Maintain

Install and maintain

NetApp
February 13, 2026

Table of Contents

Maintain

Maintain AFF C800 hardware

Maintain the hardware of your AFF C800 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 AFF C800 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the AFF C800 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 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.
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.
Drive	A drive is a device that provides the physical storage media for data.
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 - AFF C800

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 partner node to reinstall ONTAP on the replacement boot media in your AFF C800 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 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 - AFF C800

Before replacing the boot media in your AFF C800, 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:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/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 - AFF C800

Shut down the impaired controller in your AFF C800 storage system to prevent data loss nd maintain system stability during the automated boot media recovery process.

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 `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* 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 - AFF C800

The boot media in your AFF C800 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.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

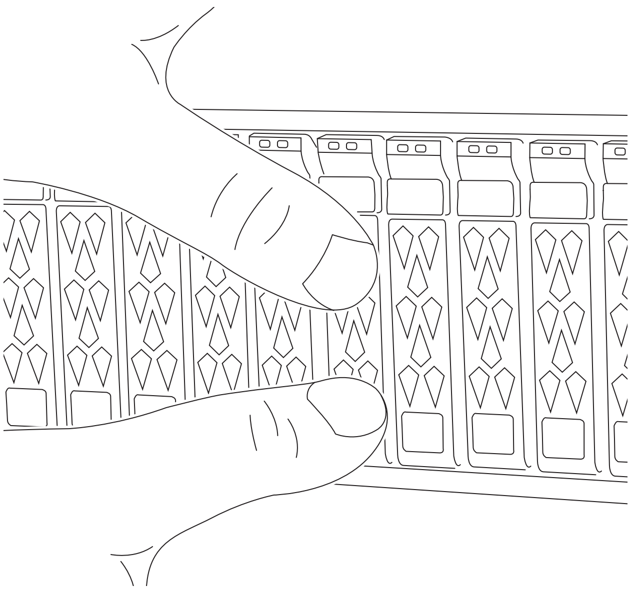

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if 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.

6. Remove the cable management device from the controller module and set it aside.
7. 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.

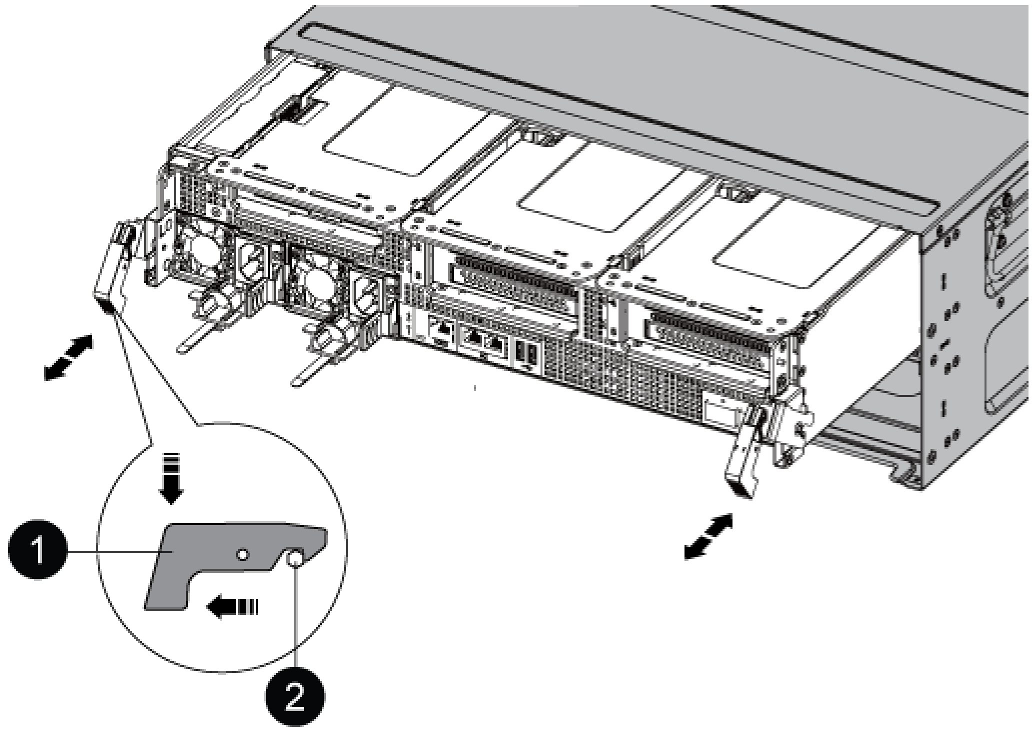

1	Locking latch
2	Locking pin

8. 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.

9. Place the controller module on a stable, flat surface, and then open the air duct:

   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.

   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

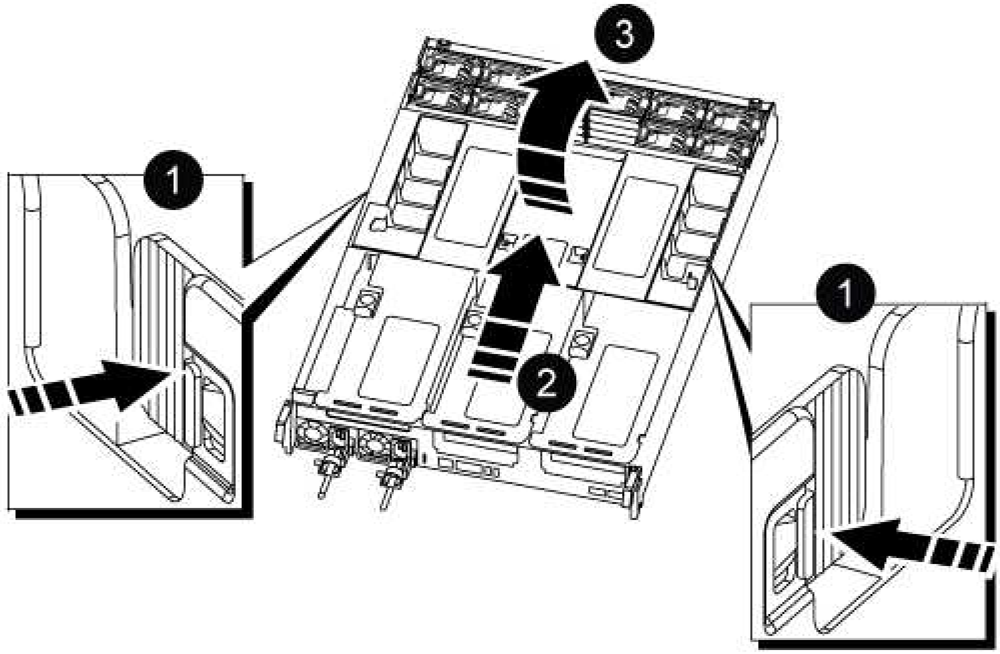

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

10. Locate the boot media in the controller module and replace it:

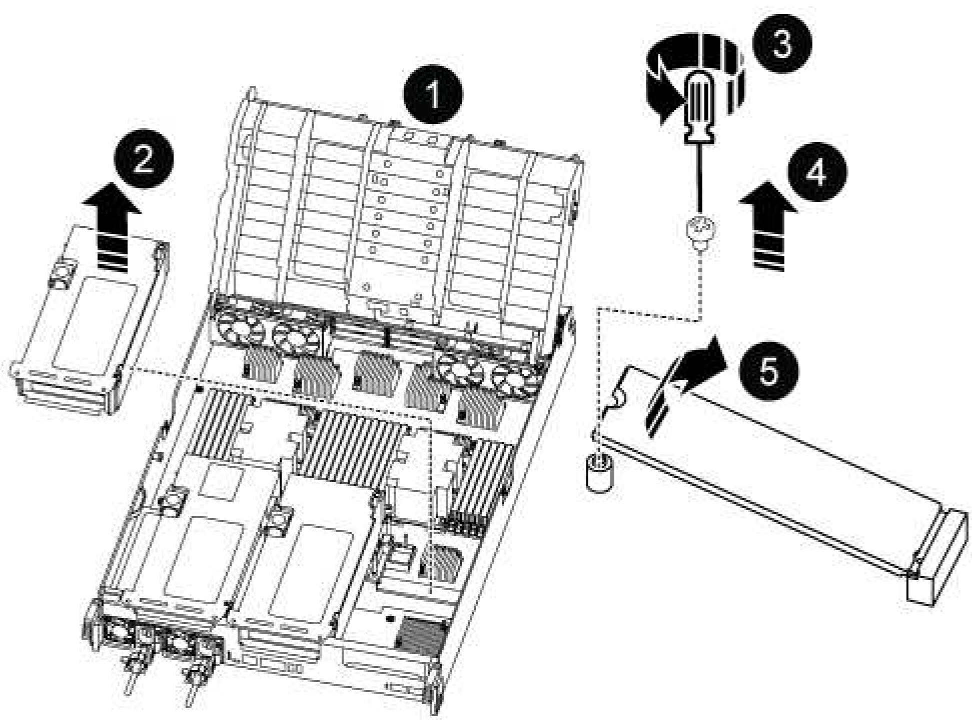

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.

b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.

11. Install the replacement boot media into the controller module:

    a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

    b. Rotate the boot media down toward the motherboard.

    c. Secure the boot media to the motherboard using the boot media screw.

       Do not over-tighten the screw or you might damage the boot media.

12. Reinstall the riser into the controller module.

13. Close the air duct:

    a. Rotate the air duct downward.

    b. Slide the air duct toward the risers until it clicks into place.

14. Install the controller module:

    a. Align the end of the controller module with the opening in the chassis, and then gently push the controller module half-way into the way into the system.

    b. Recable the controller module, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

       The controller module begins to boot and stops at the LOADER prompt.

