- iii. Copy the backup information to a separate file or your log file.

You will need this backup information if you need to manually recover OKM during the replacement procedure.

- iv. Return to admin mode:

```
set -priv admin
```

- d. You can safely shut down the impaired controller and proceed to the shutdown procedure.

Shut down the controller for manual boot media recovery - FAS8300 and FAS8700

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

If your storage system is running ONTAP 9.17.1 or later, use the [automated boot recovery procedure](#). If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

- a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays...	Then...
The LOADER prompt	Go to Remove controller module.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: storage failover takeover -ofnode impaired_node_name When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <i>y</i> .

- b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.



This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.



Do not use this procedure if your system is in a two-node MetroCluster configuration.

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

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

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

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

The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:>`
`system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt (enter system password)	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode <i>impaired_node_name</i></pre> <p>When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <i>y</i>.</p>

Option 3: Controller is in a two-node Metrocluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

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

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller...	Then...
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: <code>metrocluster switchover</code>
Has not automatically switched over, you attempted switchover with the <code>metrocluster switchover</code> command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the `metrocluster operation show` command.

```
controller_A_1::> metrocluster operation show
  Operation: heal-aggregates
    State: successful
  Start Time: 7/25/2016 18:45:55
  End Time: 7/25/2016 18:45:56
    Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate      Size Available Used% State      #Vols  Nodes      RAID
Status
----- -----
...
aggr_b2      227.1GB    227.1GB    0% online      0  mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
  State: successful
  Start Time: 7/29/2016 20:54:41
  End Time: 7/29/2016 20:54:42
  Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the boot media and prepare for manual boot recovery - FAS8300 and FAS8700

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

If your storage system is running ONTAP 9.17.1 or later, use the [automated boot recovery procedure](#). If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

[Animation - Remove the controller module](#)

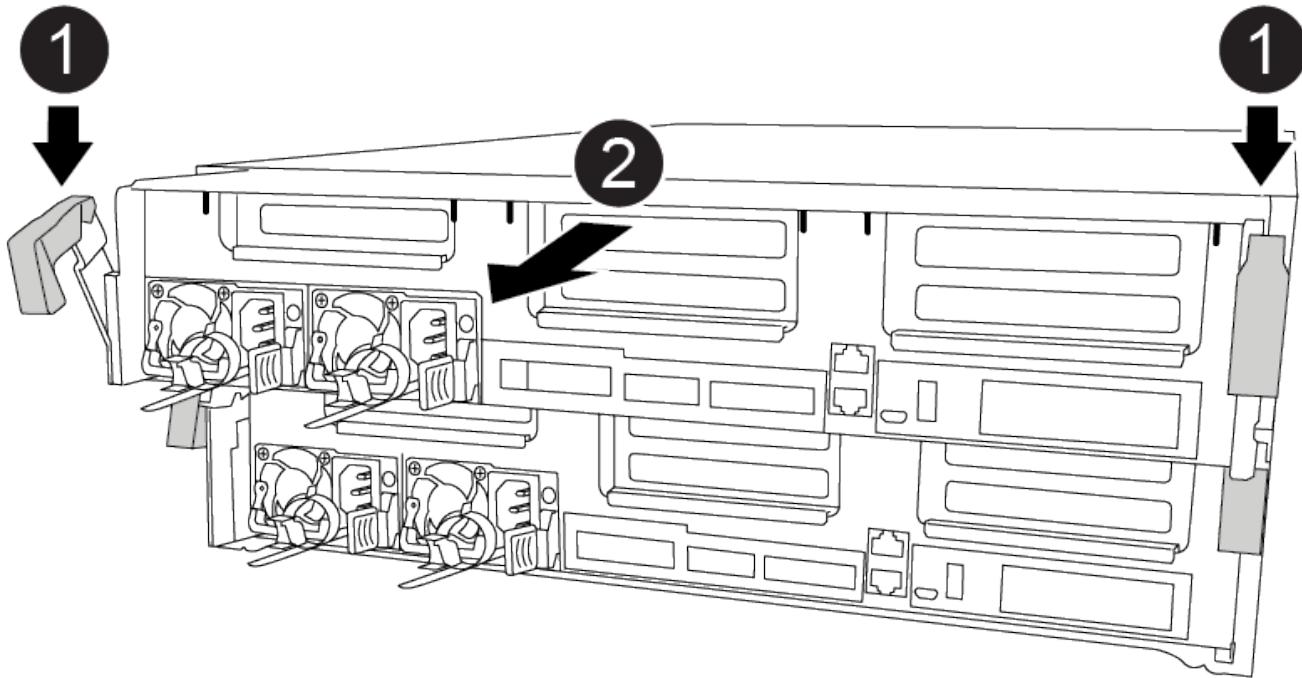
Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. 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, keeping 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.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.



1	Locking latches
2	Slide controller out of chassis

6. Slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 2: Replace the boot media

You must locate the boot media in the controller module (see the FRU map on the controller module), and then follow the directions to replace it.

Before you begin

Although the contents of the boot media is encrypted, it is a best practice to erase the contents of the boot media before replacing it. For more information, see the [Statement of Volatility](#) for your system on the NetApp Support Site.



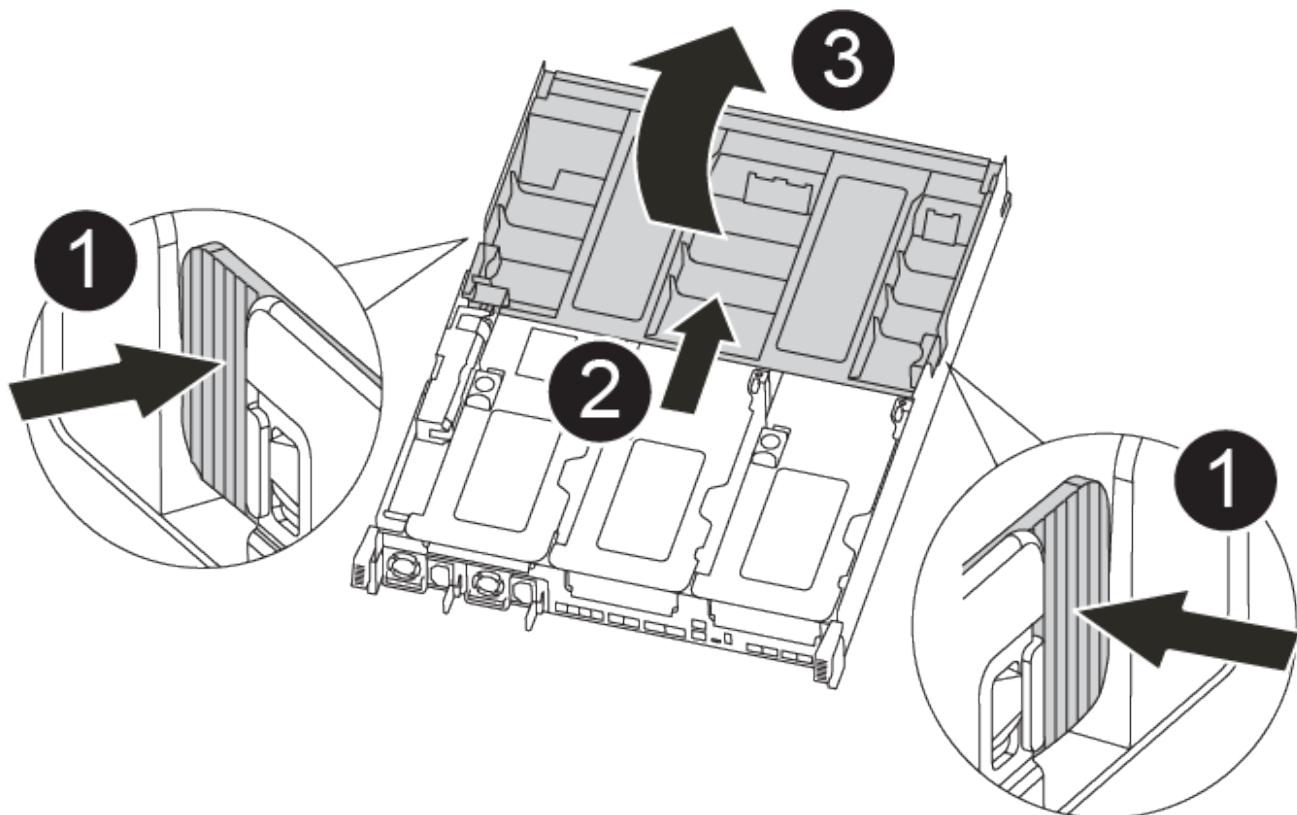
You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

You can use the following animation, illustrations, or the written steps to replace the boot media.

[Animation - Replace the boot media](#)

Steps

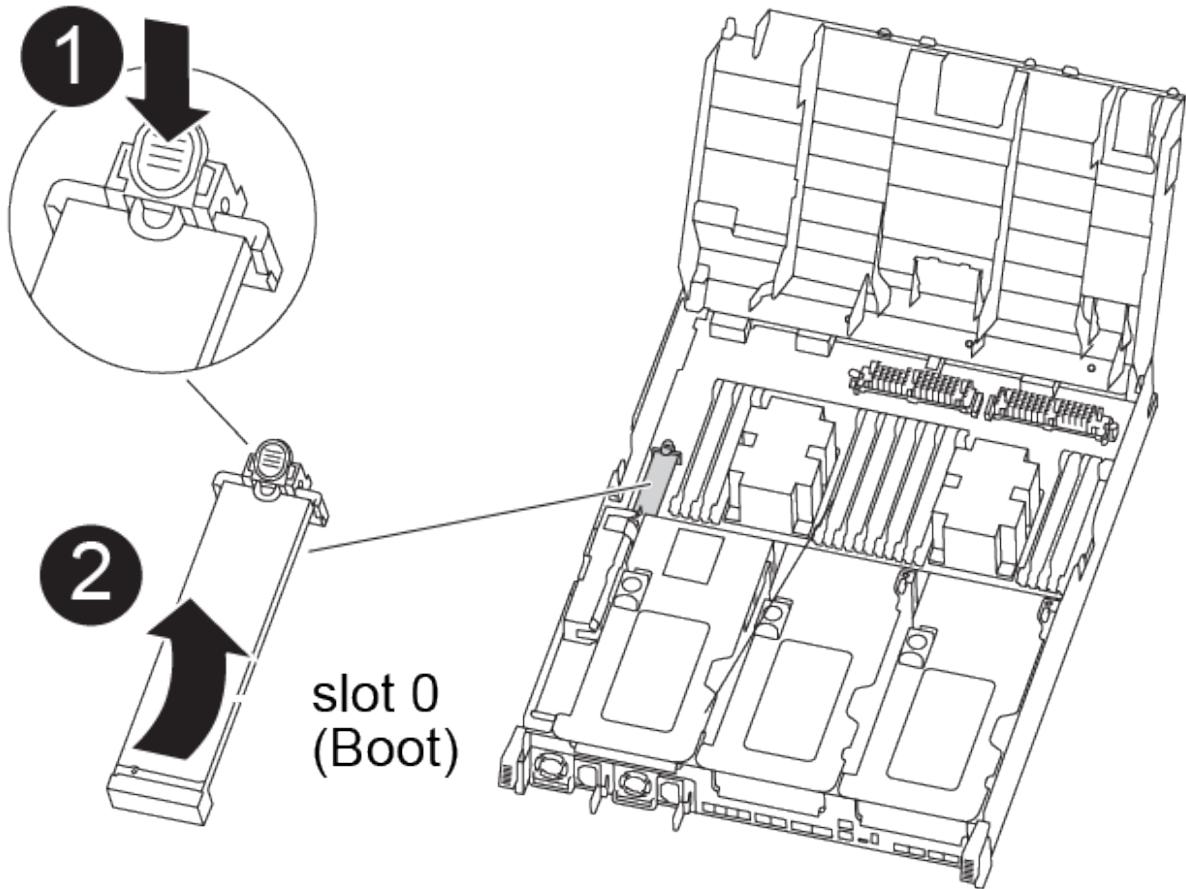
1. Open the air duct:



1	Locking tabs
2	Slide air duct toward back of controller
3	Rotate air duct up

- a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
- b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

2. Locate and remove the boot media from the controller module:



1	Press blue button
2	Rotate boot media up and remove from socket

- a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
- b. Rotate the boot media up and gently pull the boot media out of the socket.
3. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

If necessary, remove the boot media and reseat it into the socket.

5. Lock the boot media in place:
 - a. Rotate the boot media down toward the motherboard.
 - b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
 - c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
6. Close the air duct.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed does not have a boot image, so you need to transfer a boot image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
 - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
 - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the `var` file system.

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
 - a. Download the service image to your work space on your laptop.
 - b. Unzip the service image.



If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

There are two folders in the unzipped service image file:

- boot
- efi

- c. Copy the `efi` folder to the top directory on the USB flash drive.



If the service image has no `efi` folder, see [EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models](#)[^].

The USB flash drive should have the `efi` folder and the same Service Image (BIOS) version of what the impaired controller is running.

- d. Remove the USB flash drive from your laptop.
2. If you have not already done so, close the air duct.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
4. Reinstall the cable management device and recable the system, as needed.

When recabling, remember to reinstall the media converters (SFPs or QSFPs) if they were removed.

5. Insert the USB flash drive into the USB slot on the controller module.

Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB

console port.

6. Complete the installation of the controller module:

- Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latches rise when the controller module is fully seated.



Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

- Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
- Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

- If you have not already done so, reinstall the cable management device.

7. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then `halt` the controller to boot to LOADER.

8. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:

- Boot to Maintenance mode: `boot_ontap maint`
- Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t initiator adapter_name`
- Halt to return to Maintenance mode: `halt`

The changes will be implemented when the system is booted.

Manual boot media recovery from a USB drive - FAS8300 and FAS8700

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the [automated boot recovery procedure](#). If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Before you begin

- Ensure your console is connected to the impaired controller.
- Verify you have a USB flash drive with the recovery image.
- Determine if your system uses encryption. You will need to select the appropriate option in step 3 based on whether encryption is enabled.

Steps

- From the LOADER prompt on the impaired controller, boot the recovery image from the USB flash drive:

boot_recovery

The recovery image is downloaded from the USB flash drive.

2. When prompted, enter the name of the image or press **Enter** to accept the default image displayed in brackets.
3. Restore the var file system using the procedure for your ONTAP version:

ONTAP 9.16.0 or earlier

Complete the following steps on the impaired controller and partner controller:

- a. **On the impaired controller:** Press Y when you see Do you want to restore the backup configuration now?
- b. **On the impaired controller:** If prompted, press Y to overwrite `/etc/ssh/ssh_host_ecdsa_key`.
- c. **On the partner controller:** Set the impaired controller to advanced privilege level:

```
set -privilege advanced
```

- d. **On the partner controller:** Run the restore backup command:

```
system node restore-backup -node local -target-address  
impaired_node_IP_address
```



If you see any message other than a successful restore, contact NetApp Support.

- e. **On the partner controller:** Return to admin level:

```
set -privilege admin
```

- f. **On the impaired controller:** Press Y when you see Was the restore backup procedure successful?
- g. **On the impaired controller:** Press Y when you see ...would you like to use this restored copy now?
- h. **On the impaired controller:** Press Y when prompted to reboot, then press `Ctrl-C` when you see the Boot Menu.
- i. **On the impaired controller:** Do one of the following:
 - If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.
 - If the system uses encryption, go to [Restore encryption](#).

ONTAP 9.16.1 or later

Complete the following steps on the impaired controller:

- a. Press Y when prompted to restore the backup configuration.

After the restore procedure is successful, this message displays: `syncflash_partner:`
`Restore from partner complete`

- b. Press Y when prompted to confirm that the restore backup was successful.
- c. Press Y when prompted to use the restored configuration.
- d. Press Y when prompted to reboot the node.
- e. Press Y when prompted to reboot again, then press `Ctrl-C` when you see the Boot Menu.
- f. Do one of the following:
 - If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.

- If the system uses encryption, go to [Restore encryption](#).

4. Connect the console cable to the partner controller.
5. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -fromnode local
```

6. If you disabled automatic giveback, reenable it:

```
storage failover modify -node local -auto-giveback true
```

7. If AutoSupport is enabled, restore automatic case creation:

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

Restore encryption - FAS8300 and FAS8700

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the [automated boot recovery procedure](#). If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Complete the appropriate steps to restore encryption on your system based on your key manager type. If you are unsure which key manager your system uses, check the settings you captured at the beginning of the boot media replacement procedure.

Onboard Key Manager (OKM)

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

Ensure you have the following information available:

- Cluster-wide passphrase entered while [enabling onboard key management](#)
- [Backup information for the Onboard Key Manager](#)
- Verification that you have the correct passphrase and backup data using the [How to verify onboard key management backup and cluster-wide passphrase](#) procedure

Steps

On the impaired controller:

1. Connect the console cable to the impaired controller.
2. From the ONTAP boot menu, select the appropriate option:

ONTAP version	Select this option
ONTAP 9.8 or later	<p>Select option 10.</p> <p>Show example boot menu</p> <div style="border: 1px solid #ccc; padding: 10px; width: fit-content; margin: 10px auto; border-radius: 10px;"><p>Please choose one of the following:</p><p>(1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. (10) Set Onboard Key Manager recovery secrets. (11) Configure node for external key management.</p><p>Selection (1-11)? 10</p></div>

ONTAP version	Select this option
ONTAP 9.7 and earlier	<p>Select the hidden option <code>recover_onboard_keymanager</code></p> <p>Show example boot menu</p> <div style="border: 1px solid #ccc; padding: 10px; border-radius: 10px; background-color: #f9f9f9;"> <p style="margin: 0;">Please choose one of the following:</p> <ul style="list-style-type: none"> (1) Normal Boot. (2) Boot without /etc/rc. (3) Change password. (4) Clean configuration and initialize all disks. (5) Maintenance mode boot. (6) Update flash from backup config. (7) Install new software first. (8) Reboot node. (9) Configure Advanced Drive Partitioning. <p style="margin: 0;">Selection (1-19)?</p> <p style="margin: 0;"><code>recover_onboard_keymanager</code></p> </div>

3. Confirm that you want to continue the recovery process when prompted:

Show example prompt

This option must be used only in disaster recovery procedures. Are you sure? (y or n) :

4. Enter the cluster-wide passphrase twice.

While entering the passphrase, the console does not show any input.

Show example prompt

Enter the passphrase for onboard key management:

Enter the passphrase again to confirm:

5. Enter the backup information:

- a. Paste the entire content from the BEGIN BACKUP line through the END BACKUP line, including the dashes.

Show example prompt

Enter the backup data:

```
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA  
AA  
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA  
AA  
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA  
AA  
0123456789012345678901234567890123456789012345678901  
23  
1234567890123456789012345678901234567890123456789012  
34  
2345678901234567890123456789012345678901234567890123  
45  
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA  
AA  
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAA  
AA  
AAAAAAAAAAAAAAAAAAAAAA  
AA  
-----END  
BACKUP-----
```

- b. Press Enter twice at the end of the input.

The recovery process completes and displays the following message:

Successfully recovered keymanager secrets.

Show example prompt

```
Trying to recover keymanager secrets....  
Setting recovery material for the onboard key manager  
Recovery secrets set successfully  
Trying to delete any existing km_onboard.wkeydb file.  
  
Successfully recovered keymanager secrets.  
  
*****  
*****  
* Select option "(1) Normal Boot." to complete recovery  
process.  
*  
* Run the "security key-manager onboard sync" command to  
synchronize the key database after the node reboots.  
*****  
*****
```



Do not proceed if the displayed output is anything other than Successfully recovered keymanager secrets. Perform troubleshooting to correct the error.

6. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****
*****
* Select option "(1) Normal Boot." to complete the recovery
process.
*
*****
*****
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message:

Waiting for giveback...(Press Ctrl-C to abort wait)

On the partner controller:

8. Giveback the impaired controller:

```
storage failover giveback -fromnode local -only-cfo-aggregates true
```

On the impaired controller:

9. After booting with only the CFO aggregate, synchronize the key manager:

```
security key-manager onboard sync
```

10. Enter the cluster-wide passphrase for the Onboard Key Manager when prompted.

Show example prompt

Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the corresponding volume encryption keys (VEKs) will be restored automatically within 10 minutes. If any offline encrypted volumes are not brought online automatically, they can be brought online manually using the "volume online -vserver <vserver> -volume <volume_name>" command.



If the sync is successful, the cluster prompt is returned with no additional messages. If the sync fails, an error message appears before returning to the cluster prompt. Do not continue until the error is corrected and the sync runs successfully.