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 - AFF C800

After installing the new boot media device in your AFF C800 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.`	Encryption is not installed on the system. a. Wait for the login prompt to display. b. Log into the node and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* c. 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.cr
t
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 - AFF C800

If a component in your AFF C800 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 - AFF C800

Get started with replacing the boot media in your AFF C800 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 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 ONATP 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 - AFF C800

Before replacing the boot media in your AFF C800 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 - AFF C800

To ensure data security on your AFF C800 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 `1Ono-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	`security key-manager keystore show` • 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	`security key-manager show-key-store` • 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:

a. 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.

b. Verify that all authentication keys are restored:

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

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

c. 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:

a. 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
```

b. 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 - AFF C800

Shut down the impaired controller in your AFF C800 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

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 `y` 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 `y`.

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: System is in a MetroCluster

> 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 `y` 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 `y`.

What's next?

After shutting down the controller, you need to replace the boot media.

Replace the boot media and prepare for manual boot recovery - AFF C800

The boot media in your AFF C800 system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media, and then manually transferring the ONTAP image to the replacement boot media using 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

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

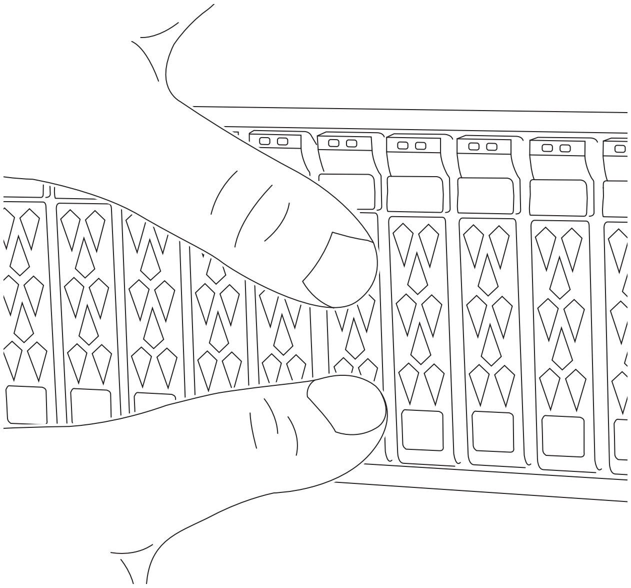

3. Unplug the controller module power supplies from the source.
4. Release the power cable retainers, and then unplug the cables from the power supplies.
5. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFP and QSFP modules (if 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.

6. Remove the cable management device from the controller module and set it aside.
7. 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.

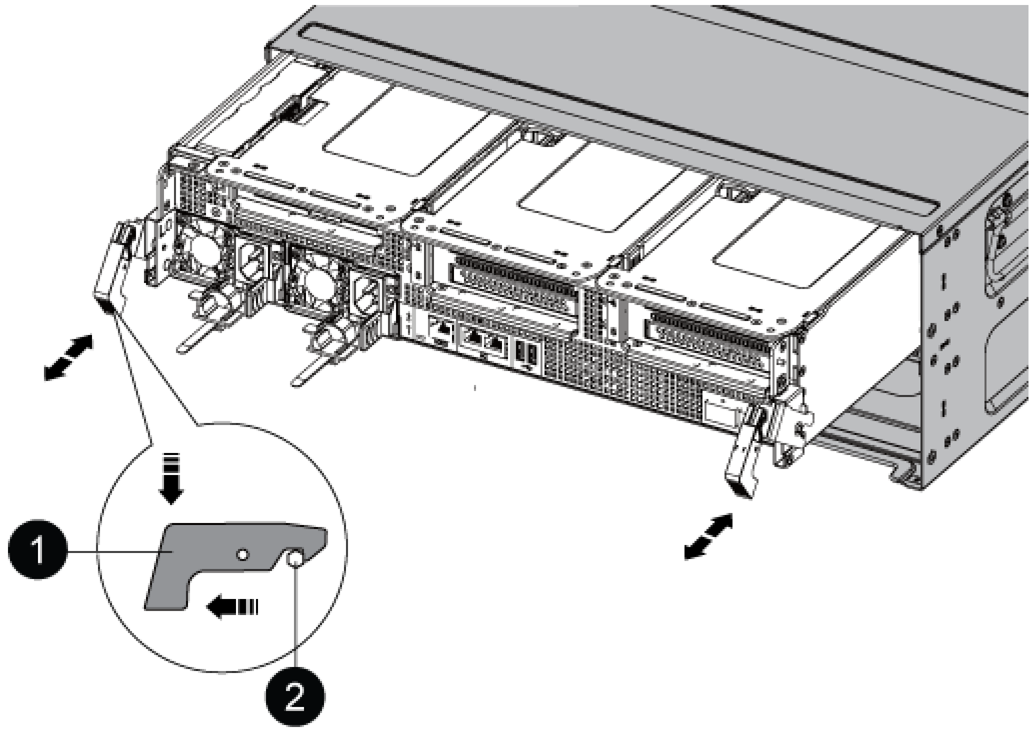

1	Locking latch
2	Locking pin

8. 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.

9. Place the controller module on a stable, flat surface, and then open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.

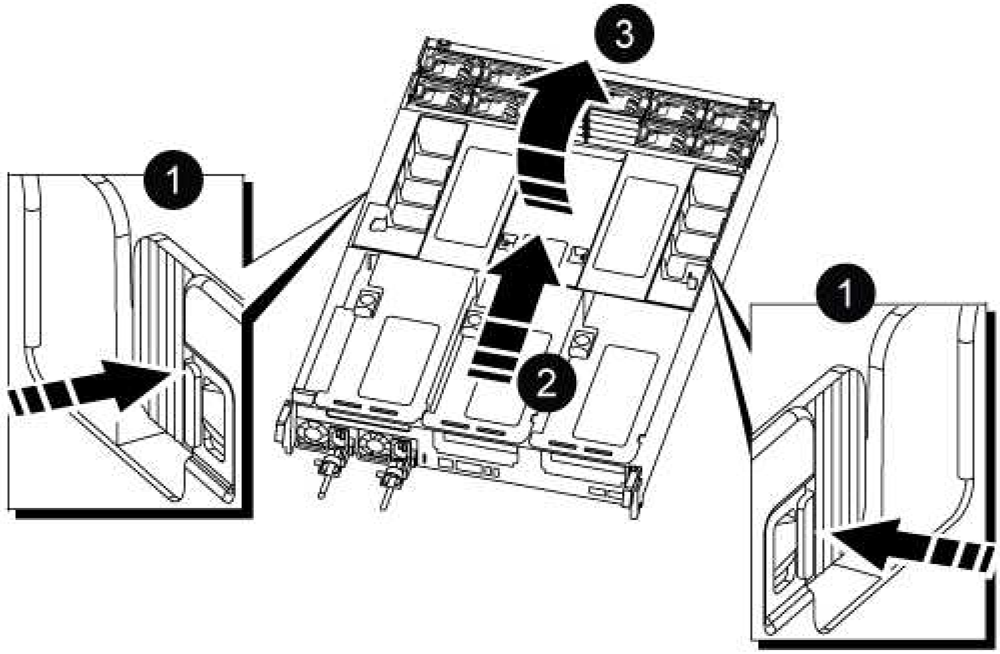

1	Air duct locking tabs
2	Slide air duct towards fan modules
3	Rotate air duct towards fan modules

Step 2: Replace the boot media

You locate the failed boot media in the controller module by removing Riser 3 on the controller module before you can replace the boot media.

You need a Phillips head screwdriver to remove the screw that holds the boot media in place.

1. Locate the boot media:

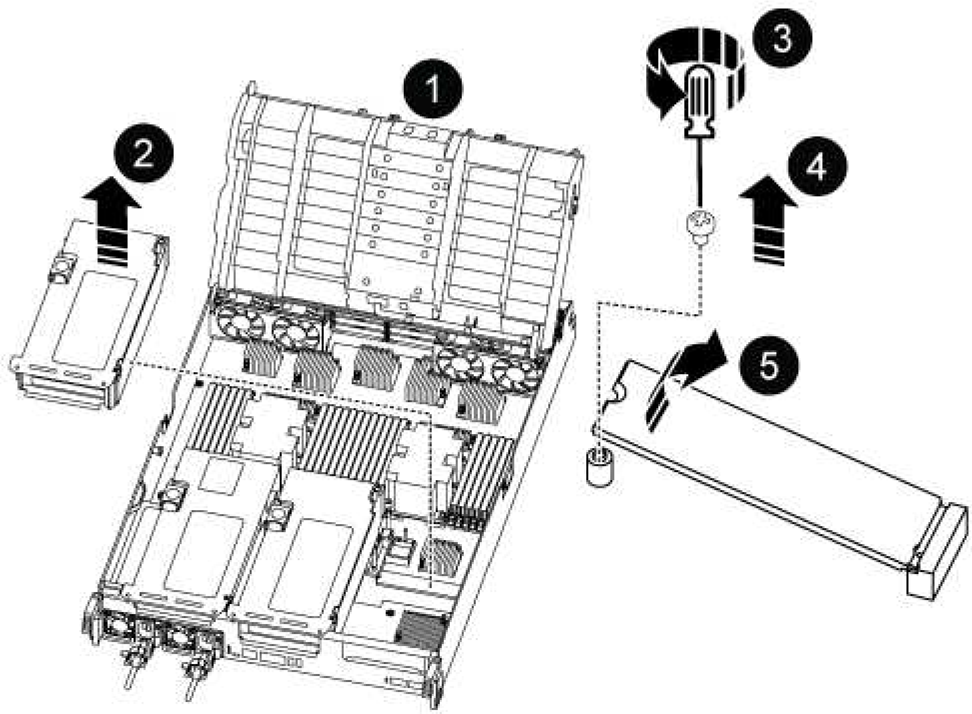

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

2. Remove the boot media from the controller module:
    a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.
    b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media straight out of the socket and set it aside.
3. Install the replacement boot media into the controller module:
    a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
    b. Rotate the boot media down toward the motherboard.
    c. Secure the boot media to the motherboard using the boot media screw.

        Do not over-tighten the screw or you might damage the boot media.
4. Reinstall the riser into the controller module.

5. Close the air duct:
    a. Rotate the air duct downward.
    b. Slide the air duct toward the risers until it clicks into place.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed is without 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 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.

Steps

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.

    a. Remove the USB flash drive from your laptop.
2. If you have not already done so, close the air duct:
    a. Swing the air duct all the way down to the controller module.
    b. Slide the air duct toward the risers until the locking tabs click into place.

c. Inspect the air duct to make sure that it is properly seated and locked into place.

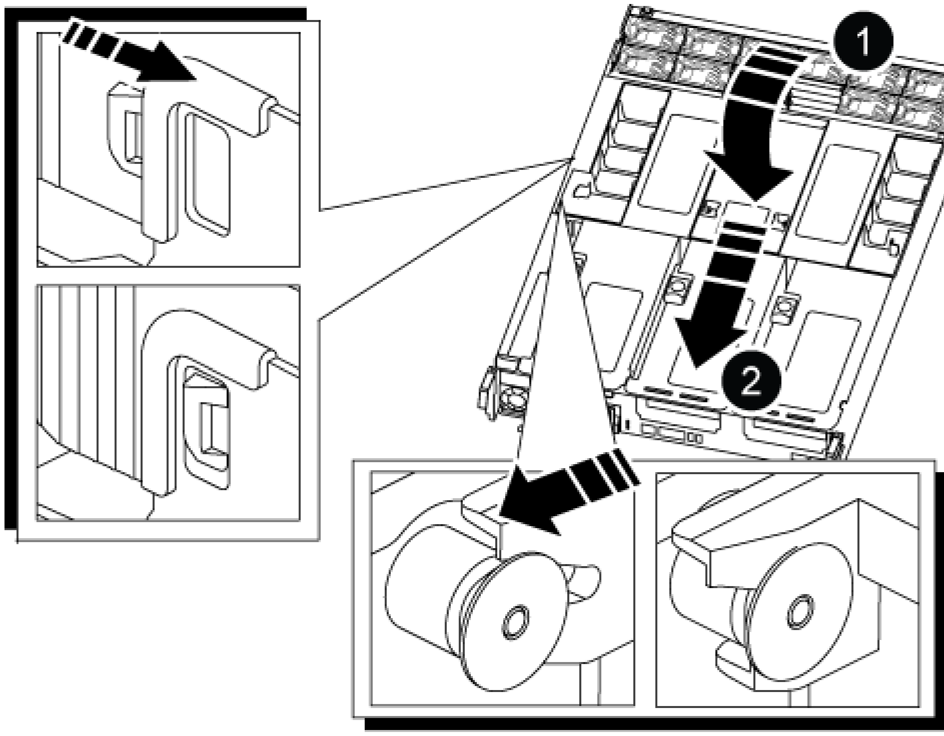

1	Air duct
2	Risers

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. Gently push the controller module all the way into the system until the controller module locking hooks begin to rise, firmly push on the locking hooks to finish seating the controller module, and then swing the locking hooks into the locked position over the pins on the controller module.
7. 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.

8. 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.

What’s next?

After replacing the boot media, you need to boot the recovery image.

Manual boot media recovery from a USB drive - AFF C800

After installing the new boot media device in your AFF C800 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

1. 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
```