11. Verify that all keys are synced:

```
security key-manager key query -restored false
```

The command should return no results. If any results appear, repeat the sync command until no results are returned.

On the partner controller:

12. Giveback the impaired controller:

```
storage failover giveback -fromnode local
```

13. Restore automatic giveback if you disabled it:

```
storage failover modify -node local -auto-giveback true
```

14. If AutoSupport is enabled, restore automatic case creation:

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

External Key Manager (EKM)

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

Gather the following files from another cluster node or from your backup:

- /cfcard/kmip/servers.cfg file or the KMIP server address and port
- /cfcard/kmip/certs/client.crt file (client certificate)
- /cfcard/kmip/certs/client.key file (client key)
- /cfcard/kmip/certs/CA.pem file (KMIP server CA certificates)

Steps

On the impaired controller:

1. Connect the console cable to the impaired controller.
2. Select option 11 from the ONTAP boot menu.

Show example boot menu

```
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 11
```

3. Confirm you have gathered the required information when prompted:

Show example prompt

```
Do you have a copy of the /cfcard/kmip/certs/client.crt file?
{y/n}
Do you have a copy of the /cfcard/kmip/certs/client.key file?
{y/n}
Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
```

4. Enter the client and server information when prompted:

- a. Enter the client certificate (client.crt) file contents, including the BEGIN and END lines.
- b. Enter the client key (client.key) file contents, including the BEGIN and END lines.
- c. Enter the KMIP server CA(s) (CA.pem) file contents, including the BEGIN and END lines.
- d. Enter the KMIP server IP address.
- e. Enter the KMIP server port (press Enter to use the default port 5696).

Show example

```
Enter the client certificate (client.crt) file contents:  
-----BEGIN CERTIFICATE-----  
<certificate_value>  
-----END CERTIFICATE-----
```

```
Enter the client key (client.key) file contents:  
-----BEGIN RSA PRIVATE KEY-----  
<key_value>  
-----END RSA PRIVATE KEY-----
```

```
Enter the KMIP server CA(s) (CA.pem) file contents:  
-----BEGIN CERTIFICATE-----  
<certificate_value>  
-----END CERTIFICATE-----
```

```
Enter the IP address for the KMIP server: 10.10.10.10  
Enter the port for the KMIP server [5696]:
```

```
System is ready to utilize external key manager(s).  
Trying to recover keys from key servers....  
kmip_init: configuring ports  
Running command '/sbin/ifconfig e0M'  
..  
..  
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

The recovery process completes and displays the following message:

```
Successfully recovered keymanager secrets.
```

Show example

```
System is ready to utilize external key manager(s).  
Trying to recover keys from key servers....  
Performing initialization of OpenSSL  
Successfully recovered keymanager secrets.
```

5. Select option 1 from the boot menu to continue booting into ONTAP.

Show example prompt

```
*****
*****
* Select option "(1) Normal Boot." to complete the recovery
process.
*
*****
*****
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it:

```
storage failover modify -node local -auto-giveback true
```

7. If AutoSupport is enabled, restore automatic case creation:

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

Return the failed boot media to NetApp - FAS8300 and FAS8700

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Replace the caching module - FAS8300 and FAS8700

You must replace the caching module in the controller module when your system registers a single AutoSupport (ASUP) message that the module has gone offline; failure to do so results in performance degradation.



The Ver2 controller module has only one caching module socket in the FAS8300. FAS8700 does not have a VER2 controller module. The caching module functionality is not impacted by the socket removal.

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

Step 1: Shut down the impaired controller

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

Option 1: Most configurations

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

About this tasks

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller.

Synchronize a node with the cluster

You might want to erase the contents of your caching module before replacing it.

Steps

1. Although data on the caching module is encrypted, you might want to erase any data from the impaired caching module and verify that the caching module has no data:
 - a. Erase the data on the caching module: `system controller flash-cache secure-erase run -node node_name localhost -device-id device_number`

 Run the `system controller flash-cache show` command if you don't know the Flash Cache device ID.

- b. Verify that the data has been erased from the caching module: `system controller flash-cache secure-erase show`

2. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:
`system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:*> system node autosupport invoke -node * -type all -message  
MAINT=2h
```

3. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
4. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <code>y</code> .
System prompt or password prompt (enter system password)	Take over or halt the impaired controller: <code>storage failover takeover -ofnode impaired_node_name</code> When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <code>y</code> .

Option 2: Controller is in a two-node MetroCluster

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

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller...	Then...
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: <code>metrocluster switchover</code>
Has not automatically switched over, you attempted switchover with the <code>metrocluster switchover</code> command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoed` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the `metrocluster operation show` command.

```
controller_A_1::> metrocluster operation show
Operation: heal-aggregates
State: successful
Start Time: 7/25/2016 18:45:55
End Time: 7/25/2016 18:45:56
Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate      Size Available Used% State    #Vols  Nodes
RAID Status
-----
-----
...
aggr_b2      227.1GB   227.1GB   0% online      0  mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-veto`s parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
  State: successful
  Start Time: 7/29/2016 20:54:41
  End Time: 7/29/2016 20:54:42
  Errors: -
```

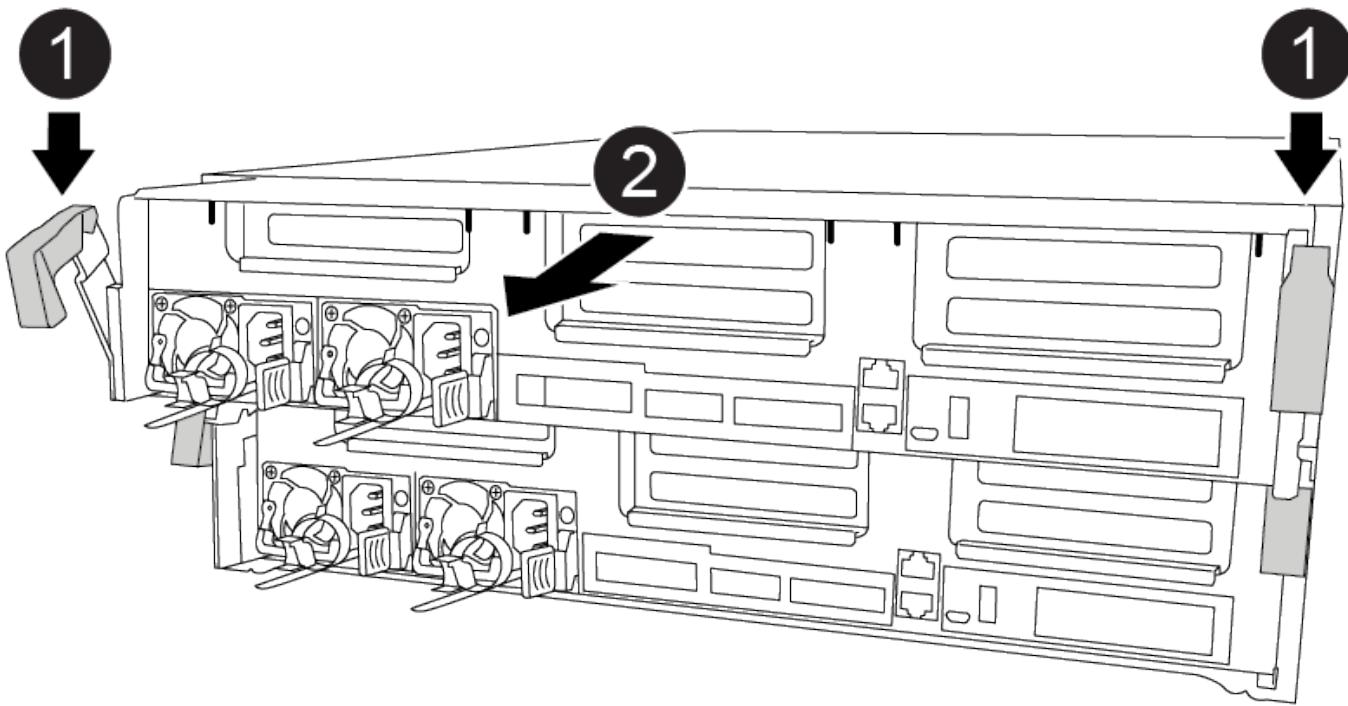
8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

[Animation - Remove the controller module](#)



Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. 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, keeping 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.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace a caching module

To replace a caching module, referred to as the Flash Cache on the label on your controller, locate the slot inside the controller and follow the specific sequence of steps. See the FRU map on the controller module for the location of the Flash Cache.



Slot 6 is only available in FAS8300 VER2 Controller.

Your storage system must meet certain criteria depending on your situation:

- It must have the appropriate operating system for the caching module you are installing.
- It must support the caching capacity.
- Although the contents of the caching module is encrypted, it is a best practice to erase the contents of the module before replacing it. For more information, see the [Statement of Volatility](#) for your system on the NetApp Support Site.

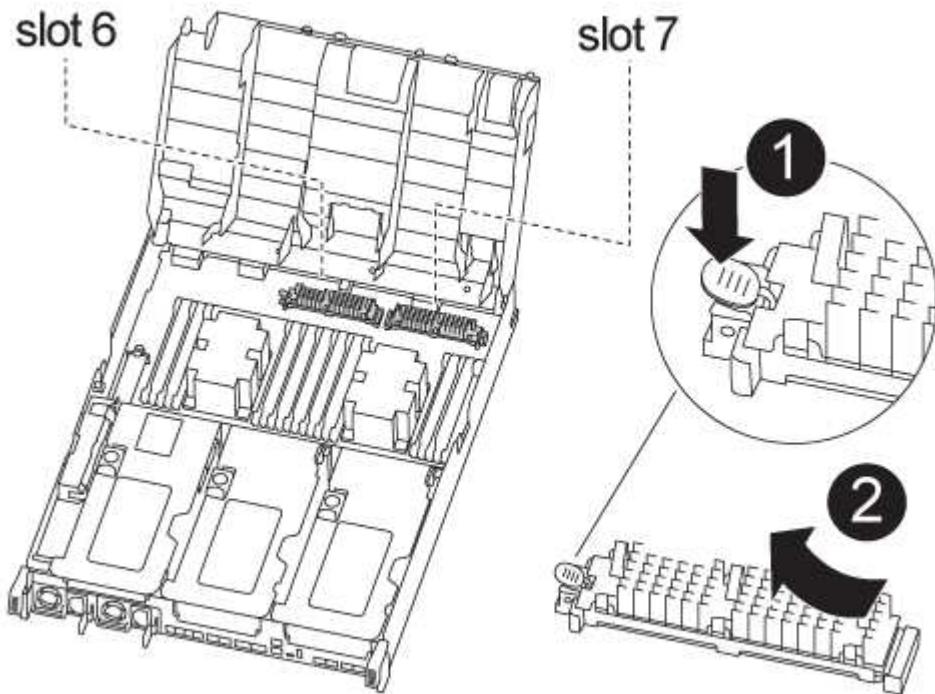


You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

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

You can use the following animation, illustration, or the written steps to replace a caching module.

[Animation - Replace the caching module](#)



Steps

1. If you are not already grounded, properly ground yourself.
2. Open the air duct:
 - a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
 - b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
3. Using the FRU map on the controller module, locate the failed caching module and remove it:

Depending on your configuration, there may be zero, one, or two caching modules in the controller module. Use the FRU map inside the controller module to help locate the caching module.

- a. Press the blue release tab.

The caching module end rises clear of the release tab.

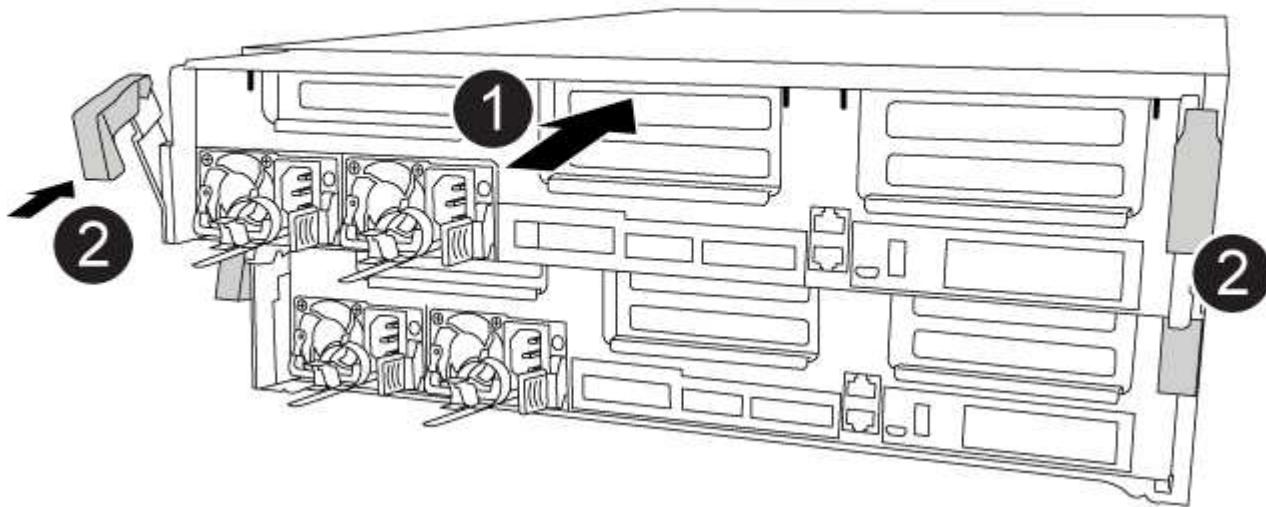
- b. Rotate the caching module up and slide it out of the socket.
4. Install the replacement caching module:
 - a. Align the edges of the replacement caching module with the socket and gently insert it into the socket.
 - b. Rotate the caching module downward toward the motherboard.
 - c. Placing your finger at the end of the caching module by the blue button, firmly push down on the caching module end, and then lift the locking button to lock the caching module in place.
5. Close the air duct:
 - a. Rotate the air duct down to the controller module.
 - b. Slide the air duct toward the risers to lock it in place.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

[Animation - Install the controller module](#)



Steps

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

 Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.



You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:

- a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

- Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

- b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

- c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

- d. If you have not already done so, reinstall the cable management device.

- e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.



If your system stops at the boot menu, select the option to boot to LOADER.

- f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`

3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 7: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the enabled state: `metrocluster node show`

```
cluster_B::> metrocluster node show

DR          Configuration DR
Group Cluster Node  State   Mirroring Mode
-----  -----  -----  -----
-----  -----
1   cluster_A
      controller_A_1 configured   enabled   heal roots
completed
      cluster_B
      controller_B_1 configured   enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster          Configuration State   Mode
-----  -----  -----  -----
Local: cluster_B configured      switchover
Remote: cluster_A configured    waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster          Configuration State   Mode
-----  -----  -----  -----
Local: cluster_B configured      normal
Remote: cluster_A configured    normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 8: Complete the replacement process

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Chassis

Overview of chassis replacement - FAS8300 and FAS8700

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

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is disruptive. For a two-controller cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - FAS8300 and FAS8700

Option 1: Most configurations

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see [Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base](#).

Before you begin

- Make sure you have the necessary permissions and credentials:
 - Local administrator credentials for ONTAP.
 - BMC accessibility for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
 - Perform additional [system health checks](#).
 - Upgrade ONTAP to a recommended release for the system.
 - Resolve any [Active IQ Wellness Alerts and Risks](#). Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

```
system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
```

5. Identify the SP/BMC address of all cluster nodes:

```
system service-processor show -node * -fields address
```

6. Exit the cluster shell:

```
exit
```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true
```



For clusters using SnapMirror synchronous operating in StrictSync mode: system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict -sync-warnings true

9. Enter **y** for each controller in the cluster when you see:

Warning: Are you sure you want to halt node <node_name>? {y|n}:

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster configuration

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

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller...	Then...
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: <code>metrocluster switchover</code>

If the impaired controller...	Then...
Has not automatically switched over, you attempted switchover with the metrocluster switchover command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the metrocluster heal -phase aggregates command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the metrocluster heal command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
  Operation: heal-aggregates
    State: successful
  Start Time: 7/25/2016 18:45:55
  End Time: 7/25/2016 18:45:56
    Errors: -
```

5. Check the state of the aggregates by using the storage aggregate show command.

```
controller_A_1::> storage aggregate show
Aggregate      Size Available Used% State      #Vols  Nodes      RAID
Status
-----
...
aggr_b2      227.1GB    227.1GB    0% online      0  mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the metrocluster heal -phase root-aggregates command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
    State: successful
  Start Time: 7/29/2016 20:54:41
  End Time: 7/29/2016 20:54:42
  Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Move and replace hardware - FAS8300 and FAS8700

Move the fans, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. 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, keeping 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.

4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move the fans

To move the fan modules to the replacement chassis when replacing the chassis, you must perform a specific sequence of tasks.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

The fan module moves a little bit away from the chassis.

4. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.



The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

5. Set the fan module aside.
6. Repeat the preceding steps for any remaining fan modules.
7. Insert the fan module into the replacement chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.
10. Repeat these steps for the remaining fan modules.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

Steps

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install 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.



Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Complete the installation of the controller module:
 - a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.



Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

- b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
- c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

- d. If you have not already done so, reinstall the cable management device.
- e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.



If your system stops at the boot menu, select the option to boot to LOADER.

- f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.
- g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

If your system stops at the boot menu, select the option to boot to LOADER.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - FAS8300 and FAS8700

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

Steps

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

- a. Set the HA state for the chassis: `ha-config modify chassis HA-state`

The value for HA-state can be one of the following:

- ha
- mcc
- mcc-2n
- mccip
- non-ha

- b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the enabled state: `metrocluster node show`

```
cluster_B::> metrocluster node show

DR          Configuration  DR
Group Cluster Node  State      Mirroring Mode
-----  -----
-----  -----
1      cluster_A
          controller_A_1 configured  enabled  heal roots
completed
      cluster_B
          controller_B_1 configured  enabled  waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`

3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster          Configuration State      Mode
-----
Local: cluster_B configured          switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster          Configuration State      Mode
-----
Local: cluster_B configured          normal
Remote: cluster_A configured        normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Complete the replacement process

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Controller

Overview of controller module replacement - FAS8300 and FAS8700

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in a MetroCluster configuration, you must review the section [Choosing the correct recovery procedure](#) to determine whether you should use this procedure.

If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be

used to provide nondisruptive operation during the replacement.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.
- It is important that you apply the commands in these steps on the correct systems:
 - The *impaired* controller is the controller that is being replaced.
 - The *replacement node* is the new controller that is replacing the impaired controller.
 - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text file.

This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - FAS8300 and FAS8700

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

Option 1: Most systems

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

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, **quorum status** of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

Steps

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

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter *y* when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a two-node MetroCluster

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

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller...	Then...
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: <code>metrocluster switchover</code>
Has not automatically switched over, you attempted switchover with the <code>metrocluster switchover</code> command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
  Operation: heal-aggregates
    State: successful
  Start Time: 7/25/2016 18:45:55
  End Time: 7/25/2016 18:45:56
  Errors: -
```

5. Check the state of the aggregates by using the storage aggregate show command.

```
controller_A_1::> storage aggregate show
  Aggregate      Size Available Used% State  #Vols  Nodes
  RAID  Status
  -----  -----
  ...
  aggr_b2      227.1GB    227.1GB    0% online      0 mcc1-a2
  raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the metrocluster heal -phase root-aggregates command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the metrocluster heal command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the metrocluster operation show command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
    State: successful
  Start Time: 7/29/2016 20:54:41
  End Time: 7/29/2016 20:54:42
  Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the controller module hardware - FAS8300 and FAS8700

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.



The Ver2 controller module has only one caching module socket (Slot 6) in the FAS8300. FAS8700 does not have a VER2 controller module. The caching module functionality is not impacted by the socket removal.