What's next?

After booting the recovery image, you need to restore encryption on the boot media.

Restore encryption - AFF C800

Restore encryption on the replacement boot media in your AFF C800 system to ensure continued data protection. The replacement process involves verifying key availability, reapplying encryption settings, and confirming secure access to your data.

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:

<table>
<tr><th>ONTAP version</th><th>Select this option</th></tr>
<tr><td>ONTAP 9.8 or later</td><td>Select option 10.

Show example boot menu
<pre>Please choose one of the following:

(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)? 10</pre></td></tr>
</table>

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(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.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

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:

---------------------------BEGIN
BACKUP---------------------------
01234567890123456789012345678901234567890123456789012345678901
23
12345678901234567890123456789012345678901234567890123456789012
34
23456789012345678901234567890123456789012345678901234567890123
45
34567890123456789012345678901234567890123456789012345678901234
56
45678901234567890123456789012345678901234567890123456789012345
67
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
```

```
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
01234567890123456789012345678901234567890123456789012345678901
23
12345678901234567890123456789012345678901234567890123456789012
34
23456789012345678901234567890123456789012345678901234567890123
45
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
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`

What's next?

After restoring encryption on the boot media, you need to return the failed part to NetApp.

Return the failed boot media to NetApp - AFF C800

If a component in your AFF C800 storage system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Chassis

Chassis replacement workflow - AFF C800

Get started with replacing the chassis of your AFF C800 storage system by reviewing the replacement requirements, shutting down the controllers, replacing the chassis, and verifying system operations.

1 **Review the chassis replacement requirements**

Review the chassis replacement requirements, including system compatibility, required tools, ONTAP credentials, and component functionality verification.

2 **Prepare for the chassis replacement**

Prepare for the chassis replacement by locating the system, gathering credentials and tools, verifying the replacement chassis, and labeling cables.

3 **Shut down the controllers**

Shut down the controllers to perform chassis maintenance safely.

4 **Replace the chassis**

Move the components from the impaired chassis to the replacement chassis.

5 **Complete the chassis replacement**

Complete the replacement by booting the controllers, performing giveback, and returning the failed chassis to NetApp.

Requirements to replace the chassis - AFF C800

Before replacing the chassis in your AFF C800 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have local administrator credentials for ONTAP, the correct replacement chassis, and the necessary tools.

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Review the following requirements.

- Make sure all other components in the system are functioning properly; if not, contact NetApp support for assistance.
- Obtain local administrator credentials for ONTAP if you don’t have them.
- Make sure that you have the necessary tools and equipment for the replacement.
- You can use the chassis replacement procedure with all versions of ONTAP supported by your system.
- The chassis replacement procedure is written with the assumption that you are moving the bezel, NVMe drives, and controller modules to the new chassis, and that the replacement chassis is a new component from NetApp.

- The chassis replacement procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

What's next?

After reviewing the requirements, prepare to replace the chassis.

Prepare to replace the chassis - AFF C800

Prepare to replace the impaired chassis in your AFF C800 system by identifying the impaired chassis, verifying the replacement components, and labeling the cables and controller modules.

Steps

1. Connect to the serial console port to interface with and monitor the system.
2. Turn on the controller's Location LED:
   a. Use the `system controller location-led show` command to display the current state of the Location LED.
   b. Turn on the Location LED:

      ```
      system controller location-led modify -node node1 -state on
      ```

      The Location LED remains lit for 30 minutes.
3. Before opening the packaging, examine the packaging label and verify the following:
   - Component part number
   - Part description
   - Quantity in the box
4. Remove the contents from the packaging and save the packaging for returning the failed component to NetApp.
5. Label all cables connected to the storage system. This ensures proper recabling later in this procedure.
6. Ground yourself if not already grounded.

What's next?

After you've prepared to replace your AFF C800 chassis hardware, you need to shut down the controllers.

Shut down the controllers - AFF C800

Shut down the controllers in your AFF C800 storage system to prevent data loss and ensure system stability when replacing the chassis.

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 accessability 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.

What’s next?

After shutting down the controllers, replace the chassis.

Replace the chassis - AFF C800

Replace the chassis in your AFF C800 system when a hardware failure requires it. The replacement process involves removing the controllers, moving the drives to the replacement chassis, removing the impaired chassis, installing the replacement chassis, and reinstalling the chassis components.

Step 1: Remove the controller modules from the old chassis

Remove the controller modules from the old chassis.

Steps

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

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 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.

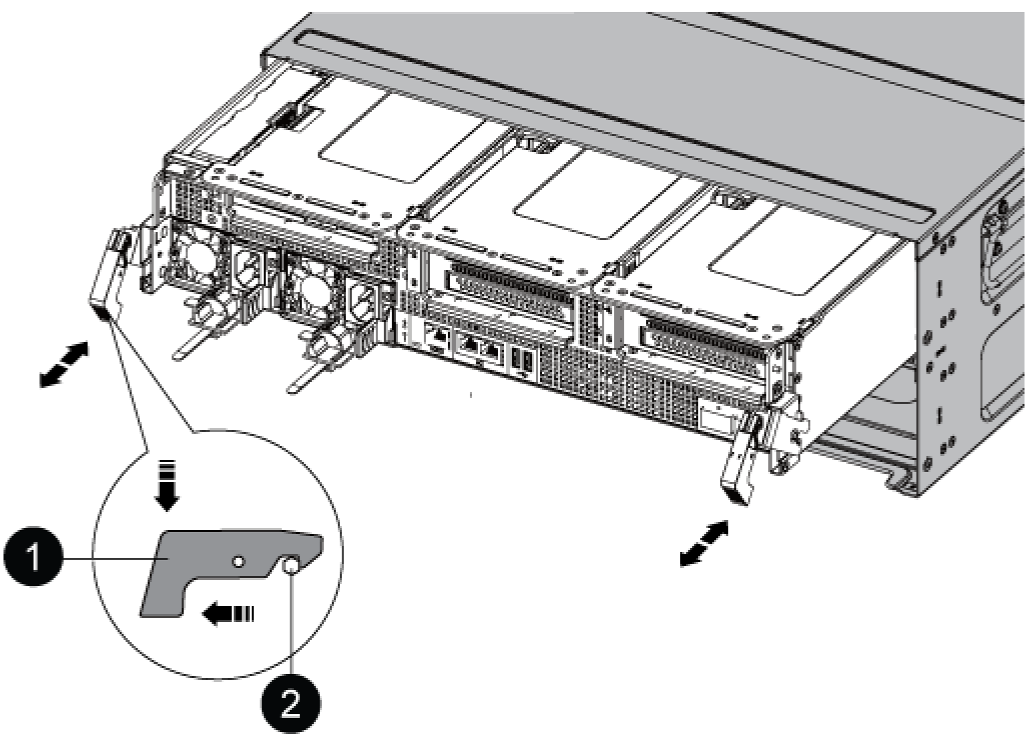

1	Locking latch
2	Locking pin

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 drives to the new chassis

Move the drives from the old chassis to the new chassis.