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

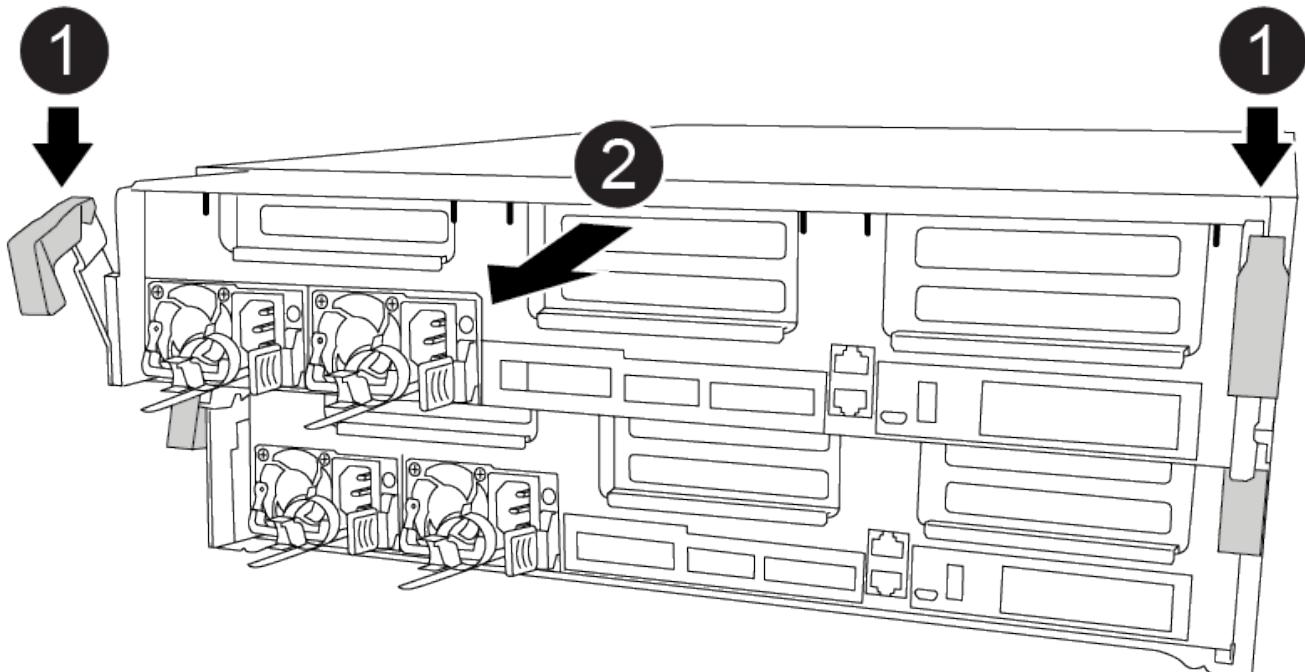
You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

[Animation - Remove the controller module](#)

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. 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, keeping 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.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.



The controller module moves slightly out of the chassis.

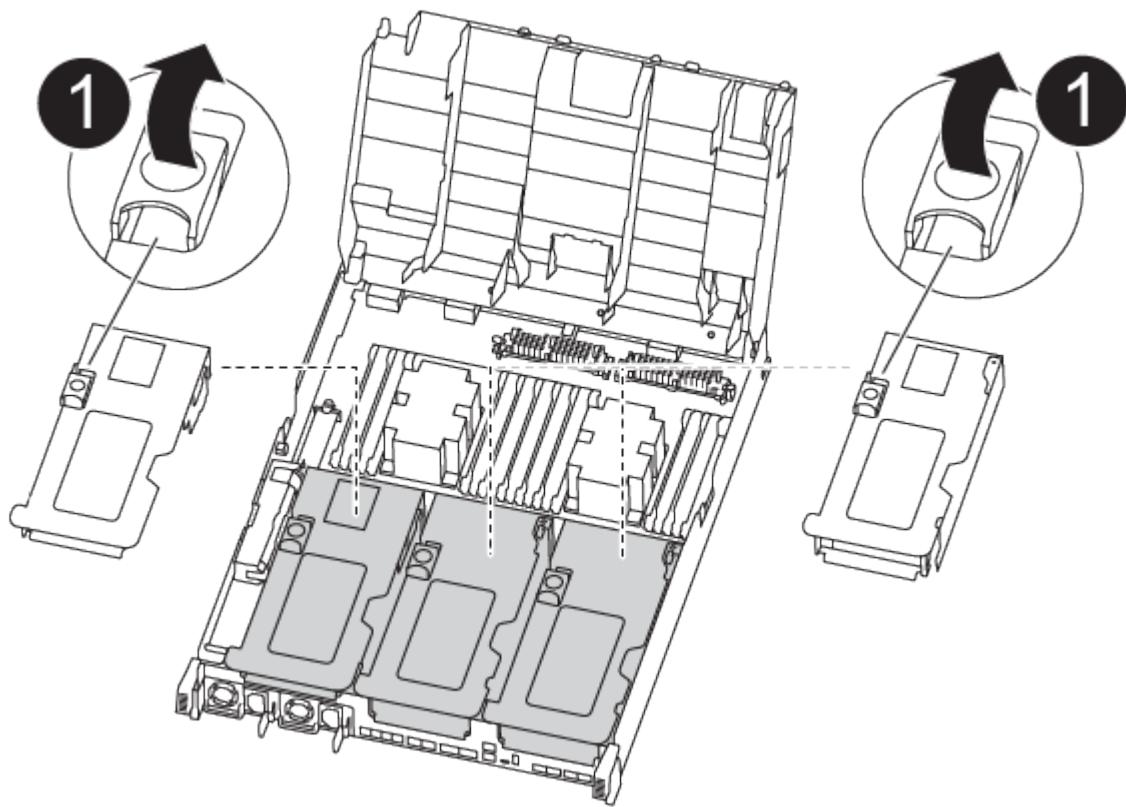
6. Slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

8. On the replacement controller module, open the air duct and remove the empty risers from the controller module using the animation, illustration, or the written steps:

[Removing the empty risers from the replacement controller module](#)



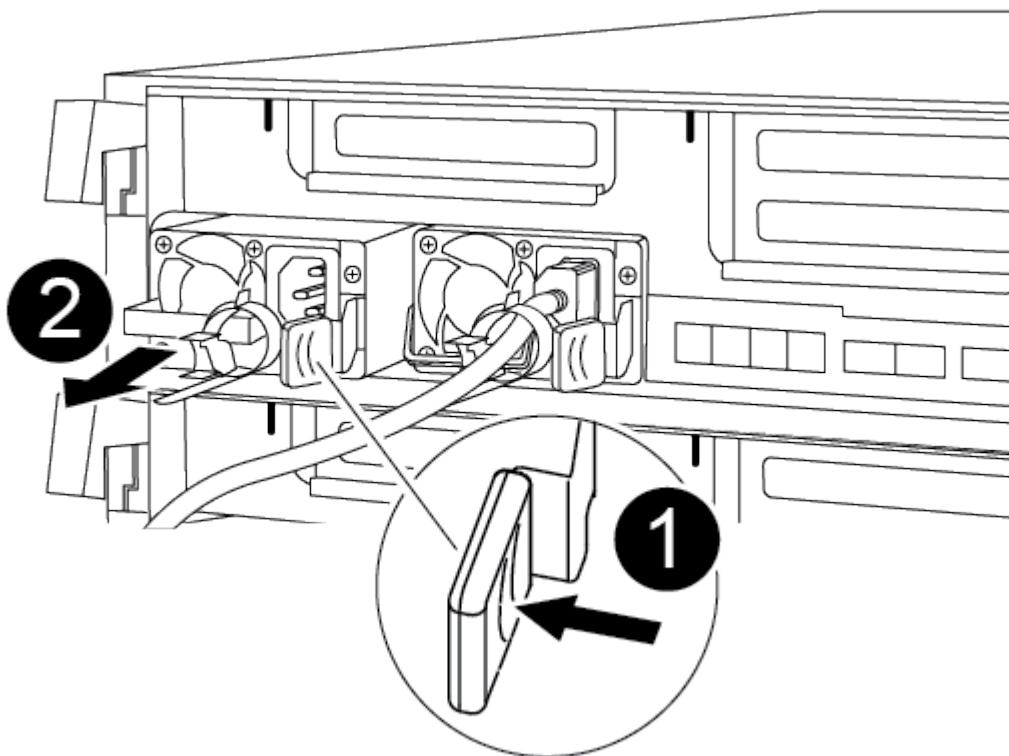
- a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
- b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
- c. Rotate the riser locking latch on the left side of riser 1 up and toward air duct, lift the riser up, and then set it aside.
- d. Repeat the previous step for the remaining risers.

Step 2: Move the power supplies

You must move the power supply from the impaired controller module to the replacement controller module when you replace a controller module.

You can use the following animation, illustration, or the written steps to move the power supplies to the replacement controller module.

Animation - Move the power supplies



1. Remove the power supply:
 - a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.
 - b. Press the blue locking tab to release the power supply from the chassis.
 - c. Using both hands, pull the power supply out of the chassis, and then set it aside.
2. Move the power supply to the new controller module, and then install it.
3. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

The power supplies will only properly engage with the internal connector and lock in place one way.



To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

4. Repeat the preceding steps for any remaining power supplies.

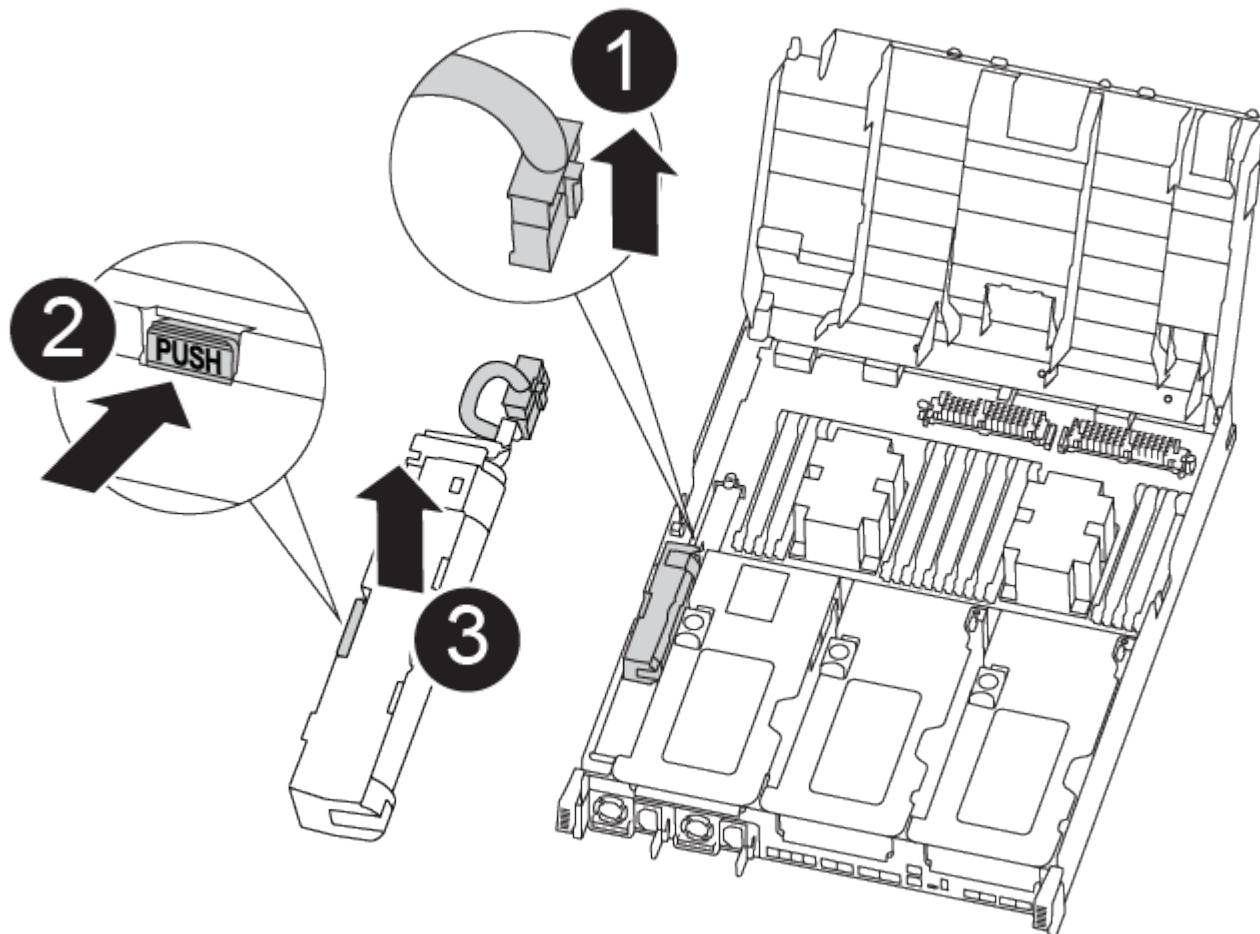
Step 3: Move the NVDIMM battery

To move the NVDIMM battery from the impaired controller module to the replacement controller module, you must perform a specific sequence of steps.

You can use the following animation, illustration, or the written steps to move the NVDIMM battery from the impaired controller module to the replacement controller module.

Animation - Move the NVDIMM battery

1. Open the air duct:
 - a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
 - b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the NVDIMM battery in the controller module.



1. 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.
2. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
3. Move the battery to the replacement controller module.
4. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.



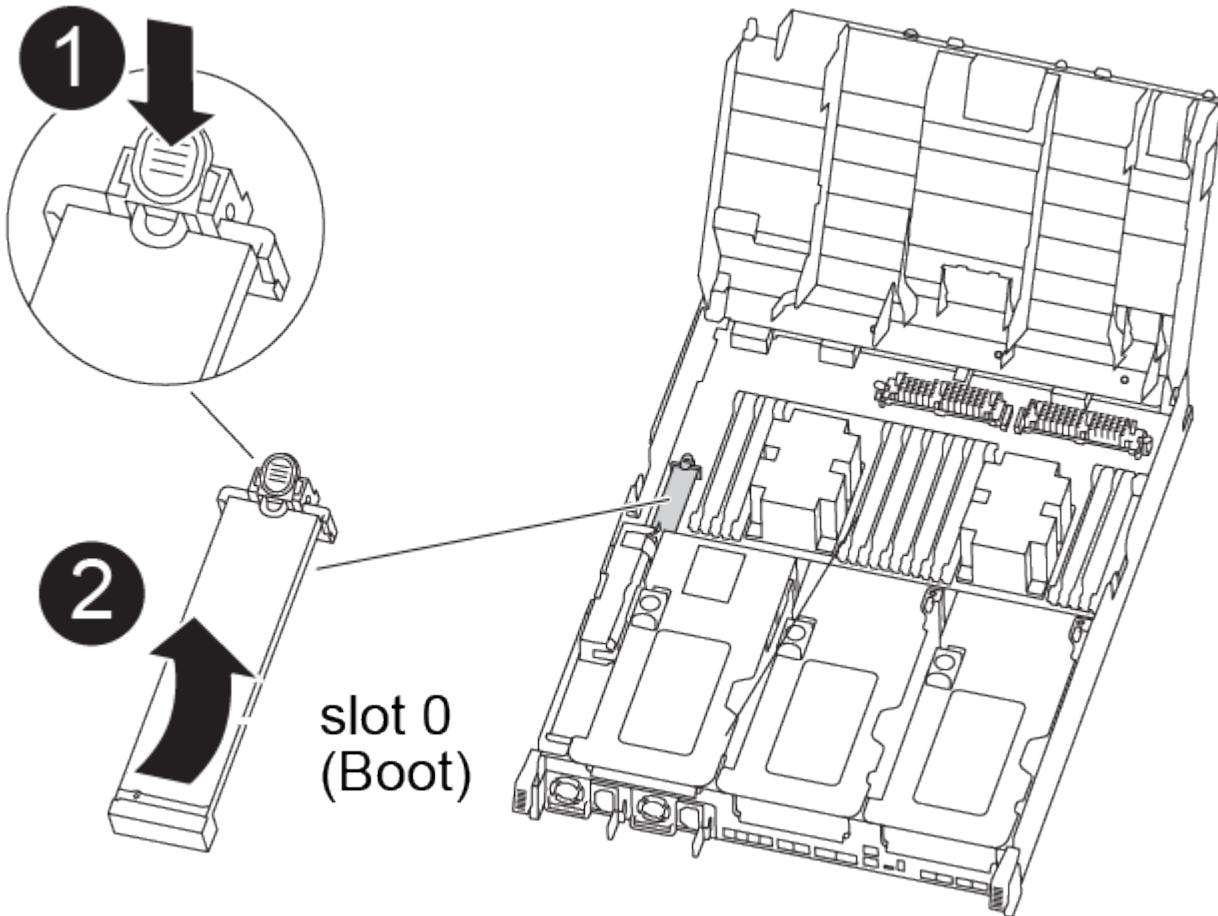
Do not plug the battery cable back into the motherboard until instructed to do so.

Step 4: Move the boot media

You must locate the boot media, and then follow the directions to remove it from the impaired controller module and insert it into the replacement controller module.

You can use the following animation, illustration, or the written steps to move the boot media from the impaired controller module to the replacement controller module.

[Animation - Move the boot media](#)



1. Locate and remove the boot media from the controller module:
 - a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
 - b. Rotate the boot media up and gently pull the boot media out of the socket.
2. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
3. Check the boot media to make sure that it is seated squarely and completely in the socket.
If necessary, remove the boot media and reseat it into the socket.
4. Lock the boot media in place:
 - a. Rotate the boot media down toward the motherboard.
 - b. Press the blue locking button so that it is in the open position.
 - c. Placing your fingers at the end of the boot media by the blue button, firmly push down on the boot media end to engage the blue locking button.

Step 5: Move the PCIe risers and mezzanine card

As part of the controller replacement process, you must move the PCIe risers and mezzanine card from the impaired controller module to the replacement controller module.

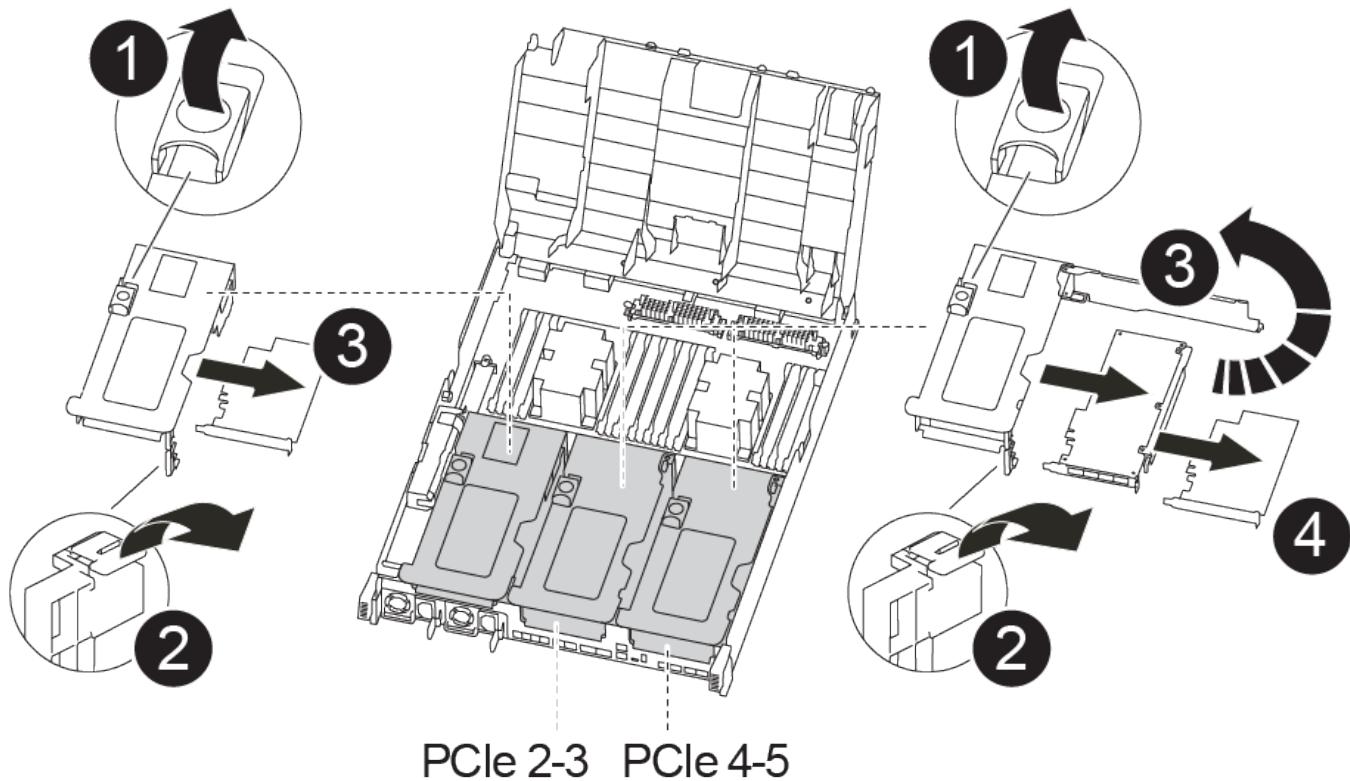
You can use the following animations, illustrations, the FUR map on the system, or the written steps to move the PCIe risers and mezzanine card from the impaired controller module to the replacement controller module.