Steps

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button at the top of the carrier face below the LEDs.
   b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

      The drive should disengage from the chassis, allowing it to slide free of the chassis.

      > When removing a drive, always use two hands to support its weight.

      > Drives are fragile. Handle them as little as possible to prevent damage to them.

3. Align the drive from the old chassis with the same bay opening in the new chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate upward.

5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it up and against the drive holder.

   Be sure to close the cam handle slowly so that it aligns correctly with the front of the drive carrier. It clicks when it is secure.

6. Repeat the process for the remaining drives in the system.

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

Replace the impaired chassis in the equipment rack or system cabinet with the new 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 into the new chassis

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 reinstallation of the controller module:
   a. 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. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them 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 by pressing `Ctrl-C`.
4. Repeat the preceding steps to install the second controller into the new chassis.

What's next?

After you have replaced the impaired AFF C800 chassis and reinstalled the components, you need to complete the chassis replacement

Complete the restoration and replacement process - AFF C800

Reboot the controllers, verify system health, and return the failed part to NetApp to complete the final step in the AFF C800 chassis replacement procedure.

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

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`
      - `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.
4. Reinstall the bezel on the front of the system.

Step 2: Verify storage system health

After the controller giveback completes, verify system health using Active IQ Config Advisor. Address any issues found.

Step 3: 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.

Controller

Controller replacement workflow - AFF C800

Replace the controller in your AFF C800 storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.

1 **Review the requirements to replace the controller**

Review the controller replacement requirements, including system compatibility, required tools, ONTAP credentials, and component functionality verification.

2

Shut down the impaired controller

Shut down or take over the impaired controller so that the healthy controller continues to serve data from the impaired controller storage.

3

Replace the controller

Remove the impaired controller, move the FRU components to the replacement controller module, and install the replacement controller module in the enclosure.

4

Restore and verify the system configuration

Verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

5

Recable and give back the controller

Recable the controller and transfer the ownership of storage resources back to the replacement controller.

6

Complete controller replacement

Verify the LIFs, check cluster health, and return the failed part to NetApp.

Requirements to replace the controller - AFF C800

Before replacing the controller of your AFF C800 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements for replacing the controller module.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller being replaced (referred to in this procedure as the impaired controller).
- Do not use this procedure for controller upgrades. Refer to Choose your controller hardware upgrade procedure for guidance.
- If your system is in a MetroCluster configuration, review Choosing the correct recovery procedure to determine whether to use this procedure.
- Replace the failed component with the field-replaceable unit (FRU) you received from NetApp.
- Replace the controller module with a controller module of the same model type. You cannot upgrade your system by replacing the controller module.
- You cannot change drives or drive shelves as part of this procedure.

- The boot device is located on the System Management module installed in the back of the system. You do not need to move the boot device when replacing a controller module.
- Understand the controller terminology used in this procedure:
  - The *impaired* controller is the controller being replaced.
  - The *replacement* controller is the new controller replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- Capture the controller's console output to a text log file.

  This provides a record of the procedure to troubleshoot any issues during the replacement process.

What's next?

After you've reviewed the requirements to replace your AFF C800 controller, you need to shut down the impaired controller.

Shut down the impaired controller - AFF C800

Shut down the controller in your AFF C800 storage system to prevent data loss and ensure system stability when replacing the controller.

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 `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

What's next?

After you've shut down the controller, you need to replace the controller.

Replace the controller module hardware - AFF C800

Replace the controller in your AFF C800 system when a hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

Step 1: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

   Video - Confirm drive seating

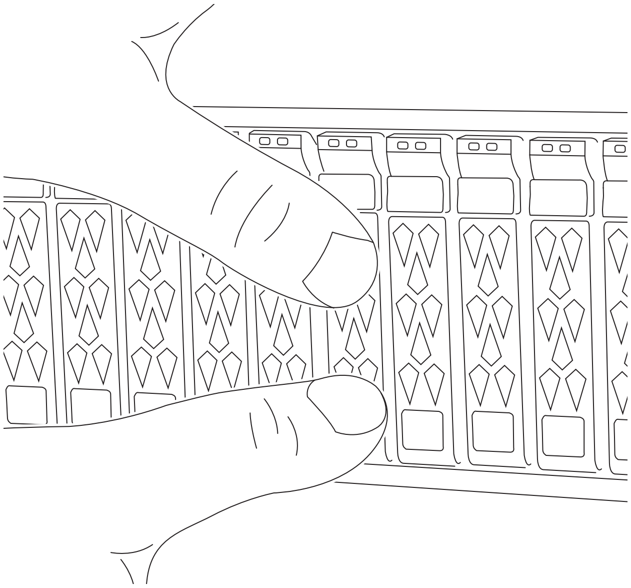

3. Check the controller drives based on the system status:

   a. On the healthy controller, check if any active RAID group is in a degraded state, failed state, or both:

   ```
   storage aggregate show -raidstatus !*normal*
   ```

   ▪ If the command returns `There are no entries matching your query.` continue to go to the next sub-step to check for missing drives.

   ▪ If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

   ```
   system node autosupport invoke -node * -type all -message
   '<message_name>'
   ```

   b. Check for missing drive issues for both the file system or spare drives:

   ```
   event log show -severity * -node * -message-name *disk.missing*
   ```

   ▪ If the command returns `There are no entries matching your query.` continue to go to the next step.

   ▪ If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

   ```
   system node autosupport invoke -node * -type all -message
   '<message_name>'
   ```

4. Remove the power cable retainers, then unplug the cables from the power supplies.

5. Loosen the hook and loop strap on the cable management device. Unplug the system cables and SFP/QSFP modules (if needed) from the controller module. Note each cable’s location.

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

6. Remove the cable management device from the controller module and set it aside.
7. 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.

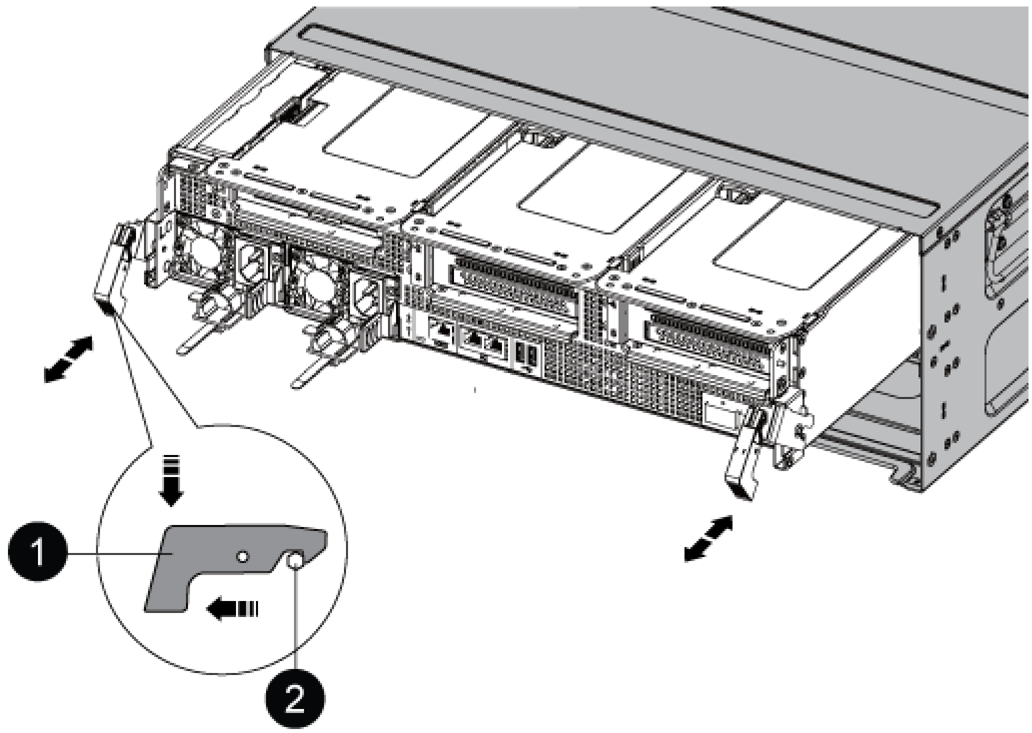

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a stable, flat surface.

   Support the bottom of the controller module while sliding it out of the chassis.

Step 2: Move the power supplies

Move the power supplies to the replacement controller module.

Steps

1. Rotate the cam handle such that it can be used to pull power supply out of the controller module while pressing the locking tab.

The power supply is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

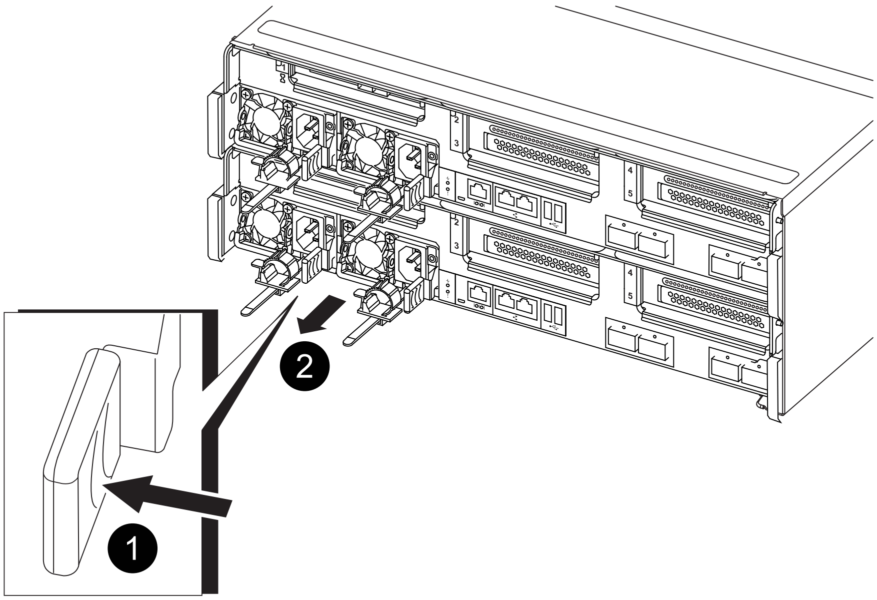

1	Blue power supply locking tab
2	Power supply

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.

Step 3: Move the fans

Move the fan modules to the replacement controller module.

Steps

1. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

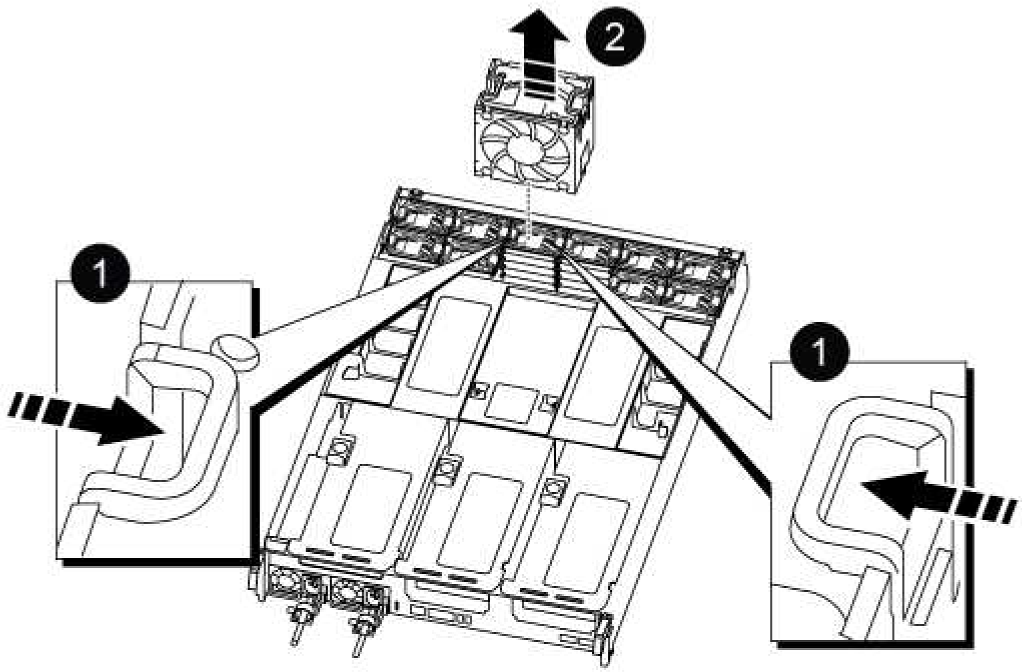

1	Fan locking tabs
2	Fan module

2. Move the fan module to the replacement controller module, and then install the fan module by aligning its edges with the opening in the controller module, and then sliding the fan module into the controller module until the locking latches click into place.
3. Repeat these steps for the remaining fan modules.

Step 4: Move the NVDIMM battery

Move the NVDIMM battery to the replacement controller module.

Steps

1. Open the air duct cover and locate the NVDIMM battery in the riser.

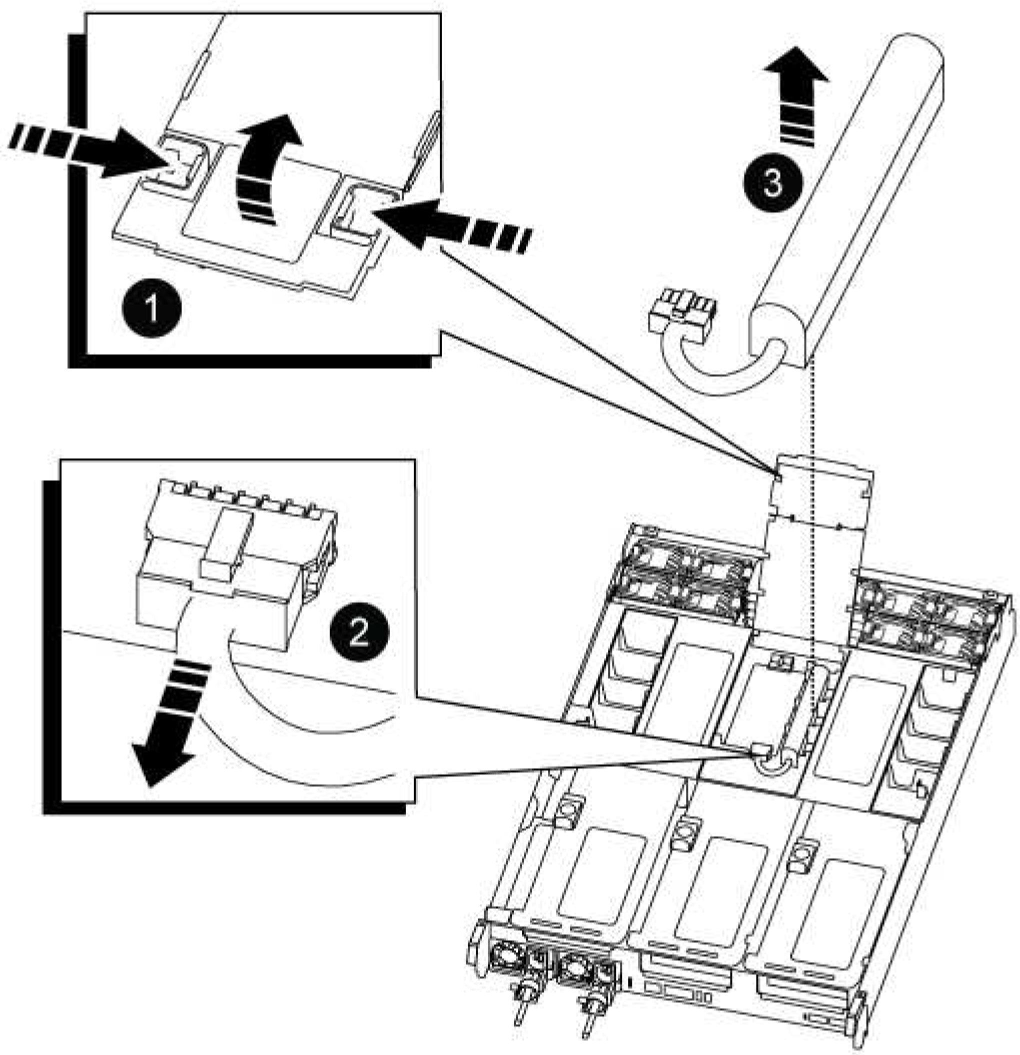

1	Air duct riser
2	NVDIMM battery plug
3	NVDIMM battery pack

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

2. 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.
3. Grasp the battery and lift the battery out of the air duct and controller module.
4. Move the battery pack to the replacement controller module and then install it in the NVDIMM air duct:
    a. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
    b. Plug the battery plug into the riser socket and make sure that the plug locks into place.

Step 5: Remove the PCIe risers

Remove the PCIe risers from the impaired controller module.

Steps

1. Remove the PCIe riser from the 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 the fan modules.

      The riser raises up slightly from the controller module.
   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

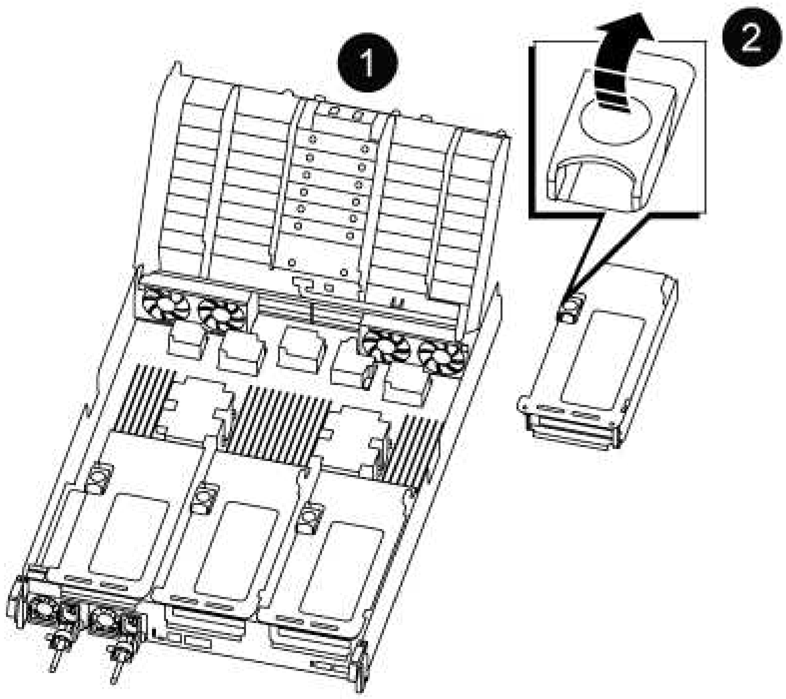

1	Air duct
2	Riser 1 (left riser), Riser 2 (middle riser), and 3 (right riser) locking latches

2. Repeat the preceding step for the remaining risers in the impaired controller module.
3. Repeat the above steps with the empty risers in the replacement controller and put them away.