You do not have to remove the PCIe cards from the risers. Transfer the risers, with the PCIe cards still installed, to the replacement controller module.

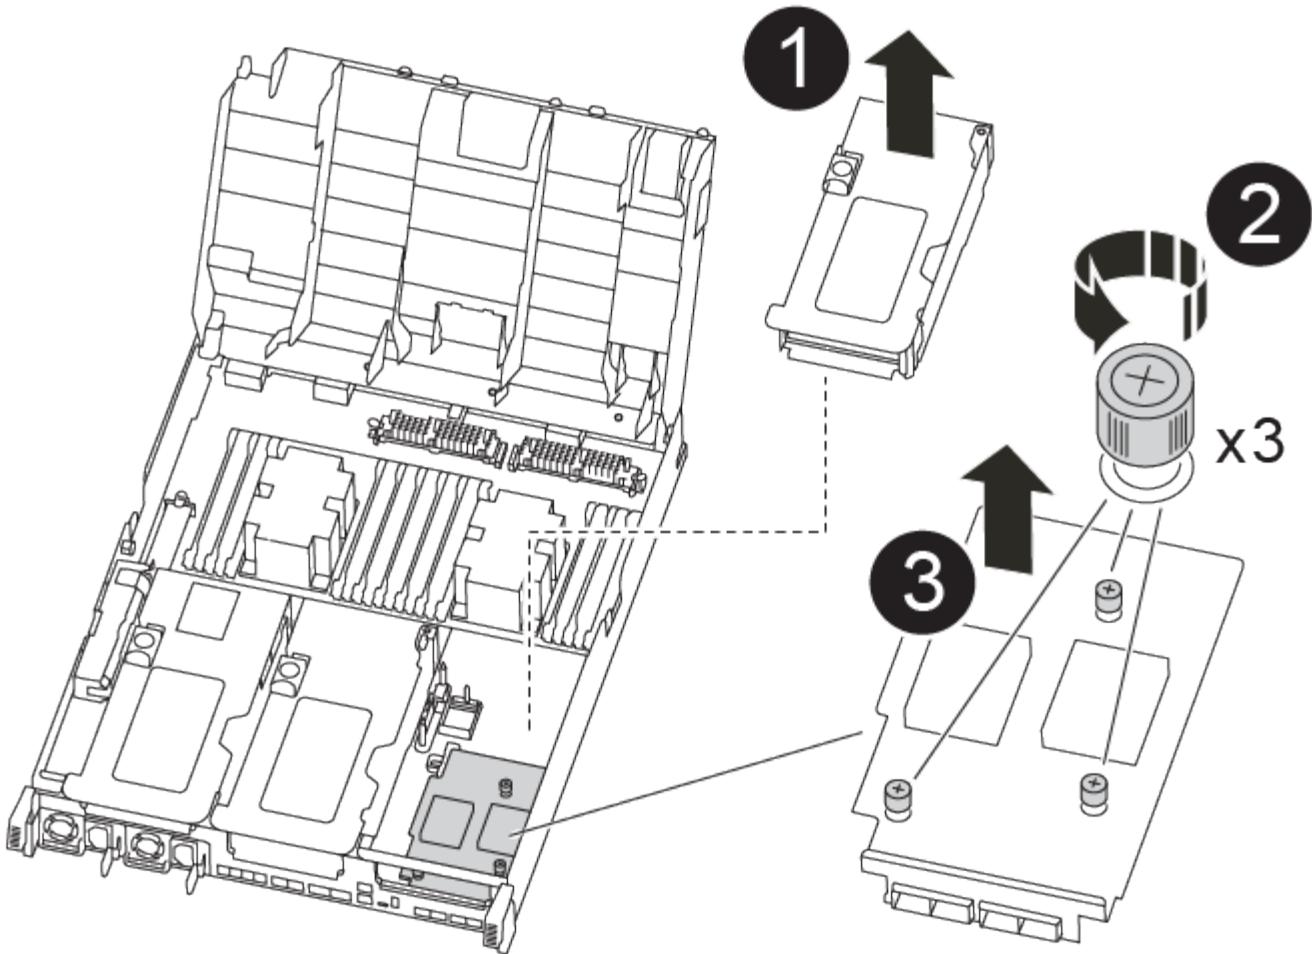
Moving PCIe riser 1 and 2 (left and middle risers):

[Animation - Move PCI risers 1 and 2](#)



Moving the mezzanine card and riser 3 (right riser):

[Animation - Move the mezzanine card and riser 3](#)



1. Move PCIe risers one and two from the impaired controller module to the replacement controller module:
 - a. Remove any SFP or QSFP modules that might be in the PCIe cards.
 - b. Rotate the riser locking latch on the left side of the riser up and toward air duct.

The riser raises up slightly from the controller module.

 - c. Lift the riser up, and then move it to the replacement controller module.
 - d. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins, push the riser squarely into the socket on the motherboard, and then rotate the latch down flush with the sheet metal on the riser.
 - e. Repeat this step for riser number 2.
2. Remove riser number 3, remove the mezzanine card, and install both into the replacement controller module:
 - a. Remove any SFP or QSFP modules that might be in the PCIe cards.
 - b. Rotate the riser locking latch on the left side of the riser up and toward air duct.

The riser raises up slightly from the controller module.

 - c. Lift the riser up, and then set it aside on a stable, flat surface.
 - d. Loosen the thumbscrews on the mezzanine card, and gently lift the card directly out of the socket, and then move it to the replacement controller module.

e. Install the mezzanine in the replacement controller and secure it with the thumbscrews.

f. Install the third riser in the replacement controller module.

Step 6: Move caching modules

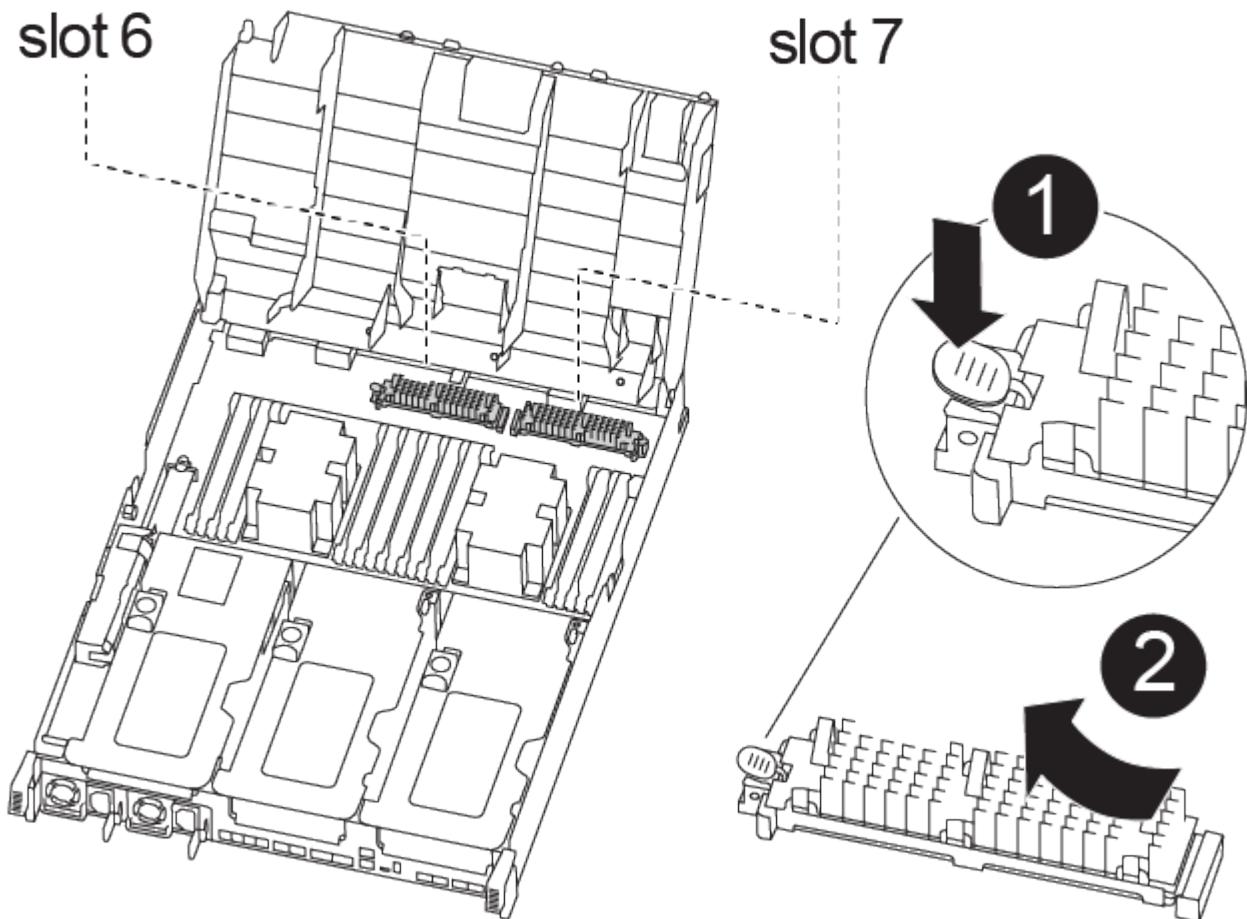
You must move the caching modules from the impaired controller modules to the replacement controller module when replacing a controller module.



The Ver2 controller module has only one caching module socket in the FAS8300. FAS8700 does not have a VER2 controller module. The caching module functionality is not impacted by the socket removal.

You can use the following animation, illustration, or the written steps to move caching modules to the new controller module.

Animation - Move the caching modules



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

2. Move the caching modules from the impaired controller module to the replacement controller module:

- Press the blue release tab at the end of the caching module, rotate the module up, and then remove the module from the socket.
- Move the caching module to the same socket on the replacement controller module.

- c. Align the edges of the caching module with the socket and gently insert the module as far into the socket as it will go.
- d. Rotate the caching module downward toward the motherboard.
- e. Placing your finger at the end of the caching module by the blue button, firmly push down on the caching module end, and then lift the locking button to lock the caching module in place.

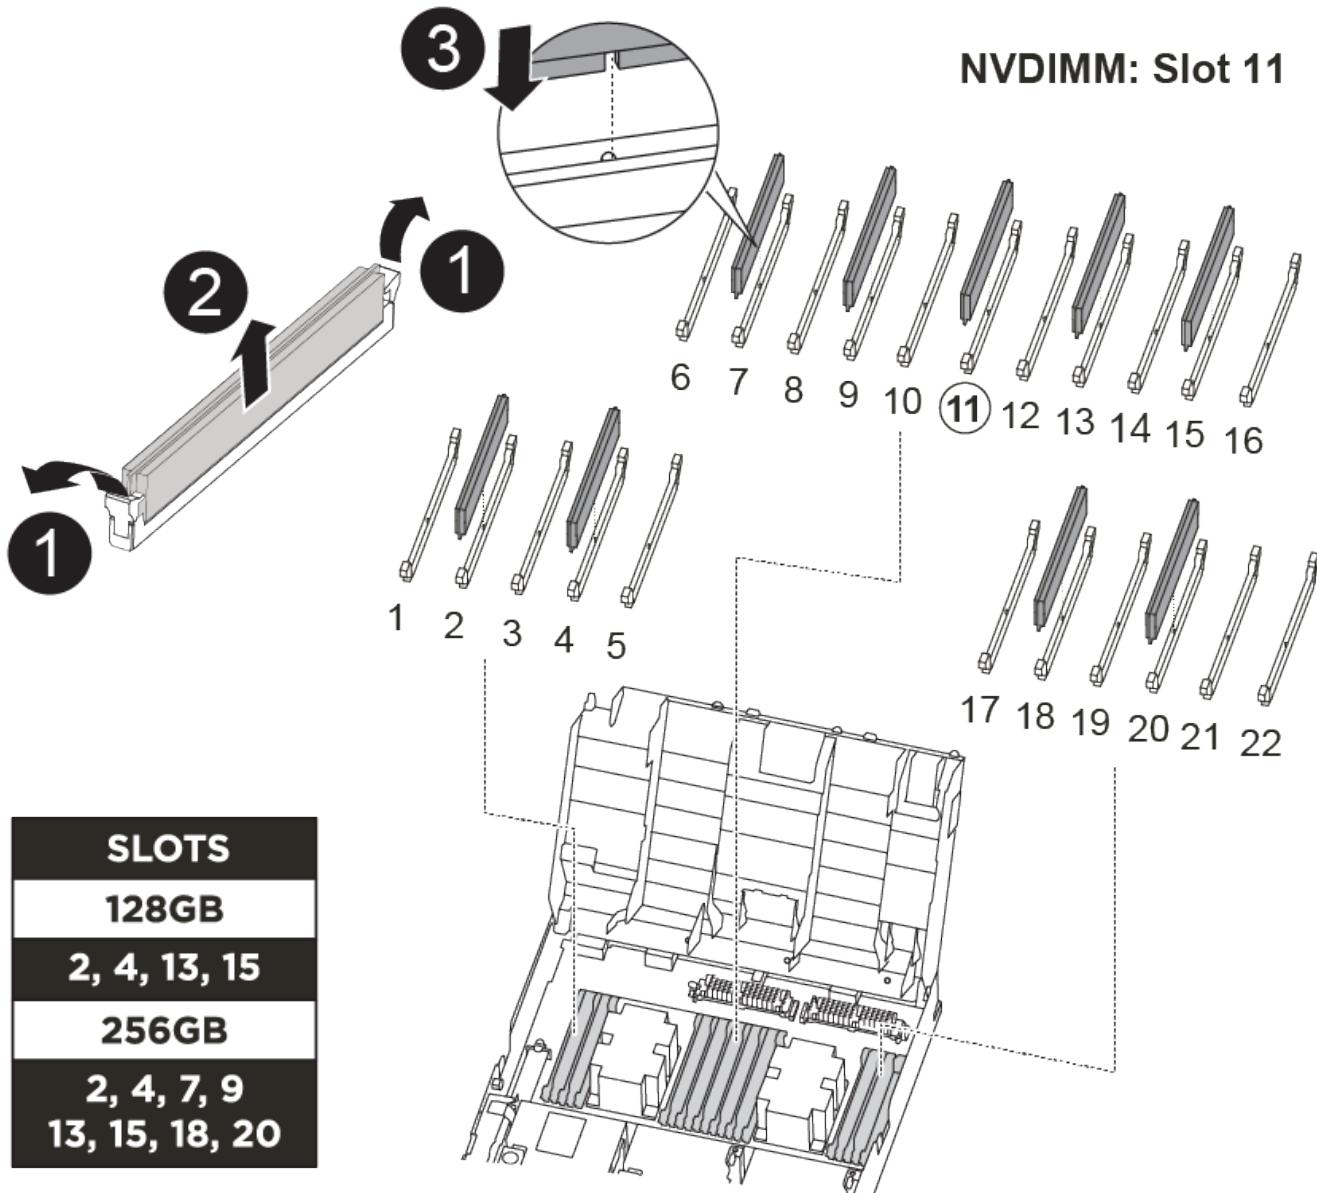
Step 7: Move the DIMMs

You need to locate the DIMMs, and then move them from the impaired controller module to the replacement controller module.

You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

You can use the following animation, illustration, or the written steps to move the DIMMs from the impaired controller module to the replacement controller module.

[Animation - Move the DIMMs](#)



1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Verify that the NVDIMM battery is not plugged into the new controller module.
4. Move the DIMMs from the impaired controller module to the replacement controller module:



Make sure that you install the each DIMM into the same slot it occupied in the impaired controller module.

- a. Eject the DIMM from its slot by slowly pushing apart the DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.



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

- b. Locate the corresponding DIMM slot on the replacement controller module.
- c. Make sure that the DIMM ejector tabs on the DIMM socket are in the open position, and then insert the DIMM squarely into the socket.

The DIMMs fit tightly in the socket, but should go in easily. If not, realign the DIMM with the socket and reinsert it.

- d. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.
- e. Repeat these substeps for the remaining DIMMs.

5. Plug the NVDIMM battery into the motherboard.

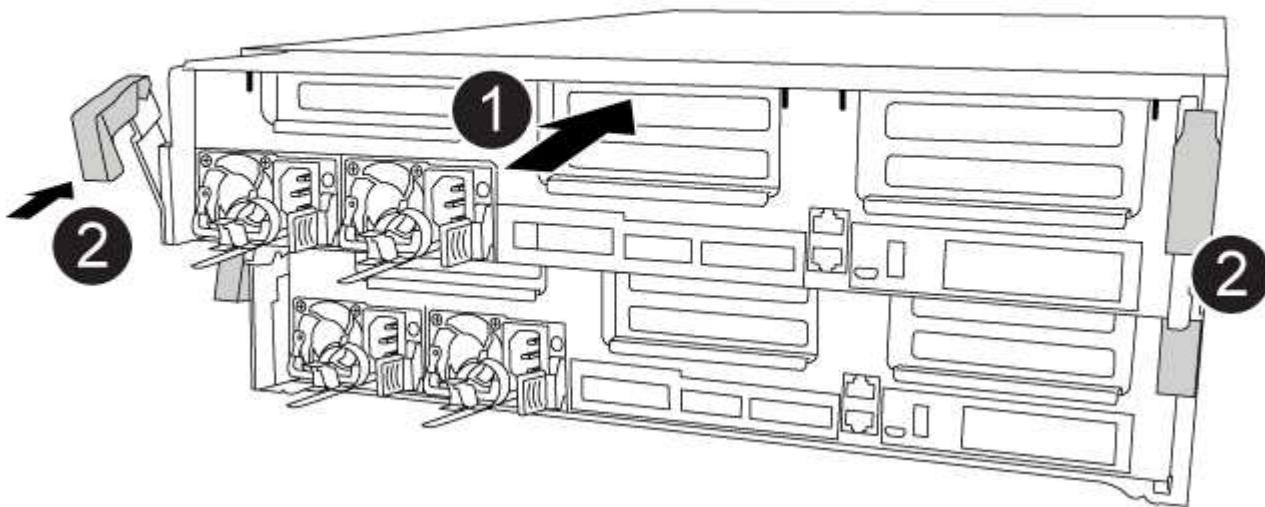
Make sure that the plug locks down onto the controller module.

Step 8: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis, and then boot it to Maintenance mode.

You can use the following animation, illustration, or the written steps to install the replacement controller module in the chassis.

[Animation - Install the controller module](#)



1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.



Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.



You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:

- a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

- Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

- b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

- c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

- d. If you have not already done so, reinstall the cable management device.

- e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.



If your system stops at the boot menu, select the option to boot to LOADER.

- f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

- g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

If your system stops at the boot menu, select the option to boot to LOADER.

Restore and verify the system configuration - FAS8300 and FAS8700

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.

- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.

2. On the *healthy* node, check the system time: `cluster date show`

The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`

5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`

6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the HA state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same HA state: `ha-config show`

The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the HA state for the controller module: `ha-config modify controller ha-state`

The value for `ha-state` can be one of the following:

- ha
- mcc
- mcc-2n
- mccip
- non-ha

3. If the displayed system state of the controller module does not match your system configuration, set the HA state for the controller module: `ha-config modify controller ha-state`

4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - FAS8300 and FAS8700

You must complete a series of tasks before restoring your system to full operation.

Step 1: Recable the system

Verify the controller module's storage and network connections by using [Active IQ Config Advisor](#).

Steps

1. Download and install Config Advisor.
2. Enter the information for the target system, and then click Collect Data.
3. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
4. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. In a stand-alone system, you must manually reassign the ID to the disks.