Step 6: Move system DIMMs

Move the system DIMMS to the replacement controller module.

Steps

1. 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.

2. Eject the DIMM from its slot by slowly pushing apart the two 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.

3. Locate the slot where you are installing the DIMM.
4. Insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

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

5. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
6. Repeat these steps for the remaining DIMMs.

Step 7: Move the NVDIMMs

Move the NVDIMMS to the replacement controller module.

Steps

1. Locate the NVDIMMs on your controller module.

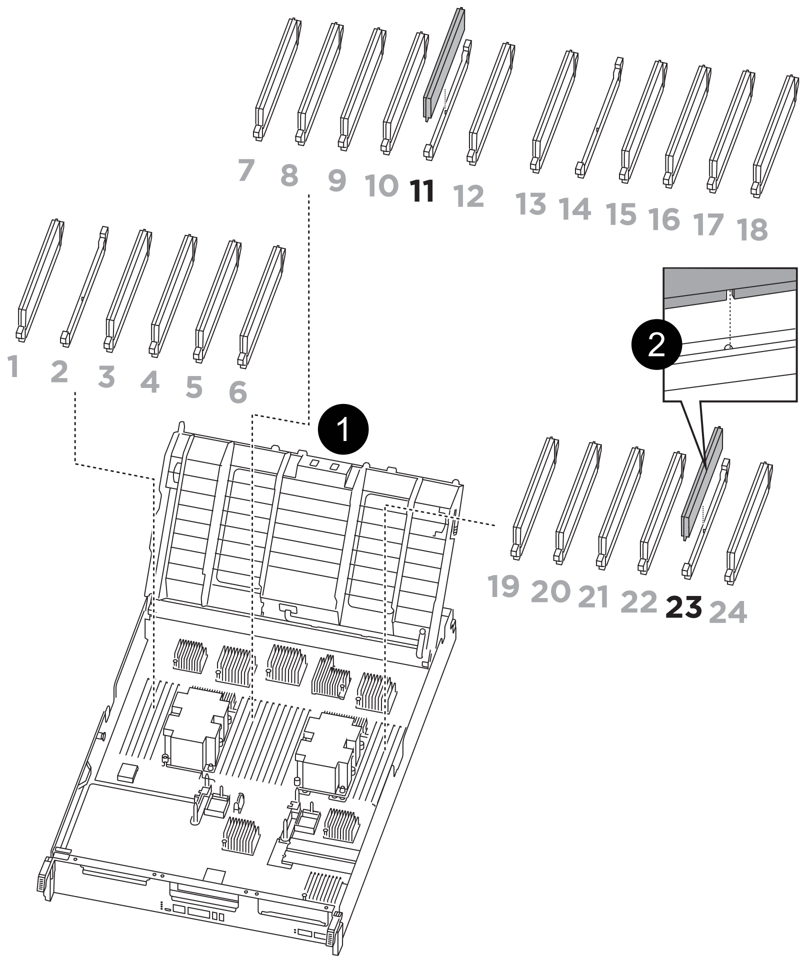

- NVDIMM: SLOTS 11 & 23

1	Air duct
2	NVDIMMs

2. Note the orientation of the NVDIMM in the socket so that you can insert the NVDIMM in the replacement controller module in the proper orientation.
3. 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.

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. Repeat the preceding steps to move the other NVDIMM.

Step 8: Move the boot media

Move the boot media to the replacement controller module.

Steps

1. Locate the boot media under Riser 3.

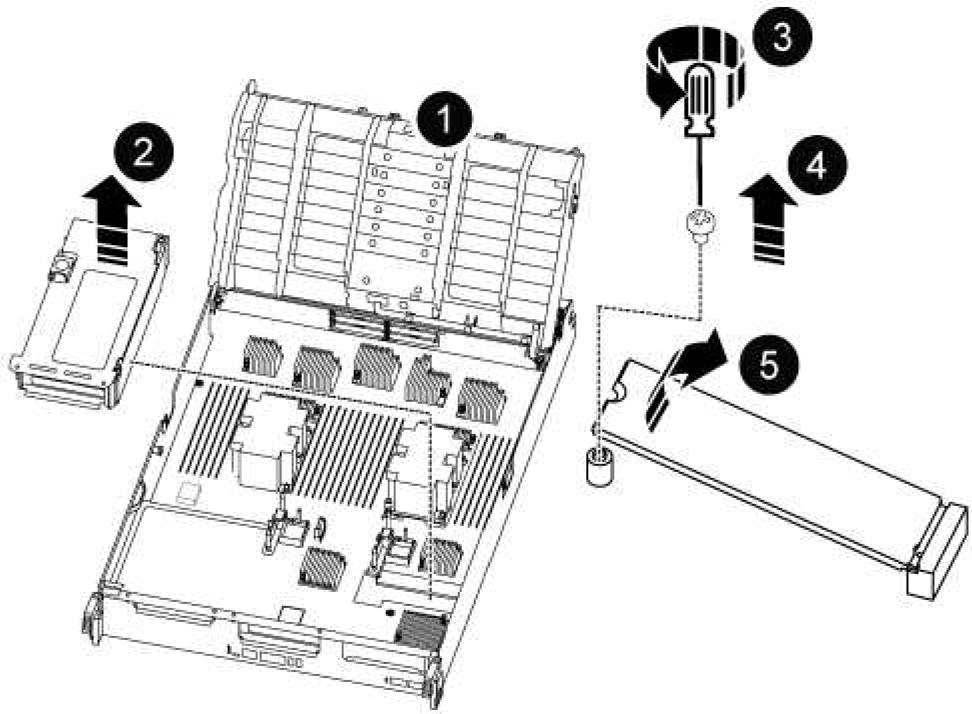

1	Air duct
2	Riser 3
3	Phillips #1 screwdriver
4	Boot media screw
5	Boot media

2. Remove the boot media from the controller module:
   a. Using a #1 Phillips head screwdriver, remove the screw holding down the boot media and set the screw aside in a safe place.
   b. Grasping the sides of the boot media, gently rotate the boot media up, and then pull the boot media

straight out of the socket and set it aside.

3. Move the boot media to the new controller module and install it:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the motherboard.
   c. Secure the boot media to the motherboard using the boot media screw.

      Do not over-tighten the screw or you might damage the boot media.

Step 9: Install the PCIe risers

Install the risers in the replacement controller module.

Steps

1. Install the riser into the replacement controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

   d. Reinsert any SFP or QSFP modules that were removed from the PCIe cards.
2. Repeat the preceding step for the remaining PCIe risers.

Step 10: Install the controller module

Reinstall the controller module and reboot it.

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 system, as needed.

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

3. Complete the reinstallation of the controller module:
   a. 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. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them 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.

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

   `storage failover giveback -ofnode impaired_node_name`.

5. If automatic giveback was disabled, reenable it:

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

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

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

What's next?

After you've replaced the impaired AFF C800 controller, you need to restore the system configuration

Restore and verify the system configuration - AFF C800

Return control of storage resources to the replacement controller so your AFF C800 system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption or Onboard Key Manager (OKM) encryption.

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

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 chassis

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`

   ◦ `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`

What's next?

After you've transferred the ownership of storage resources back to the replacement controller, you need to complete the controller replacement procedure.

Recable the system and reassign disks - AFF C800

Return control of storage resources to the replacement controller so your AFF C800 system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption or Onboard Key Manager (OKM) encryption.

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

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

Steps

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:`boot_ontap`
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 content 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
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The *replacement* controller is the current owner of the disks on the disaster site.

     Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

10. If your system is in a MetroCluster configuration, verify that each controller is configured: `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 that the expected volumes are present for each controller: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - AFF C800

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: 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 - AFF C800

Replace a DIMM in your AFF C800 system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

Before you begin

- Make sure all other components in the system are functioning properly; if not, you must contact technical support.
- Make sure you replace the failed component with a replacement component you received from NetApp.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller

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 `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

   Video - Confirm drive seating

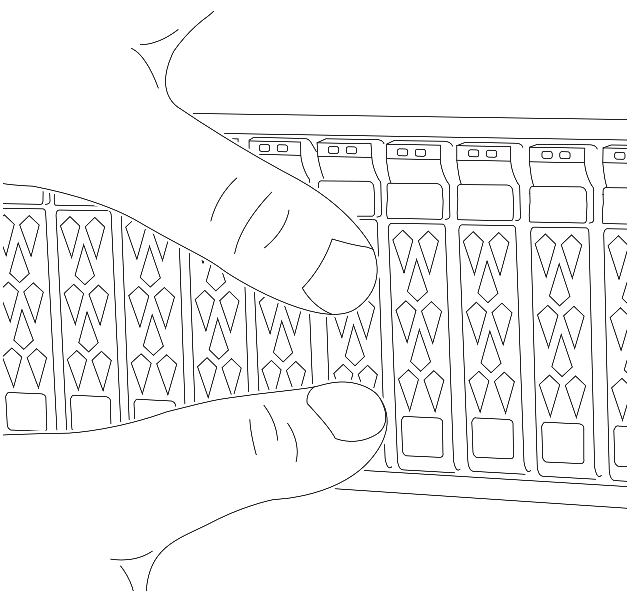

3. Check the controller drives based on the system status:
   a. On the healthy controller, check if any active RAID group is in a degraded state, failed state, or both:

   ```
   storage aggregate show -raidstatus !*normal*
   ```

   - If the command returns `There are no entries matching your query.` continue to go to the next sub-step to check for missing drives.
   - If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message
'<message_name>'
```

b. Check for missing drive issues for both the file system or spare drives:

```
event log show -severity * -node * -message-name *disk.missing*
```

- If the command returns `There are no entries matching your query.` continue to go to the next step.
- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message
'<message_name>'
```

4. Remove the power cable retainers, then unplug the cables from the power supplies.
5. Loosen the hook and loop strap on the cable management device. Unplug the system cables and SFP/QSFP modules (if needed) from the controller module. Note each cable's location.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
6. Remove the cable management device from the controller module and set it aside.
7. 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.

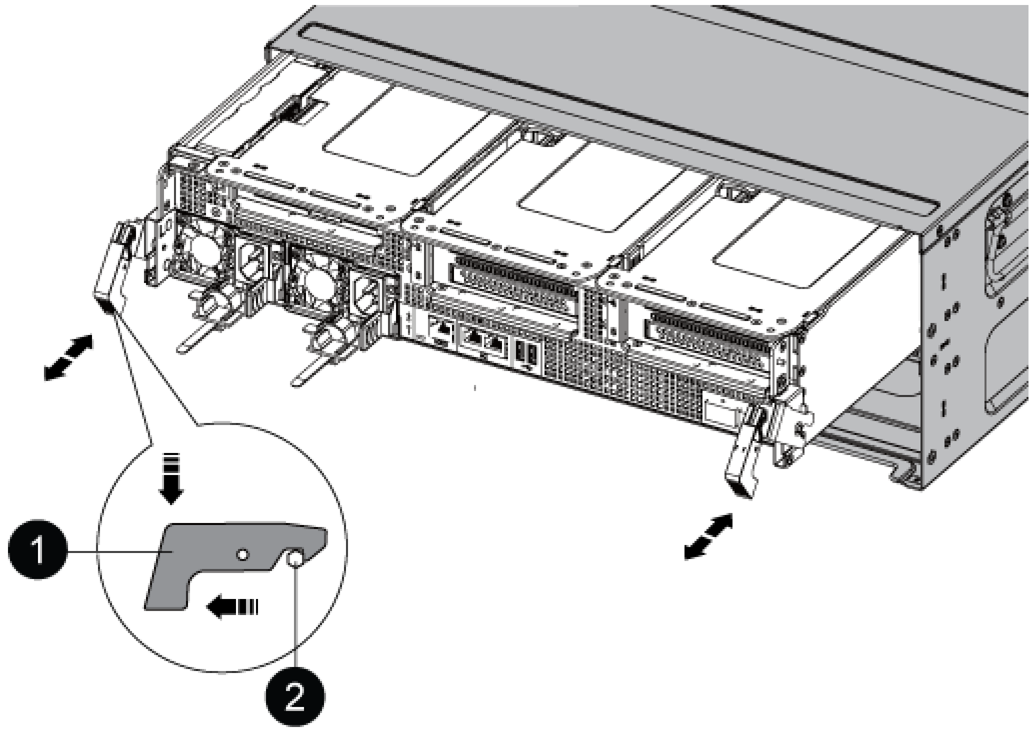

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a stable, flat surface.

   Support the bottom of the controller module while sliding it out of the chassis.

Step 3: Replace the DIMM

Replace the DIMM in the controller.

To replace a DIMM, you must locate it in the controller module using the DIMM map label on top of the air duct and then replace it following the specific sequence of steps.

1. Open the air duct:
   a. Press in the locking tabs on the sides of the air duct toward the middle of the controller module.
   b. Slide the air duct toward the fan modules, and then rotate it upward to its completely open position.
2. When removing a DIMM, unlock the locking latch on the applicable riser, and then remove the riser.

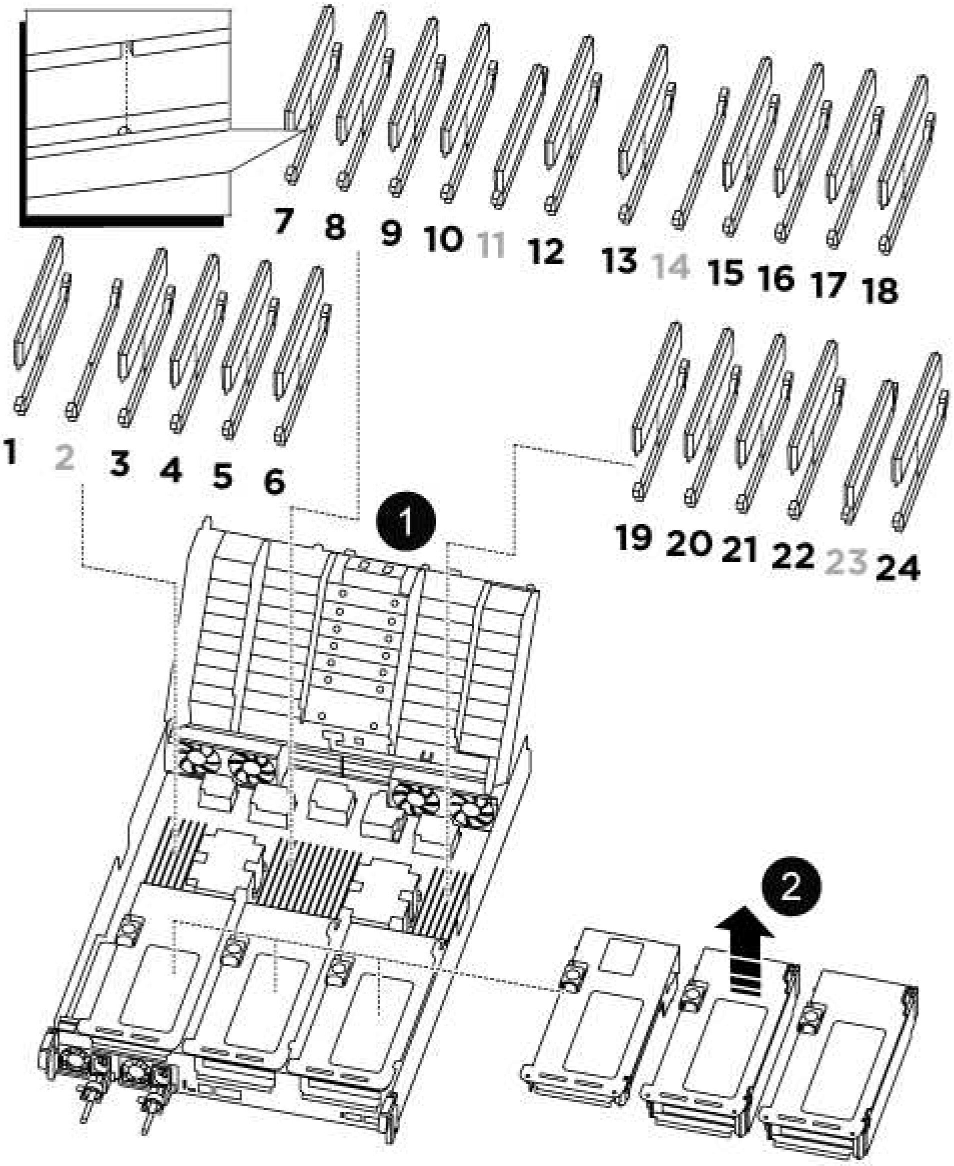

1	Air duct cover
2	Riser 1 and DIMM bank 1, and 3-6
Riser 2 and DIMM bank 7-10, 12-13, and 15-18	Riser 3 and DIMM 19 -22 and 24

Note: Slot 2 and 14 are left empty. Do not attempt to install DIMMs into these slots.

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 slot by slowly pushing apart the two 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.

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. 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. Reinstall any risers that you removed from the controller module.
9. Close the air duct.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

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 system, as needed.

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

3. Complete the reinstallation of the controller module:
   a. 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. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them 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.

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

   `storage failover giveback -ofnode impaired_node_name`.

5. If automatic giveback was disabled, reenable it:

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

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

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

Step 5: 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 SSD Drive or HDD Drive - AFF C800

Replace a drive in your AFF C800 system when a drive fails or requires an upgrade. The replacement process involves identifying the faulty drive, safely removing it, and installing a new drive to ensure continued data access and system performance.

You can replace a failed drive nondisruptively while I/O is in progress. The procedure for replacing an SSD is meant for non-spinning drives and the procedure for replacing an HDD is meant for spinning drives.

When a drive fails, the platform logs a warning message to the system console indicating which drive has failed. In addition, both the fault LED on the operator display panel and the fault LED on the failed drive are illuminated.

Before you begin

- Follow best practice and install the current version of the Disk Qualification Package (DQP) before replacing a drive.
- Identify the failed drive by running the `storage disk show -broken` command from the system console.

  The failed drive appears in the list of failed drives. If it does not, you should wait, and then run the command again.

  > Depending on the type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

- Determine whether SED authentication is enabled.