You must use the correct procedure for your configuration:

Controller redundancy	Then use this procedure...
HA pair	Option 1: Verify the system ID change on an HA system
Two-node MetroCluster configuration	Option 2: Manually reassign the system ID on systems in a two-node MetroCluster configuration

Option 1: Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the *> prompt, exit Maintenance mode and go to the LOADER prompt: `halt`)
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:
3. Wait until the `Waiting for giveback...` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```

node1> `storage failover show`  

                                         Takeover  

Node          Partner      Possible    State Description  

-----  -----  -----  

-----  

node1          node2      false      System ID changed on  

partner (Old:  

151759755, New:  

151759706), In takeover  

node2          node1      -          Waiting for giveback  

(HA mailboxes)

```

4. From the healthy controller, verify that any coredumps are saved:
 - a. Change to the advanced privilege level: `set -privilege advanced`
 You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (`*>`).
 - b. Save any coredumps: `system node run -node local-node-name partner savecore`
 - c. Wait for the `'savecore'` command to complete before issuing the giveback.
 You can enter the following command to monitor the progress of the `savecore` command: `system node run -node local-node-name partner savecore -s`
 - d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
 - [Restore onboard key management encryption keys](#)
 - [Restore external key management encryption keys](#)
6. Give back the controller:
 - a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`
 The `replacement` controller takes back its storage and completes booting.
 If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

 If the giveback is vetoed, you can consider overriding the vetoes.

[Find the High-Availability Configuration Guide for your version of ONTAP 9](#)

 - b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`
 The output from the `storage failover show` command should not include the System ID changed

on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`  
  
Disk  Aggregate Home  Owner  DR Home  Home ID      Owner ID  DR Home ID  
Reserver  Pool  
-----  -----  -----  -----  -----  -----  
-----  ---  
1.0.0  aggr0_1  node1 node1  -      1873775277 1873775277  -  
1873775277 Pool0  
1.0.1  aggr0_1  node1 node1      1873775277 1873775277  -  
1873775277 Pool0  
.  
.  
.
```

Option 2: Manually reassign the system ID on systems in a two-node MetroCluster configuration

In a two-node MetroCluster configuration running ONTAP, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems in a two-node MetroCluster configuration running ONTAP.

You must be sure to issue the commands in this procedure on the correct node:

- The *impaired* node is the node on which you are performing maintenance.
- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the DR partner of the impaired node.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by entering `Ctrl-C`, and then select the option to boot to Maintenance mode from the displayed menu.

You must enter `Y` when prompted to override the system ID due to a system ID mismatch.

2. View the old system IDs from the healthy node: ``metrocluster node show -fields node-systemid,dr-partner-systemid``

In this example, the `Node_B_1` is the old node, with the old system ID of 118073209:

```

dr-group-id cluster          node          node-systemid dr-
partner-systemid

-----
-----
```

1	Cluster_A	Node_A_1	536872914
118073209			
1	Cluster_B	Node_B_1	118073209
536872914			
2 entries were displayed.			

3. View the new system ID at the Maintenance mode prompt on the impaired node: `disk show`

In this example, the new system ID is 118065481:

```

Local System ID: 118065481
...
...
```

4. Reassign disk ownership (for FAS systems), by using the system ID information obtained from the `disk show` command: `disk reassign -s old system ID`

In the case of the preceding example, the command is: `disk reassign -s 118073209`

You can respond `Y` when prompted to continue.

5. Verify that the disks were assigned correctly: `disk show -a`

Verify that the disks belonging to the *replacement* node show the new system ID for the *replacement* node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

```

*> disk show -a
Local System ID: 118065481

      DISK      OWNER          POOL      SERIAL NUMBER      HOME
-----  -----  -----  -----  -----  -----
disk_name  system-1  (118065481)  Pool0  J8Y0TDZC  system-1
(118065481)
disk_name  system-1  (118065481)  Pool0  J8Y09DXC  system-1
(118065481)
.
.
.
```

6. From the healthy node, verify that any coredumps are saved:

- a. Change to the advanced privilege level: `set -privilege advanced`

You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (`*>`).

- b. Verify that the coredumps are saved: `system node run -node local-node-name partner savecore`

If the command output indicates that `savecore` is in progress, wait for `savecore` to complete before issuing the giveback. You can monitor the progress of the `savecore` using the `system node run -node local-node-name partner savecore -s` command. </info>

- c. Return to the admin privilege level: `set -privilege admin`

7. If the *replacement* node is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode and go to the LOADER prompt: `halt`
8. Boot the *replacement* node: `boot_ontap`
9. After the *replacement* node has fully booted, perform a switchback: `metrocluster switchback`
10. Verify the MetroCluster configuration: `metrocluster node show -fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state
```

dr-group-id	cluster node	configuration-state
1 node1_siteA	node1mcc-001	configured
1 node1_siteA	node1mcc-002	configured
1 node1_siteB	node1mcc-003	configured
1 node1_siteB	node1mcc-004	configured

4 entries were displayed.

11. Verify the operation of the MetroCluster configuration in Data ONTAP:

- a. Check for any health alerts on both clusters: `system health alert show`
- b. Confirm that the MetroCluster is configured and in normal mode: `metrocluster show`
- c. Perform a MetroCluster check: `metrocluster check run`
- d. Display the results of the MetroCluster check: `metrocluster check show`
- e. Run Config Advisor. Go to the Config Advisor page on the NetApp Support Site at support.netapp.com/NOW/download/tools/config_advisor/.

After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

12. Simulate a switchover operation:

- a. From any node's prompt, change to the advanced privilege level: `set -privilege advanced`

You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (`*>`).

- b. Perform the switchback operation with the `-simulate` parameter: `metrocluster switchover -simulate`
- c. Return to the admin privilege level: `set -privilege admin`

Complete system restoration - FAS8300 and FAS8700

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

 If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in [Post Motherboard Replacement Process to update Licensing on a AFF/FAS system](#). If you are unsure of the initial ONTAP release for your system, see [NetApp Hardware Universe](#) for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the [NetApp Support Site](#) in the My Support section under Software licenses.

 The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.
2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
 - a. Check for unused licenses: `license clean-up -unused -simulate`

b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.

- If AutoSupport is enabled, send an AutoSupport message to register the serial number.
- If AutoSupport is not enabled, call [NetApp Support](#) to register the serial number.

3. Check the health of your cluster. See the [How to perform a cluster health check with a script in ONTAP KB](#) article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the enabled state: `metrocluster node show`

```
cluster_B::> metrocluster node show

DR          Configuration  DR
Group Cluster Node  State      Mirroring Mode
-----  -----  -----
-----  -----
1      cluster_A
      controller_A_1 configured  enabled  heal roots
completed
      cluster_B
      controller_B_1 configured  enabled  waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`

3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster          Configuration State      Mode
-----
Local: cluster_B configured          switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster          Configuration State      Mode
-----
Local: cluster_B configured          normal
Remote: cluster_A configured        normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Replace a DIMM - FAS8300 and FAS8700

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

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.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage

system hardware configuration.

Option 1: Most configurations

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

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, **quorum status** of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

Steps

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

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a two-node MetroCluster

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

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller...	Then...
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: <code>metrocluster switchover</code>
Has not automatically switched over, you attempted switchover with the <code>metrocluster switchover</code> command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
  Operation: heal-aggregates
    State: successful
  Start Time: 7/25/2016 18:45:55
  End Time: 7/25/2016 18:45:56
  Errors: -
```

5. Check the state of the aggregates by using the storage aggregate show command.

```
controller_A_1::> storage aggregate show
  Aggregate      Size Available Used% State  #Vols  Nodes
  RAID  Status
  -----  -----
  ...
  aggr_b2      227.1GB    227.1GB    0% online      0  mcc1-a2
  raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the metrocluster heal -phase root-aggregates command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the metrocluster heal command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the metrocluster operation show command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
    State: successful
  Start Time: 7/29/2016 20:54:41
  End Time: 7/29/2016 20:54:42
  Errors: -
```

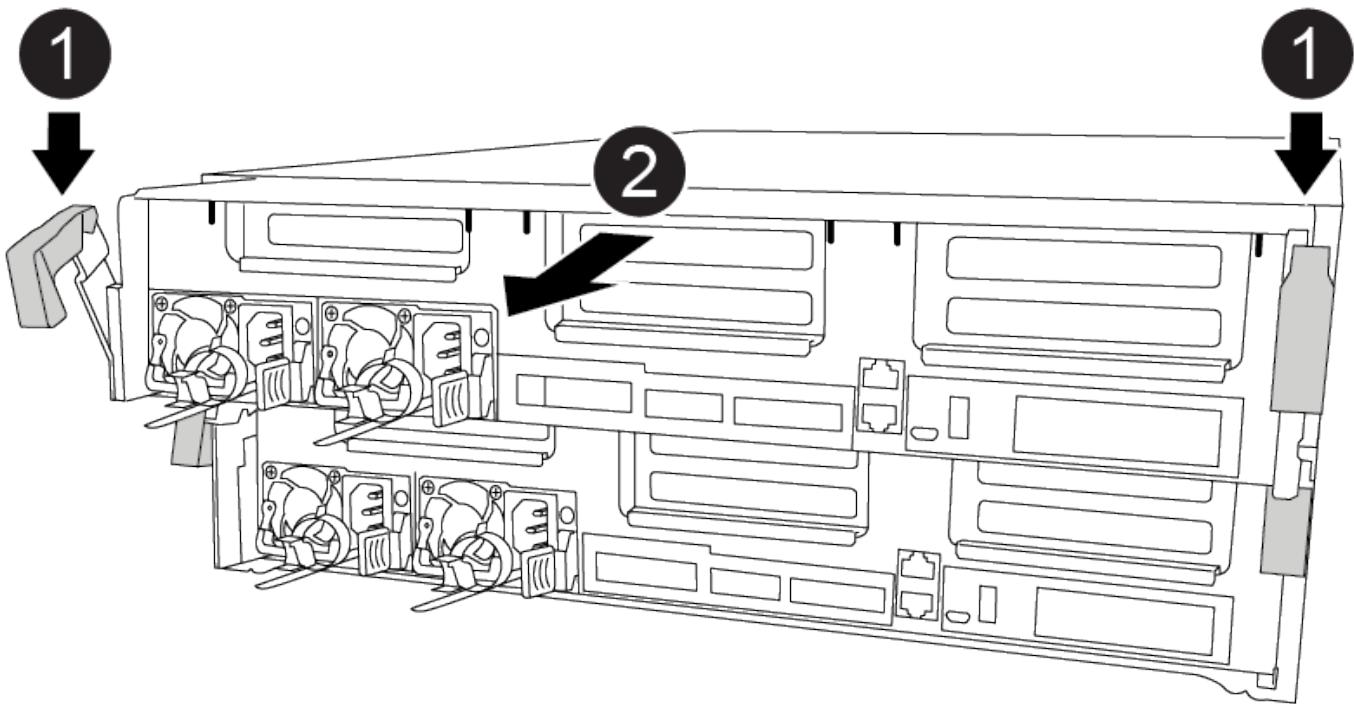
8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module



Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. 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, keeping 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.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace system DIMMs

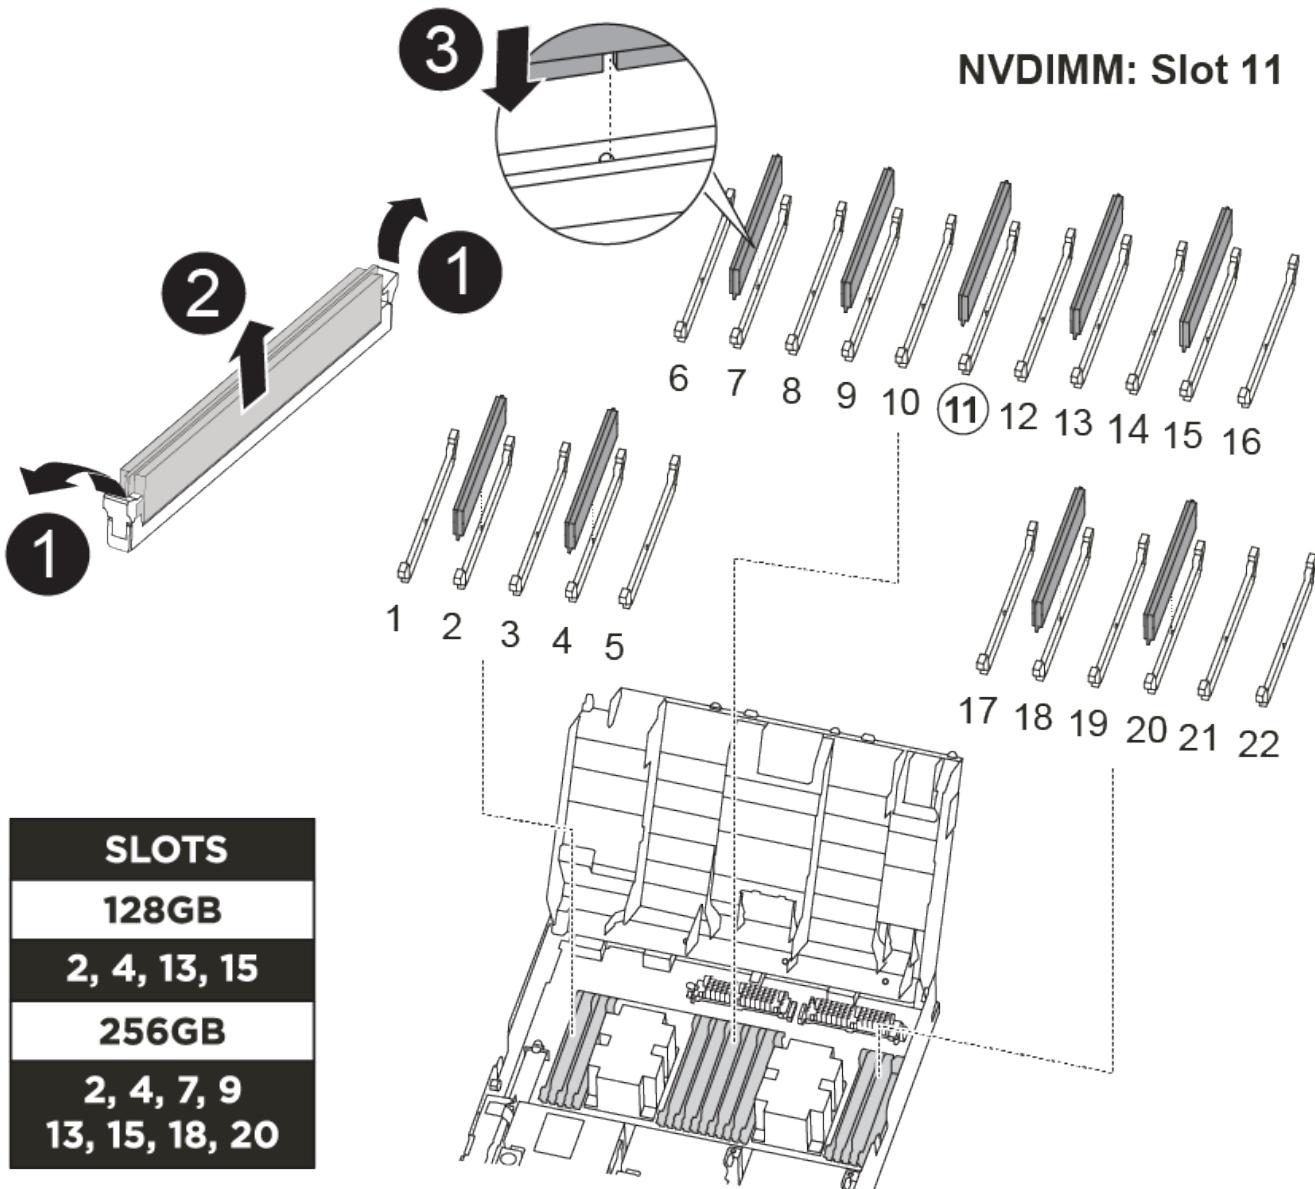
Replacing a system DIMM involves identifying the target DIMM through the associated error message, locating the target DIMM using the FRU map on the air duct, and then replacing the DIMM.

You can use the following animation, illustration, or the written steps to replace a system DIMM.



The animation and illustration shows empty slots for sockets without DIMMs. These empty sockets are populated with blanks.

[Animation - Replace a system DIMM](#)



The number and location of DIMMs in your system depends on the model of your system. Refer to FRU map on the air duct for more information.

- If you have a FAS8300 system, the system DIMMs are located in sockets 2, 4, 13, and 15.

- If you have a FAS8700 system, the system DIMMs are located in slots 2, 4, 7, 9, 13, 15, 18, and 20.
- The NVDIMM is located in slot 11.

Steps

1. Open the air duct:
 - a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
 - b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the DIMMs on your controller module.
3. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
4. Eject the DIMM from its socket by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the socket.



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

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then 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.



Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

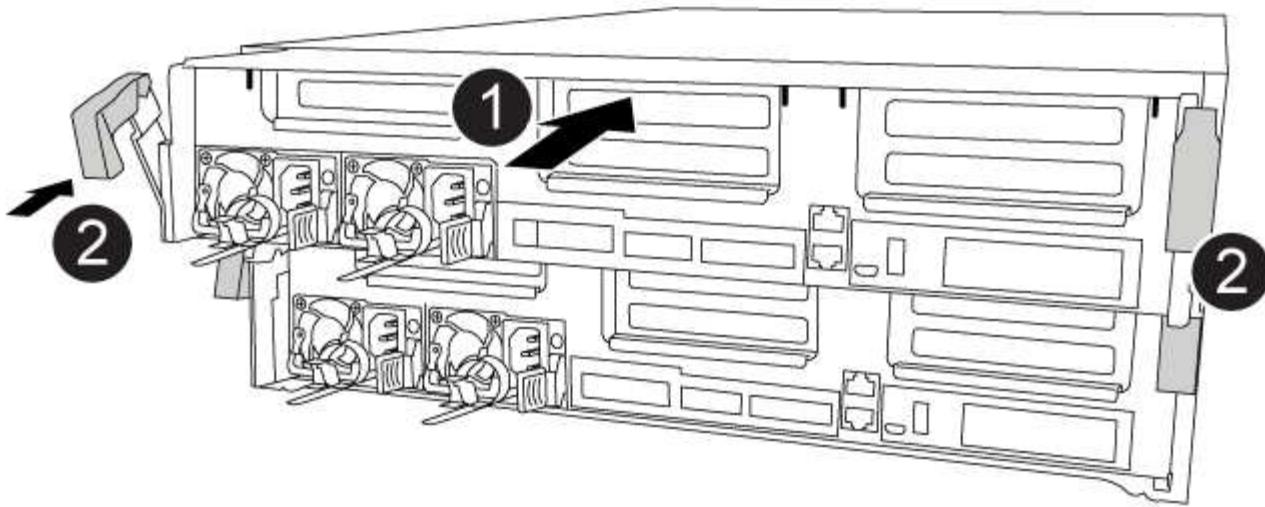
7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis.

You can use the following animation, drawing, or the written steps to install the controller module in the chassis.

[Animation - Install the controller module](#)



Steps

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

 Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

 You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
 - a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

 Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

- b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
- c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

- d. If you have not already done so, reinstall the cable management device.
- e. Interrupt the normal boot process and boot to LOADER by pressing **Ctrl-C**.



If your system stops at the boot menu, select the option to boot to LOADER.

- f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::> metrocluster node show

DR          Configuration  DR
Group Cluster Node  State      Mirroring Mode
-----  -----
-----  -----
1      cluster_A
          controller_A_1 configured  enabled  heal roots
completed
      cluster_B
          controller_B_1 configured  enabled  waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster          Configuration State    Mode
-----
Local: cluster_B configured          switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster          Configuration State    Mode
-----
Local: cluster_B configured          normal
Remote: cluster_A configured        normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Hot-swap a fan module - FAS8300 and FAS8700

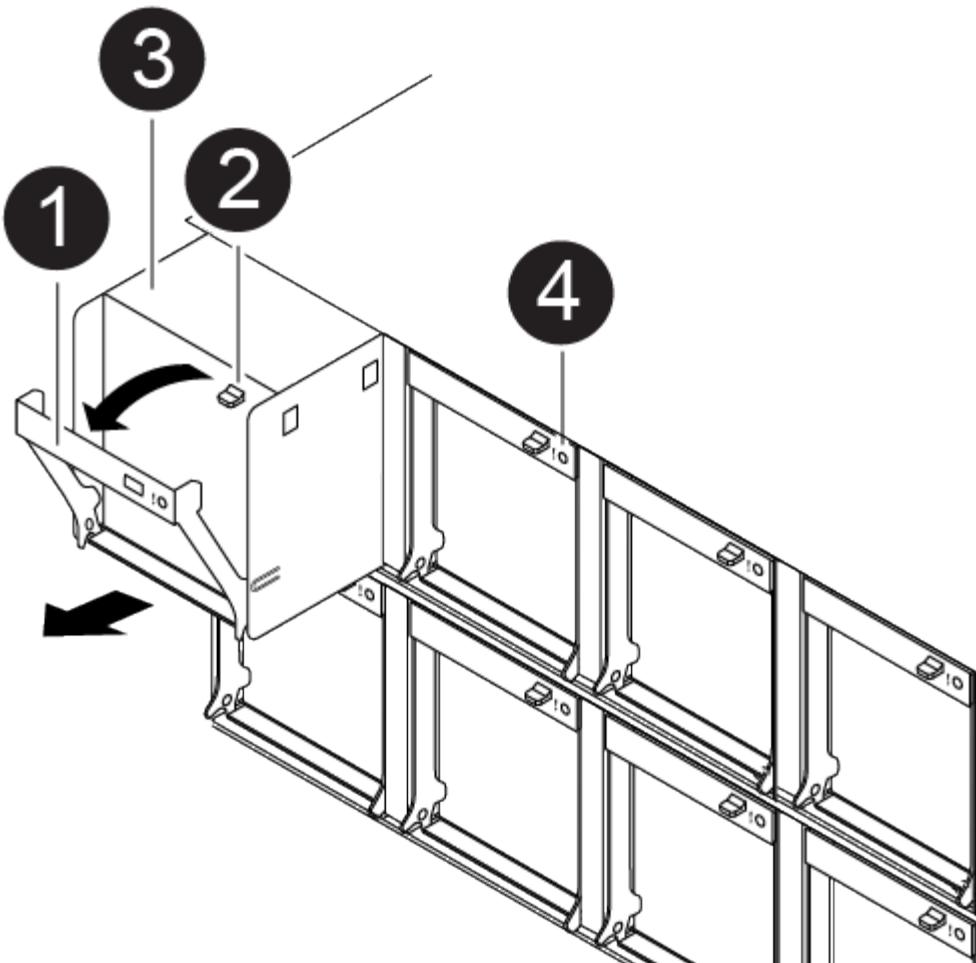
To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.



You must replace the fan module within two minutes of removing it from the chassis. System airflow is disrupted and the controller module or modules shut down after two minutes to avoid overheating.

You can use the following animation, illustration, or the written steps to hot-swap a fan module.

[Animation - Replace a fan](#)



Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

The fan module moves a little bit away from the chassis.

5. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.



The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

6. Set the fan module aside.
7. Insert the replacement fan module into the chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

The Attention LED should not be lit after the fan is seated and has spun up to operational speed.

10. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Replace an NVDIMM - FAS8300 and FAS8700

You must replace the NVDIMM in the controller module when your system registers that the flash lifetime is almost at an end or that the identified NVDIMM is not healthy in general; failure to do so causes a system panic.

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.

Step 1: Shut down the impaired controller

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

Option 1: Most configurations

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

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, **quorum status** of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

Steps

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

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a two-node MetroCluster

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

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller...	Then...
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: <code>metrocluster switchover</code>
Has not automatically switched over, you attempted switchover with the <code>metrocluster switchover</code> command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
  Operation: heal-aggregates
    State: successful
  Start Time: 7/25/2016 18:45:55
  End Time: 7/25/2016 18:45:56
  Errors: -
```

5. Check the state of the aggregates by using the storage aggregate show command.

```
controller_A_1::> storage aggregate show
  Aggregate      Size Available Used% State  #Vols  Nodes
  RAID  Status
  -----  -----
  ...
  aggr_b2      227.1GB    227.1GB    0% online      0 mcc1-a2
  raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the metrocluster heal -phase root-aggregates command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the metrocluster heal command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the metrocluster operation show command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
    State: successful
  Start Time: 7/29/2016 20:54:41
  End Time: 7/29/2016 20:54:42
  Errors: -
```

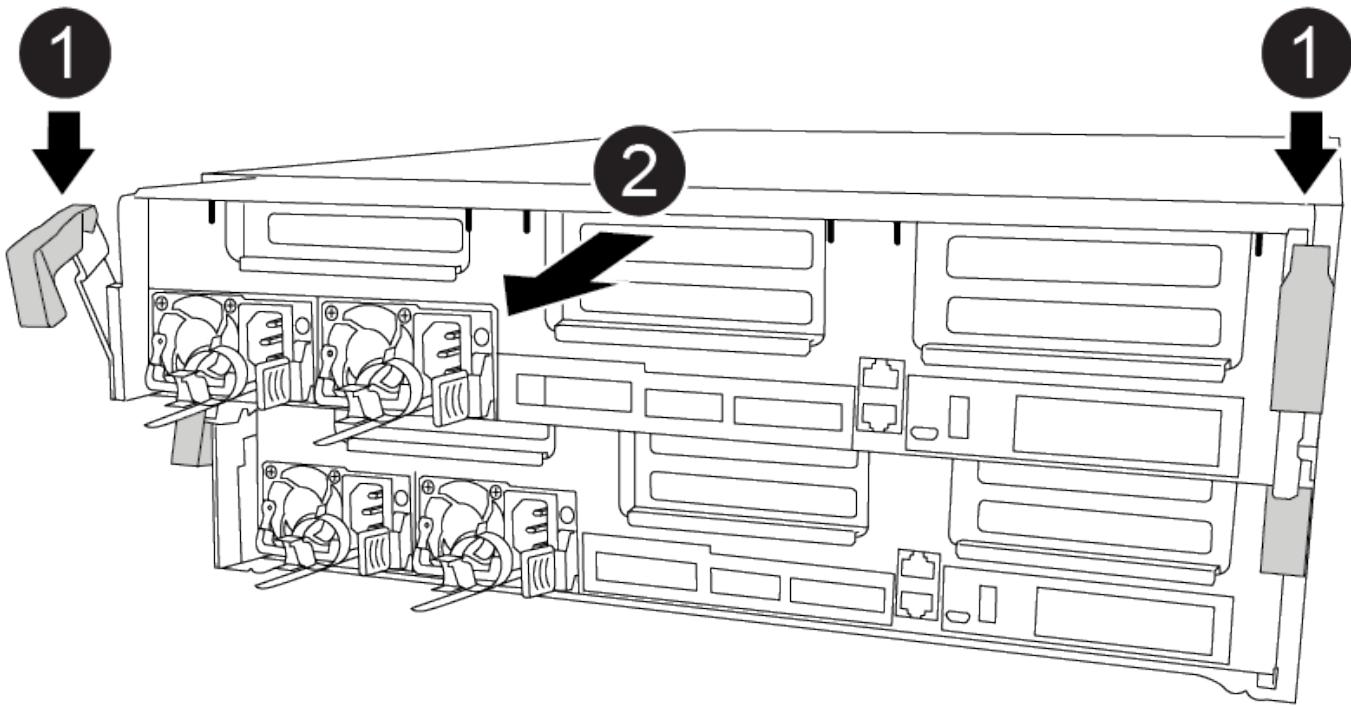
8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following , illustration, or the written steps to remove the controller module from the chassis.

[Animation - Remove the controller module](#)



Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. 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, keeping 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.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace the NVDIMM

To replace the NVDIMM, you must locate it in the controller module using the FRU map on top of the air duct the FRU Map on the top of the slot 1 riser.

- The NVDIMM LED blinks while destaging contents when you halt the system. After the destage is complete, the LED turns off.
- Although the contents of the NVDIMM is encrypted, it is a best practice to erase the contents of the NVDIMM before replacing it. For more information, see the [Statement of Volatility](#) on the NetApp Support Site.



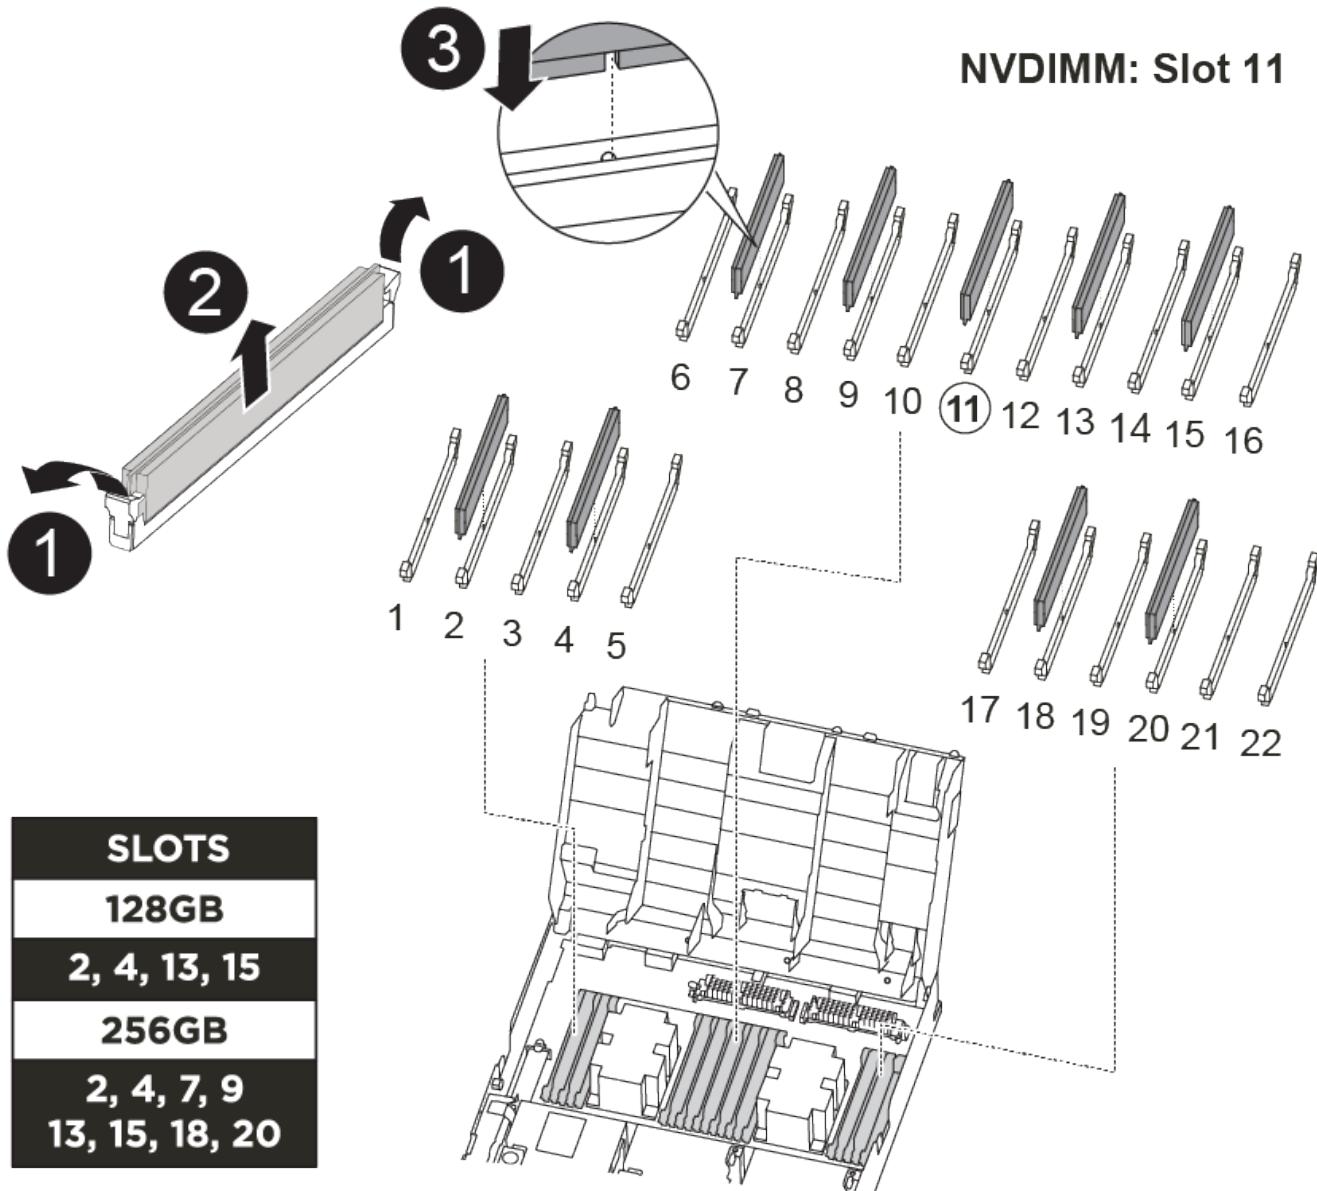
You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

You can use the following animation, illustration, or the written steps to replace the NVDIMM.



The animation and illustration show empty slots for sockets without DIMMs. These empty sockets are populated with blanks.

[Animation - Replace the NVDIMM](#)



Steps

1. Open the air duct and then locate the NVDIMM in slot 11 on your controller module.



The NVDIMM looks significantly different than system DIMMs.

2. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.



Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

3. Remove the replacement NVDIMM from the antistatic shipping bag, hold the NVDIMM by the corners, and then align it to the slot.

The notch among the pins on the NVDIMM should line up with the tab in the socket.

4. Locate the slot where you are installing the NVDIMM.

5. Insert the NVDIMM squarely into the slot.

The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.



Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.

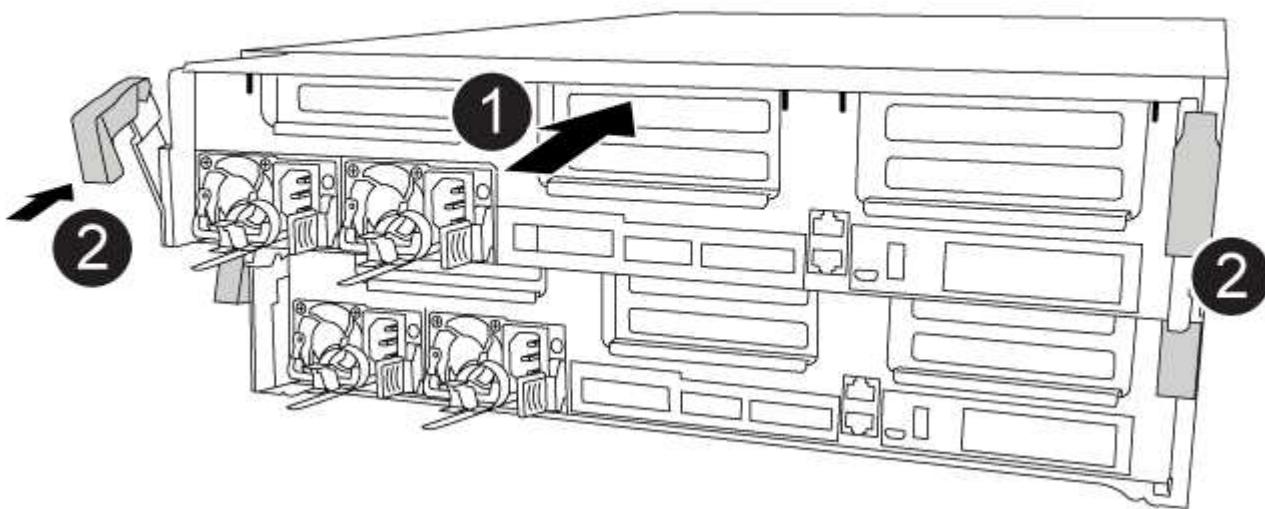
7. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module



Steps

1. If you have not already done so, close the air duct.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.



Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.



You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:

- a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

- Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

- b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

- c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

- d. If you have not already done so, reinstall the cable management device.

- e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.



If your system stops at the boot menu, select the option to boot to LOADER.

- f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode `impaired_node_name``

3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the enabled state: `metrocluster node show`

```

cluster_B::> metrocluster node show

DR          Configuration DR
Group Cluster Node  State      Mirroring Mode
-----  -----  -----
-----  -----
1      cluster_A
      controller_A_1 configured  enabled  heal roots
completed
      cluster_B
      controller_B_1 configured  enabled  waiting for
switchback recovery
2 entries were displayed.

```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```

cluster_B::> metrocluster show
Cluster      Configuration State      Mode
-----  -----  -----
Local: cluster_B configured      switchover
Remote: cluster_A configured    waiting-for-switchback

```

The switchback operation is complete when the clusters are in the `normal` state.:

```

cluster_B::> metrocluster show
Cluster      Configuration State      Mode
-----  -----  -----
Local: cluster_B configured      normal
Remote: cluster_A configured    normal

```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Replace the NVDIMM battery - FAS8300 and FAS8700

To replace the NVDIMM battery, you must remove the controller module, remove the battery, replace the battery, and then reinstall the controller module.

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

Step 1: Shut down the impaired controller

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

Option 1: Most configurations

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

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, **quorum status** of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

Steps

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

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a two-node MetroCluster

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

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller...	Then...
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: <code>metrocluster switchover</code>
Has not automatically switched over, you attempted switchover with the <code>metrocluster switchover</code> command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
  Operation: heal-aggregates
    State: successful
  Start Time: 7/25/2016 18:45:55
  End Time: 7/25/2016 18:45:56
  Errors: -
```

5. Check the state of the aggregates by using the storage aggregate show command.

```
controller_A_1::> storage aggregate show
  Aggregate      Size Available Used% State  #Vols  Nodes
  RAID  Status
  -----  -----
  ...
  aggr_b2      227.1GB    227.1GB    0% online      0  mcc1-a2
  raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the metrocluster heal -phase root-aggregates command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the metrocluster heal command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the metrocluster operation show command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
    State: successful
  Start Time: 7/29/2016 20:54:41
  End Time: 7/29/2016 20:54:42
  Errors: -
```

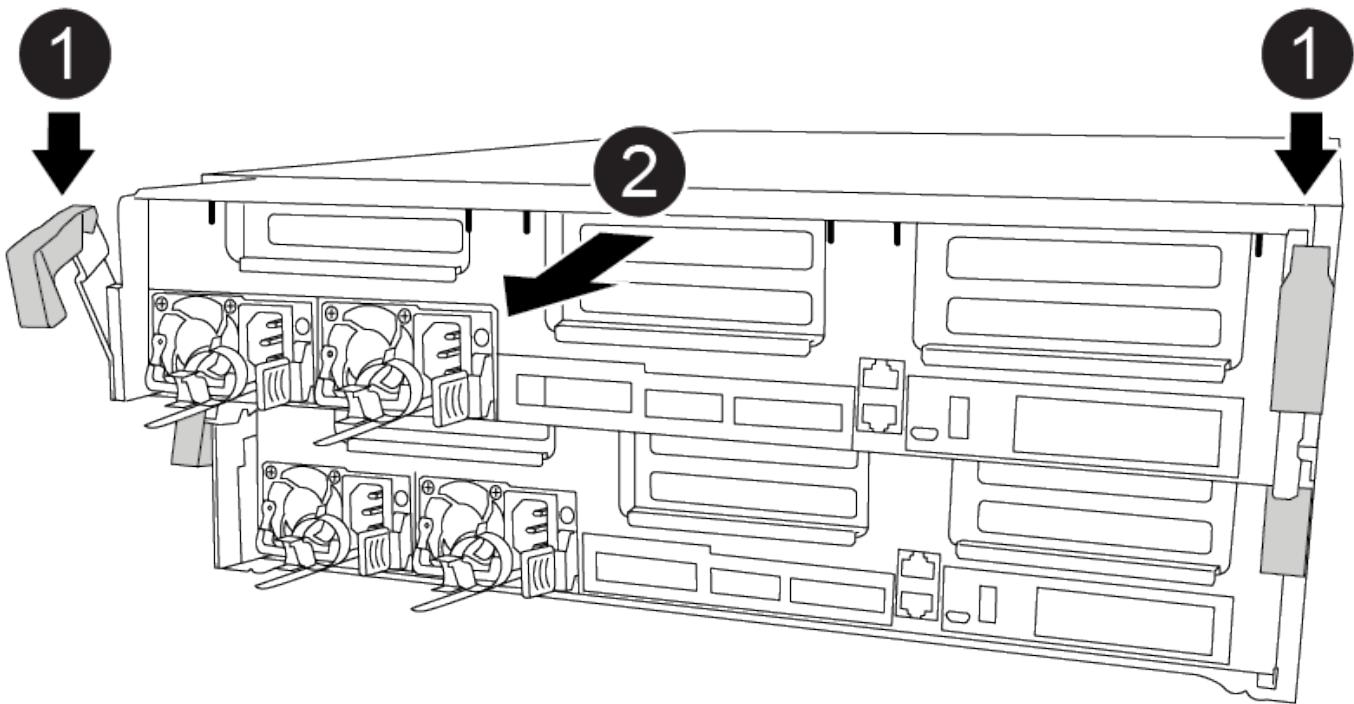
8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module



Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. 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, keeping 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.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

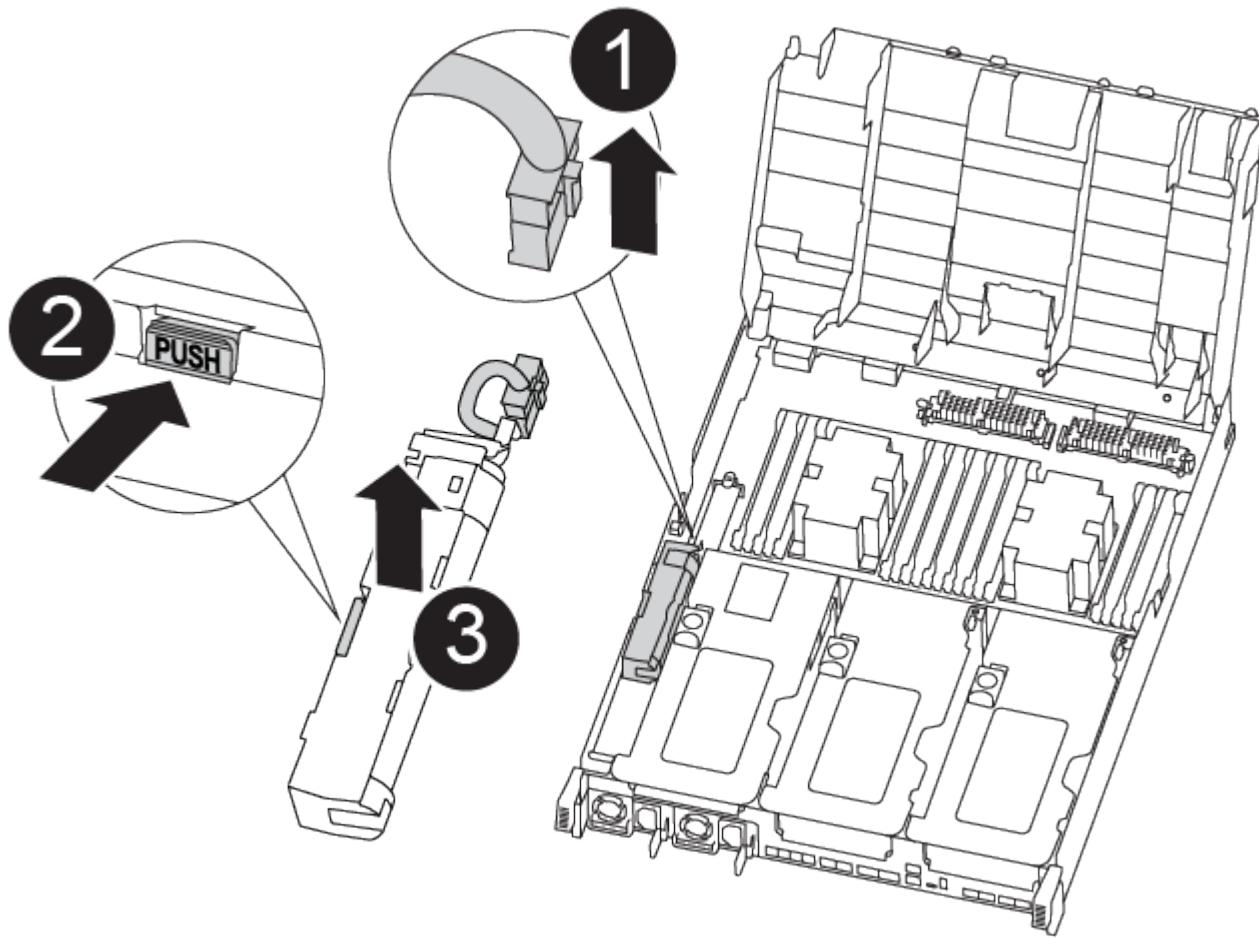
Step 3: Replace the NVDIMM battery

To replace the NVDIMM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module. See the FRU map inside the controller module to locate the NVDIMM battery.

The NVDIMM LED blinks while destaging contents when you halt the system. After the destage is complete, the LED turns off.

You can use the following animation, illustration, or the written steps to replace the NVDIMM battery.

[Animation - Replace the NVDIMM battery](#)



Steps

1. Open the air duct:
 - a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
 - b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the NVDIMM battery in the controller module.
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. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder

and controller module.

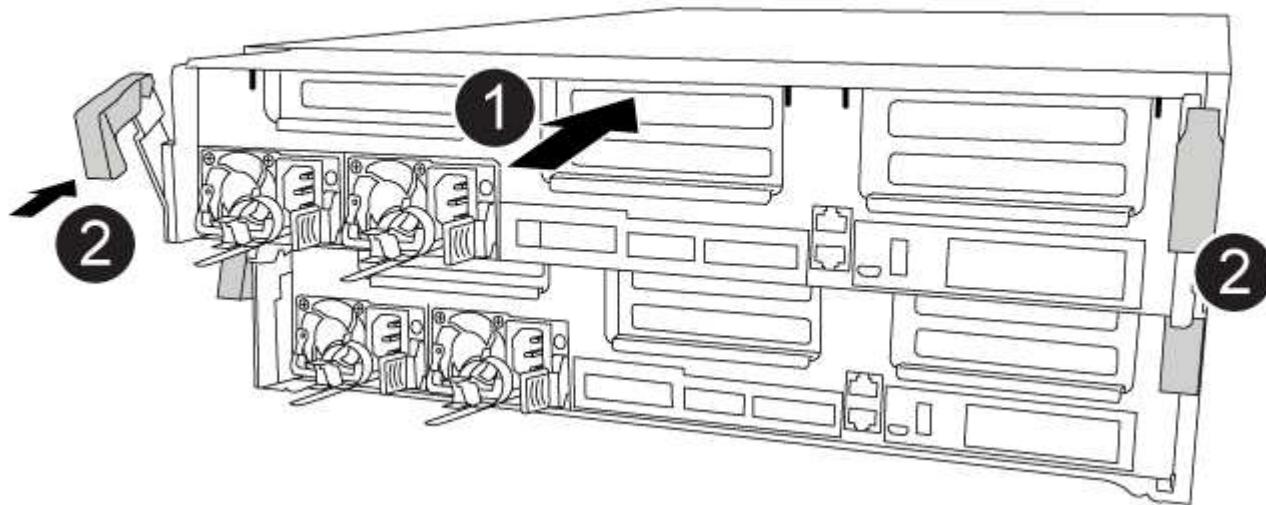
5. Remove the replacement battery from its package.
6. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.
7. Plug the battery plug back into the controller module, and then close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module



Steps

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.



Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.



You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:

- a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.



Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

- b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
- c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

- d. If you have not already done so, reinstall the cable management device.
- e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.



If your system stops at the boot menu, select the option to boot to LOADER.

- f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::> metrocluster node show

DR          Configuration DR
Group Cluster Node  State   Mirroring Mode
-----  -----  -----  -----
-----  -----
1   cluster_A
      controller_A_1 configured   enabled   heal roots
completed
      cluster_B
      controller_B_1 configured   enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster          Configuration State   Mode
-----  -----  -----  -----
Local: cluster_B configured      switchover
Remote: cluster_A configured    waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster          Configuration State   Mode
-----  -----  -----  -----
Local: cluster_B configured      normal
Remote: cluster_A configured    normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Replace a PCIe or mezzanine card - FAS8300 and FAS8700

To replace a PCIe or mezzanine card, you must disconnect the cables and any SFP and QSFP modules from the cards, replace the failed PCIe or mezzanine card, and then reassemble the cards.

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

Step 1: Shut down the impaired controller

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

Option 1: Most configurations

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

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, **quorum status** of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

Steps

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

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

- b. Enter *y* when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a two-node MetroCluster

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

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller...	Then...
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: <code>metrocluster switchover</code>
Has not automatically switched over, you attempted switchover with the <code>metrocluster switchover</code> command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
  Operation: heal-aggregates
    State: successful
  Start Time: 7/25/2016 18:45:55
  End Time: 7/25/2016 18:45:56
  Errors: -
```

5. Check the state of the aggregates by using the storage aggregate show command.

```
controller_A_1::> storage aggregate show
  Aggregate      Size Available Used% State  #Vols  Nodes
  RAID  Status
  -----  -----
  ...
  aggr_b2      227.1GB    227.1GB    0% online      0 mcc1-a2
  raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the metrocluster heal -phase root-aggregates command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the metrocluster heal command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the metrocluster operation show command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
    State: successful
  Start Time: 7/29/2016 20:54:41
  End Time: 7/29/2016 20:54:42
  Errors: -
```

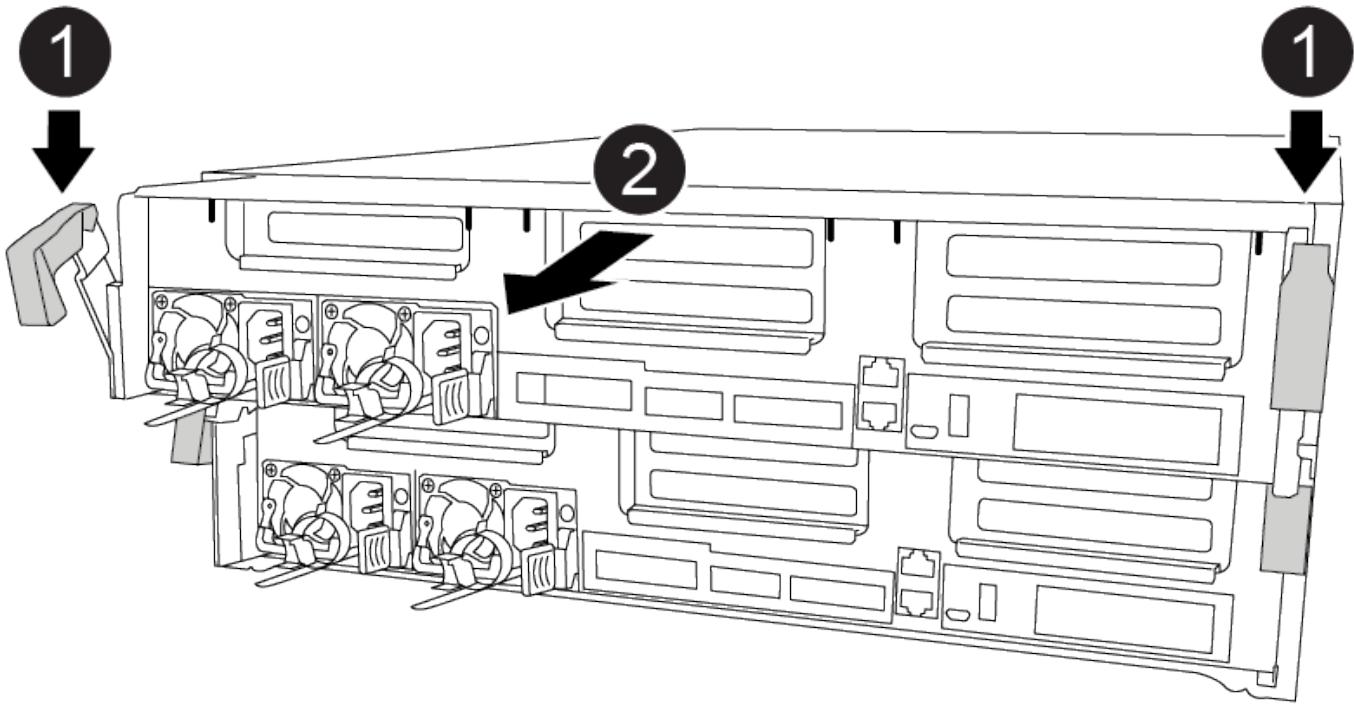
8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module



Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. 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, keeping 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.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

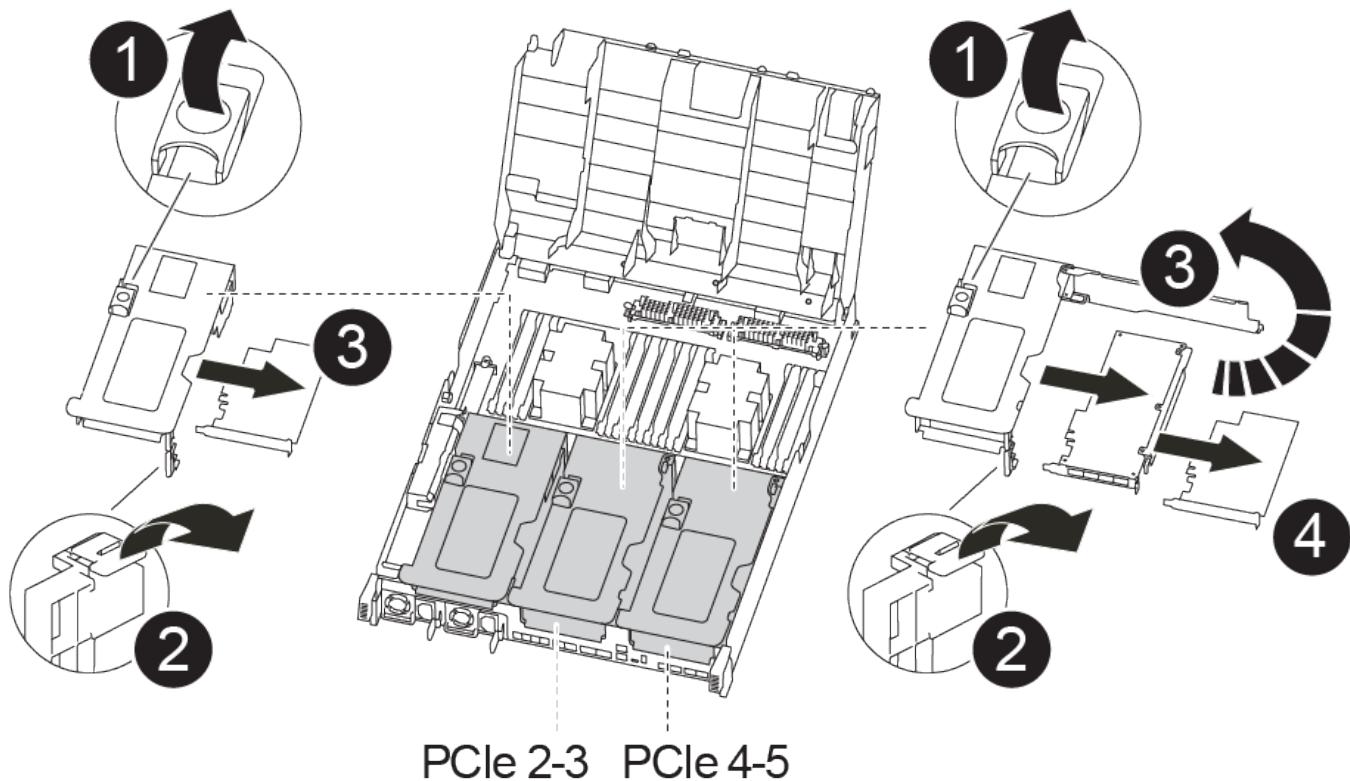
7. Place the controller module on a stable, flat surface.

Step 3: Replace a PCIe card

To replace a PCIe card, you must locate the failed PCIe card, remove the riser that contains the card from the controller module, replace the card, and then reinstall the PCIe riser in the controller module.

You can use the following animation, illustration, or the written steps to replace a PCIe card.

Animation - Replace a PCIe card



Steps

1. Remove the riser containing the card to be replaced:
 - a. Open the air duct by pressing the locking tabs on the sides of the air duct, slide it toward the back of the controller module, and then rotate it to its completely open position.
 - b. Remove any SFP or QSFP modules that might be in the PCIe cards.
 - c. Rotate the riser locking latch on the left side of the riser up and toward air duct.
The riser raises up slightly from the controller module.
 - d. Lift the riser up straight up and set it aside on a stable flat surface,
2. Remove the PCIe card from the riser:
 - a. Turn the riser so that you can access the PCIe card.
 - b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
 - c. For risers 2 and 3 only, swing the side panel up.
 - d. Remove the PCIe card from the riser by gently pushing up on the bracket and lift the card straight out of the socket.

3. Install the replacement PCIe card in the riser by aligning the card with the socket, press the card into the socket and then close the side panel on the riser, if present.

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



If you are installing a card in the bottom slot and cannot see the card socket well, remove the top card so that you can see the card socket, install the card, and then reinstall the card you removed from the top slot.

4. Reinstall the riser:

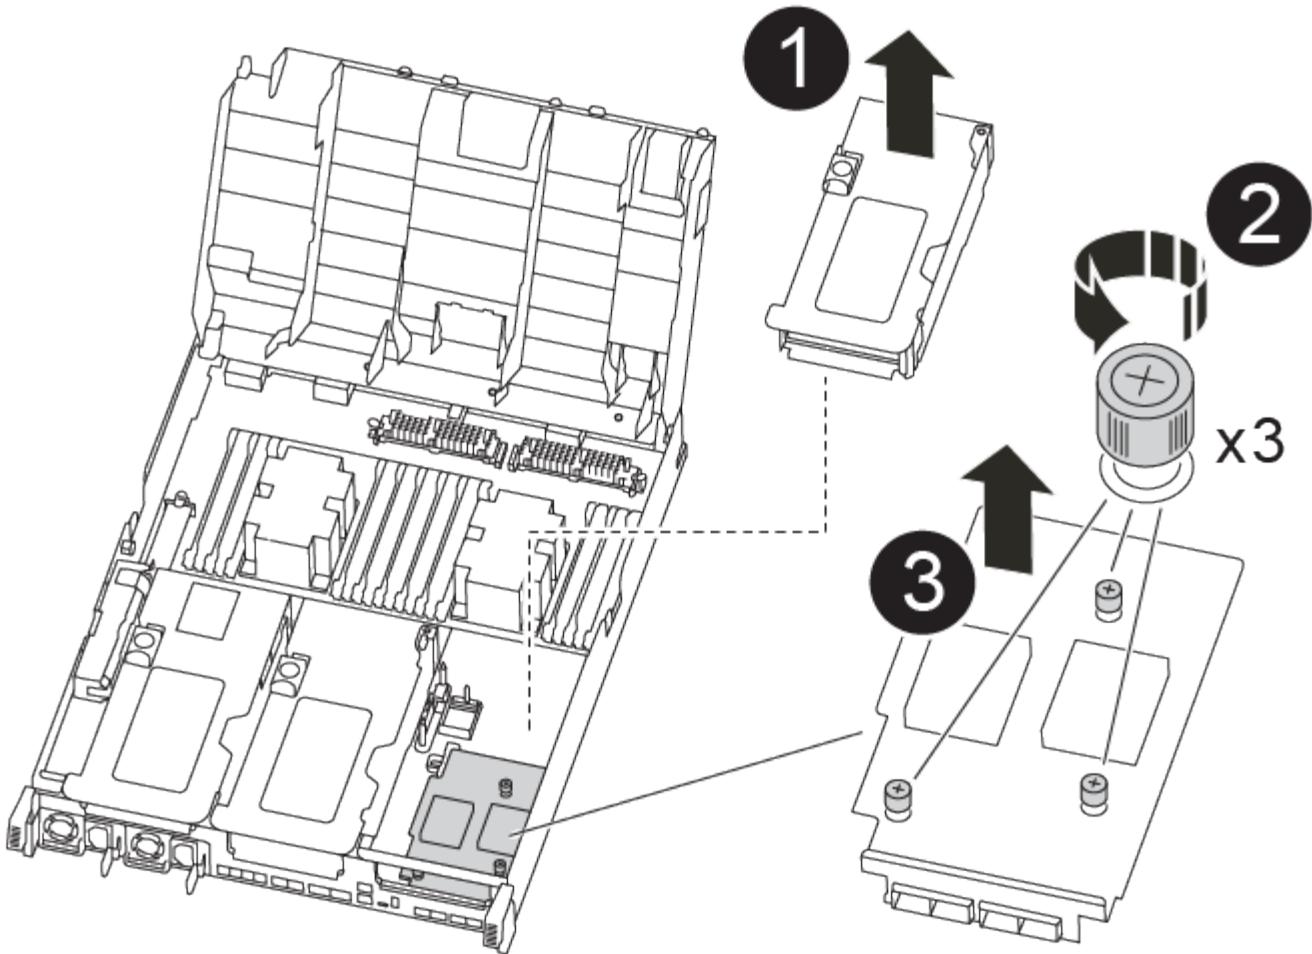
- a. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins.
- b. Push the riser squarely into the socket on the motherboard.
- c. Rotate the latch down flush with the sheet metal on the riser.

Step 4: Replace the mezzanine card

The mezzanine card is located under riser number 3 (slots 4 and 5). You must remove that riser to access the mezzanine card, replace the mezzanine card, and then reinstall riser number 3. See the FRU map on the controller module for more information.

You can use the following animation, illustration, or the written steps to replace the mezzanine card.

[Animation - Replace the mezzanine card](#)



Steps

1. Remove riser number 3 (slots 4 and 5):
 - a. Open the air duct by pressing the locking tabs on the sides of the air duct, slide it toward the back of the controller module, and then rotate it to its completely open position.
 - b. Remove any SFP or QSFP modules that might be in the PCIe cards.
 - c. Rotate the riser locking latch on the left side of the riser up and toward air duct.
 - d. Lift the riser up, and then set it aside on a stable, flat surface.
2. Replace the mezzanine card:
 - a. Remove any QSFP or SFP modules from the card.
 - b. Loosen the thumbscrews on the mezzanine card, and gently lift the card directly out of the socket and set it aside.
 - c. Align the replacement mezzanine card over the socket and the guide pins and gently push the card into the socket.
 - d. Tighten the thumbscrews on the mezzanine card.
3. Reinstall the riser:
 - a. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins.

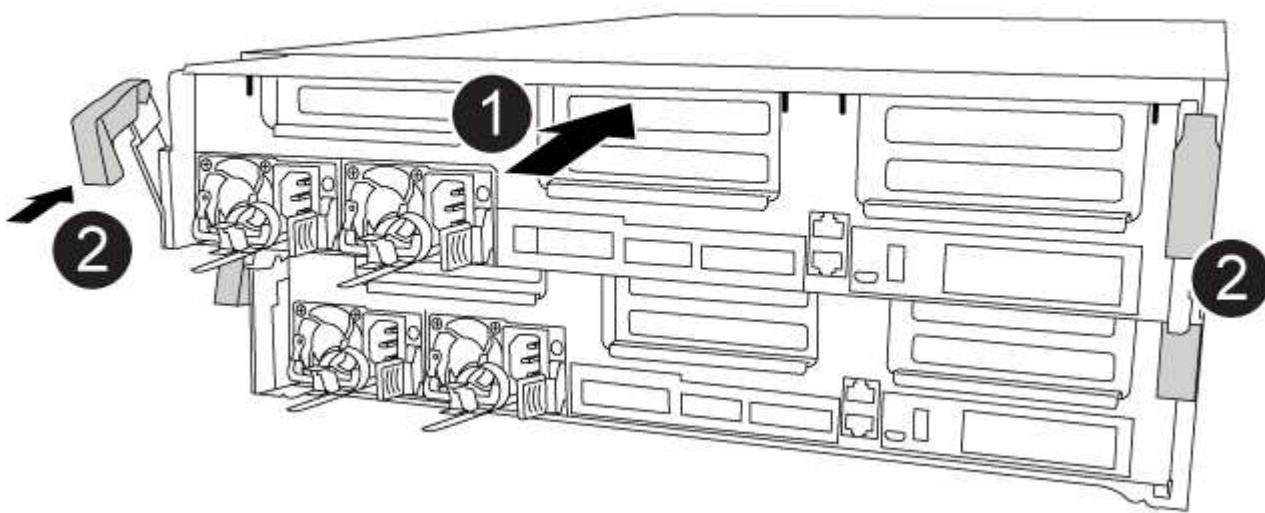
- b. Push the riser squarely into the socket on the motherboard.
- c. Rotate the latch down flush with the sheet metal on the riser.

Step 5: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

[Animation - Install the controller module](#)



Steps

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.



Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the installation of the controller module:
 - a. Using the locking latches, firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latches rise when the controller module is fully seated.



Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

- b. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

- c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

- d. If you have not already done so, reinstall the cable management device.

- e. Interrupt the normal boot process and boot to LOADER by pressing **Ctrl-C**.



If your system stops at the boot menu, select the option to boot to LOADER.

- f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.

5. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`

6. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the enabled state: `metrocluster node show`

```
cluster_B::> metrocluster node show

DR          Configuration  DR
Group Cluster Node  State      Mirroring Mode
-----  -----
-----  -----
1      cluster_A
      controller_A_1 configured  enabled  heal roots
completed
      cluster_B
      controller_B_1 configured  enabled  waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`

3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster          Configuration State      Mode
-----
Local: cluster_B configured          switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster          Configuration State      Mode
-----
Local: cluster_B configured          normal
Remote: cluster_A configured        normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 8: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Hot-swap a power supply - FAS8300 and FAS8700

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting the replacement PSU to the power source.

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.



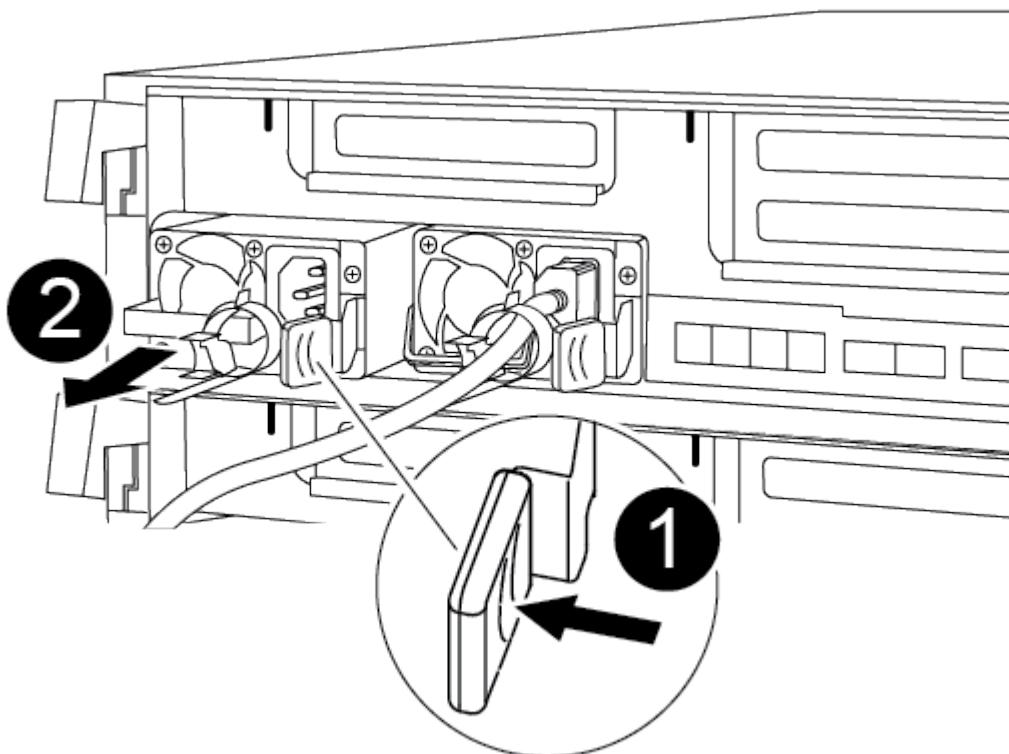
It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.



Do not mix PSUs with different efficiency ratings. Always replace like for like.

You can use the following animation, illustration, or the written steps to replace the power supply.

[Animation - Replace a power supply](#)



Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
3. Disconnect the power supply:
 - a. Open the power cable retainer, and then unplug the power cable from the power supply.

- b. Unplug the power cable from the power source.
4. Remove the power supply:
 - a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.
 - b. Press the blue locking tab to release the power supply from the chassis.
 - c. Using both hands, pull the power supply out of the chassis, and then set it aside.
5. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

The power supplies will only properly engage with the internal connector and lock in place one way.



To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

6. Rotate the cam handle so that it is flush against the power supply.

7. Reconnect the power supply cabling:

- a. Reconnect the power cable to the power supply and the power source.
- b. Secure the power cable to the power supply using the power cable retainer.

Once power is restored to the power supply, the status LED should be green.

8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Replace the real-time clock battery - FAS8300 and FAS8700

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

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

Step 1: Shut down the impaired controller

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

Option 1: Most configurations

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

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, **quorum status** of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

Steps

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

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:

- a. Enter the following command from the console of the healthy controller:

```
storage failover modify -node impaired_node_name -auto-giveback false
```

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.

If the impaired controller is displaying...	Then...
System prompt or password prompt	<p>Take over or halt the impaired controller from the healthy controller:</p> <pre>storage failover takeover -ofnode <i>impaired_node_name</i> -halt true</pre> <p>The <i>-halt true</i> parameter brings you to the LOADER prompt.</p>

Option 2: Controller is in a two-node MetroCluster

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

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller...	Then...
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: <code>metrocluster switchover</code>
Has not automatically switched over, you attempted switchover with the <code>metrocluster switchover</code> command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
  Operation: heal-aggregates
    State: successful
  Start Time: 7/25/2016 18:45:55
  End Time: 7/25/2016 18:45:56
  Errors: -
```

5. Check the state of the aggregates by using the storage aggregate show command.

```
controller_A_1::> storage aggregate show
  Aggregate      Size Available Used% State  #Vols  Nodes
  RAID  Status
  -----  -----
  ...
  aggr_b2      227.1GB    227.1GB    0% online      0  mcc1-a2
  raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the metrocluster heal -phase root-aggregates command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the metrocluster heal command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the metrocluster operation show command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
    State: successful
  Start Time: 7/29/2016 20:54:41
  End Time: 7/29/2016 20:54:42
  Errors: -
```

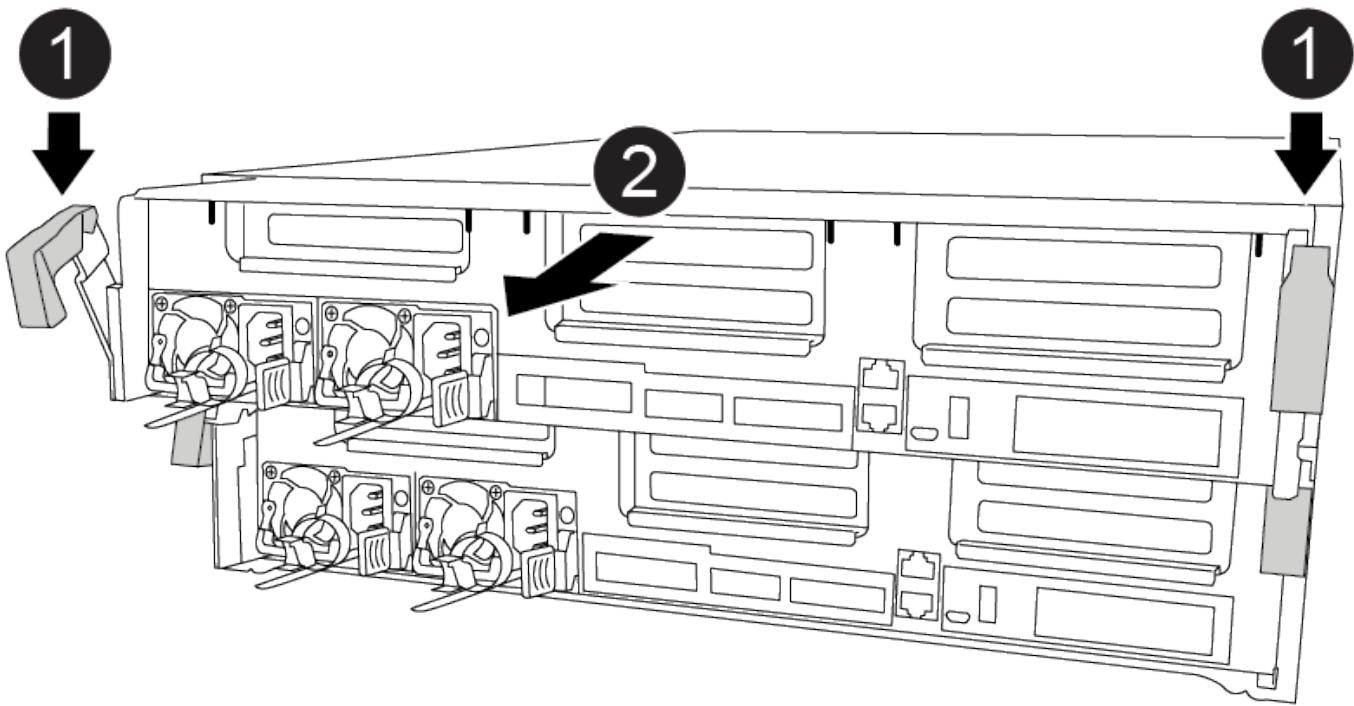
8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module



Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. 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, keeping 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.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

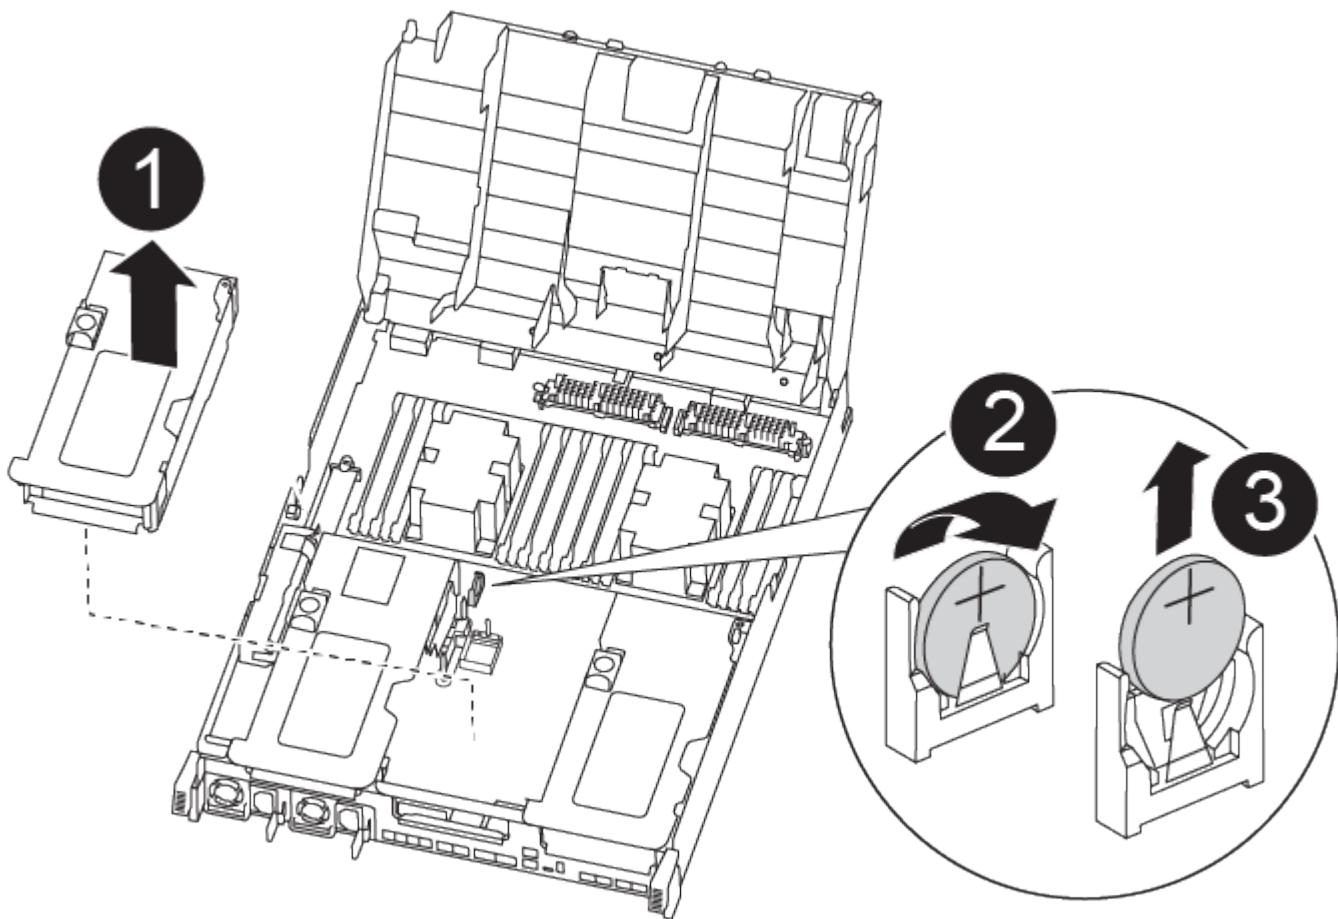
7. Place the controller module on a stable, flat surface.

Step 3: Replace the RTC battery

You need to locate the RTC battery inside the controller module, and then follow the specific sequence of steps. See the FRU map inside the controller module for the location of the RTC battery.

You can use the following animation, illustration, or the written steps to replace the RTC battery.

Animation - Replace the RTC battery



Steps

1. If you are not already grounded, properly ground yourself.
2. Open the air duct:
 - a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
 - b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
3. Locate, remove, and then replace the RTC battery:
 - a. Using the FRU map, locate the RTC battery on the controller module.
 - b. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.



Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

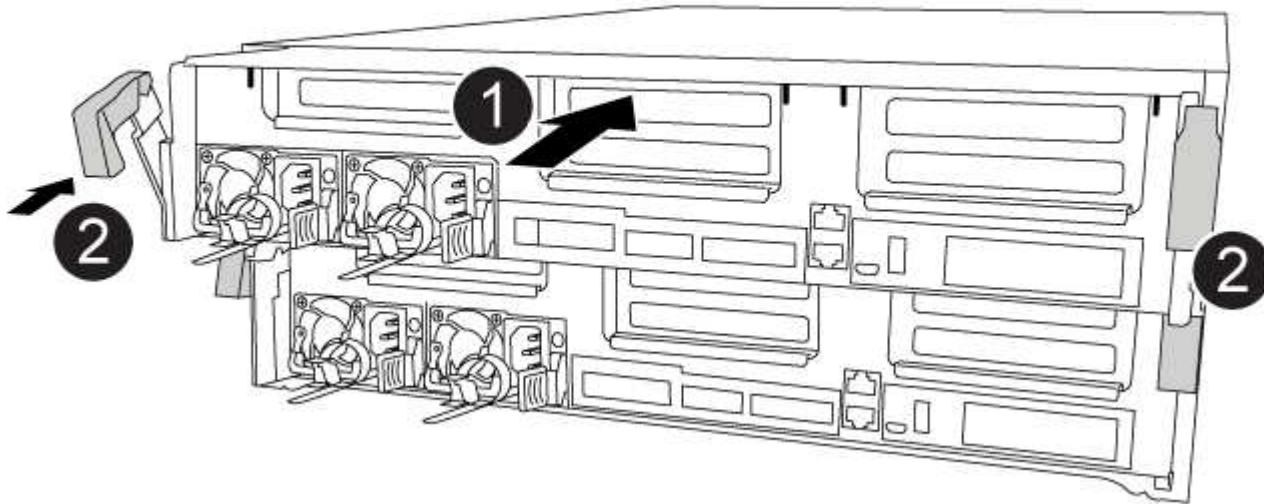
- c. Remove the replacement battery from the antistatic shipping bag.
- d. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
4. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
5. Close the air duct.

Step 4: Reinstall the controller module and sett time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

[Animation - Install the controller module](#)



Steps

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber

optic cables.

4. Complete the installation of the controller module:

- Using the locking latches, firmly push the controller module into the chassis until it meets the midplane and is fully seated.

The locking latches rise when the controller module is fully seated.



Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

- Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
- Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

- If you have not already done so, reinstall the cable management device.
- Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.



If your system stops at the boot menu, select the option to boot to LOADER.

5. Reset the time and date on the controller:

- Check the date and time on the healthy controller with the `show date` command.
- At the LOADER prompt on the target controller, check the time and date.
- If necessary, modify the date with the `set date mm/dd/yyyy` command.
- If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
- Confirm the date and time on the target controller.

6. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.

7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`

8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

- Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::> metrocluster node show

DR          Configuration DR
Group Cluster Node  State   Mirroring Mode
-----  -----  -----  -----
-----  -----
1   cluster_A
      controller_A_1 configured  enabled  heal roots
completed
      cluster_B
      controller_B_1 configured  enabled  waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster          Configuration State   Mode
-----  -----  -----  -----
Local: cluster_B configured      switchover
Remote: cluster_A configured    waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster          Configuration State   Mode
-----  -----  -----  -----
Local: cluster_B configured      normal
Remote: cluster_A configured    normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](#) page for further information.

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.