  How you replace the drive depends on how the drive is being used. If SED authentication is enabled, you must use the SED replacement instructions in the ONTAP 9 NetApp Encryption Power Guide. These Instructions describe additional steps you must perform before and after replacing an SED.

- Make sure the replacement drive is supported by your platform. See the NetApp Hardware Universe.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

About this task

- Drive firmware is automatically updated (nondisruptively) on new drives that have non current firmware versions.
- When replacing a drive, you must wait one minute between the removal of the failed drive and the insertion of the replacement drive to allow the storage system to recognize the existence of the new drive.

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.

    a. Verify whether automatic drive assignment is enabled: `storage disk option show`

       You can enter the command on either controller module.

       If automatic drive assignment is enabled, the output shows `on` in the "Auto Assign" column (for each controller module).

    b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

       You must disable automatic drive assignment on both controller modules.

2. Properly ground yourself.
3. Physically identify the failed drive.

   When a drive fails, the system logs a warning message to the system console indicating which drive failed. Additionally, the attention (amber) LED on the drive shelf operator display panel and the failed drive illuminate.

   > The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.

4. Remove the failed drive:

    a. Press the release button on the drive face to open the cam handle.

    b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.

5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.

6. Insert the replacement drive:

    a. With the cam handle in the open position, use both hands to insert the replacement drive.

    b. Push until the drive stops.

    c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

       Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.

7. Verify that the drive's activity (green) LED is illuminated.

   When the drive's activity LED is solid, it means that the drive has power. When the drive's activity LED is blinking, it means that the drive has power and I/O is in progress. If the drive firmware is automatically updating, the LED blinks.

8. If you are replacing another drive, repeat the preceding steps.
9. If you disabled automatic drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
   a. Display all unowned drives:

      ```
      storage disk show -container-type unassigned
      ```

      You can enter the command on either controller module.

   b. Assign each drive:

      ```
      storage disk assign -disk disk_name -owner node_name
      ```

      You can enter the command on either controller module.

      You can use the wildcard character to assign more than one drive at once.

   c. Reenable automatic drive assignment if needed:

      ```
      storage disk option modify -node node_name -autoassign on
      ```

      You must reenable automatic drive assignment on both controller modules.

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

    Contact NetApp Support if you need the RMA number or additional help with the replacement procedure.

Option 2: Replace HDD

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment replacement drive, if it is enabled

   > You manually assign drive ownership and then reenable automatic drive assignment later in this procedure.

   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

      If automatic drive assignment is enabled, the output shows `on` in the “Auto Assign” column (for each controller module).

   b. If automatic drive assignment is enabled, disable it: `storage disk option modify -node node_name -autoassign off`

      You must disable automatic drive assignment on both controller modules.

2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

5. Press the release button on the disk drive face.

   Depending on the storage system, the disk drives have the release button located at the top or on the left of the disk drive face.

   For example, the following illustration shows a disk drive with the release button located on the top of the disk drive face:

   The cam handle on the disk drive springs open partially and the disk drive releases from the midplane.

6. Pull the cam handle to its fully open position to unseat the disk drive from the midplane.

7. Slide out the disk drive slightly and allow the disk to safely spin down, which can take less than one minute, and then, using both hands, remove the disk drive from the disk shelf.
8. With the cam handle in the open position, insert the replacement disk drive into the drive bay, firmly pushing until the disk drive stops.

   > Wait a minimum of 10 seconds before inserting a new disk drive. This allows the system to recognize that a disk drive was removed.

   > If your platform drive bays are not fully loaded with drives, it is important to place the replacement drive into the same drive bay from which you removed the failed drive.

   > Use two hands when inserting the disk drive, but do not place hands on the disk drive boards that are exposed on the underside of the disk carrier.

9. Close the cam handle so that the disk drive is fully seated into the midplane and the handle clicks into place.

   Be sure to close the cam handle slowly so that it aligns correctly with the face of the disk drive..

10. If you are replacing another disk drive, repeat Steps 4 through 9.
11. Reinstall the bezel.
12. If you disabled automatice drive assignment in Step 1, then, manually assign drive ownership and then reenable automatic drive assignment if needed.
    a. Display all unowned drives: `storage disk show -container-type unassigned`

       You can enter the command on either controller module.

    b. Assign each drive: `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller module.

       You can use the wildcard character to assign more than one drive at once.

    c. Reenable automatic drive assignment if needed: `storage disk option modify -node node_name -autoassign on`

You must reenable automatic drive assignment on both controller modules.

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

    Contact technical support at NetApp Support, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Replace a fan - AFF C800

Replace a fan module in your AFF C800 system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the impaired controller, removing it, replacing the fan module, reinstalling the controller, and then returning the failed part to NetApp.

Step 1: Shut down the impaired controller

Shut down the impaired controller

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 `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

   Video - Confirm drive seating

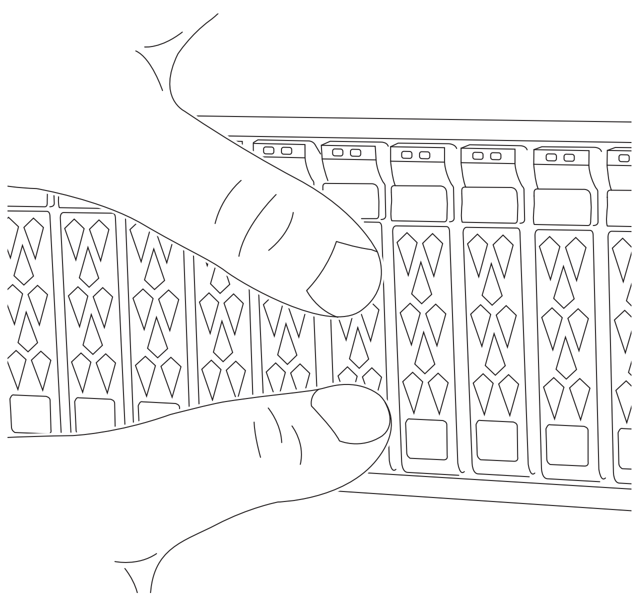

3. Check the controller drives based on the system status:
   a. On the healthy controller, check if any active RAID group is in a degraded state, failed state, or both:

```
storage aggregate show -raidstatus !*normal*
```

- If the command returns `There are no entries matching your query.` continue to go to the next sub-step to check for missing drives.
- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message
'<message_name>'
```

b. Check for missing drive issues for both the file system or spare drives:

```
event log show -severity * -node * -message-name *disk.missing*
```

- If the command returns `There are no entries matching your query.` continue to go to the next step.
- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message
'<message_name>'
```

4. Remove the power cable retainers, then unplug the cables from the power supplies.
5. Loosen the hook and loop strap on the cable management device. Unplug the system cables and SFP/QSFP modules (if needed) from the controller module. Note each cable's location.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
6. Remove the cable management device from the controller module and set it aside.
7. 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.

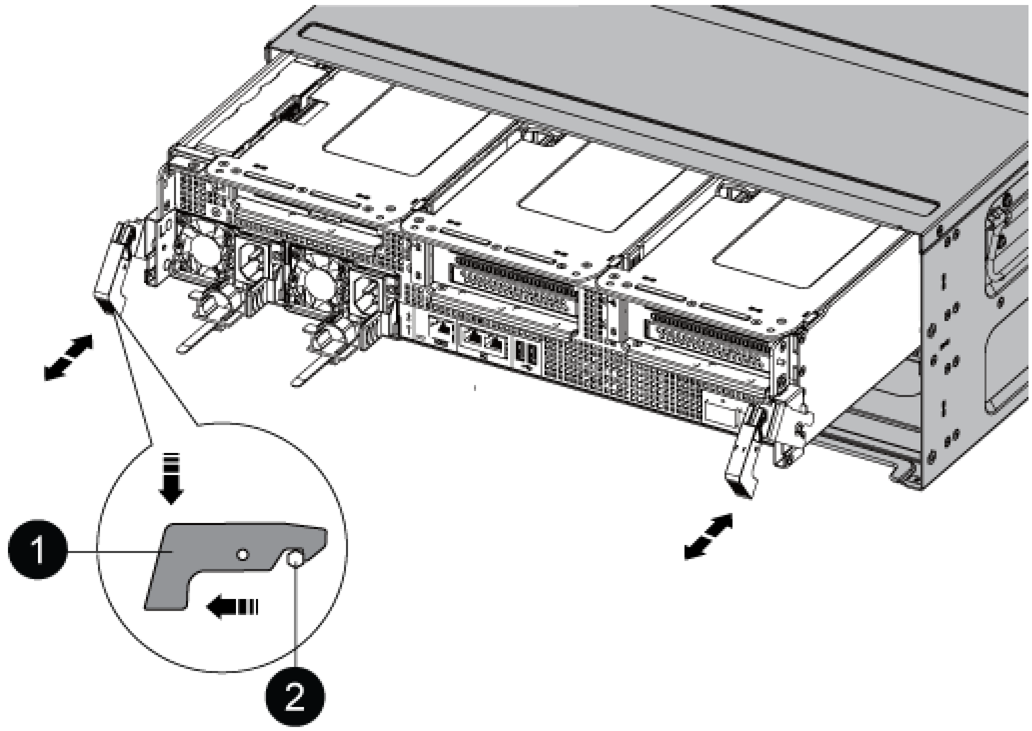

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a stable, flat surface.

   Support the bottom of the controller module while sliding it out of the chassis.

Step 3: Replace a fan

Locate the failed fan module and replace it with a new fan module.

Steps

1. Identify the fan module that you must replace by checking the console error messages or by locating the lit LED for the fan module on the motherboard.
2. Remove the fan module by pinching the locking tabs on the side of the fan module, and then lifting the fan module straight out of the controller module.

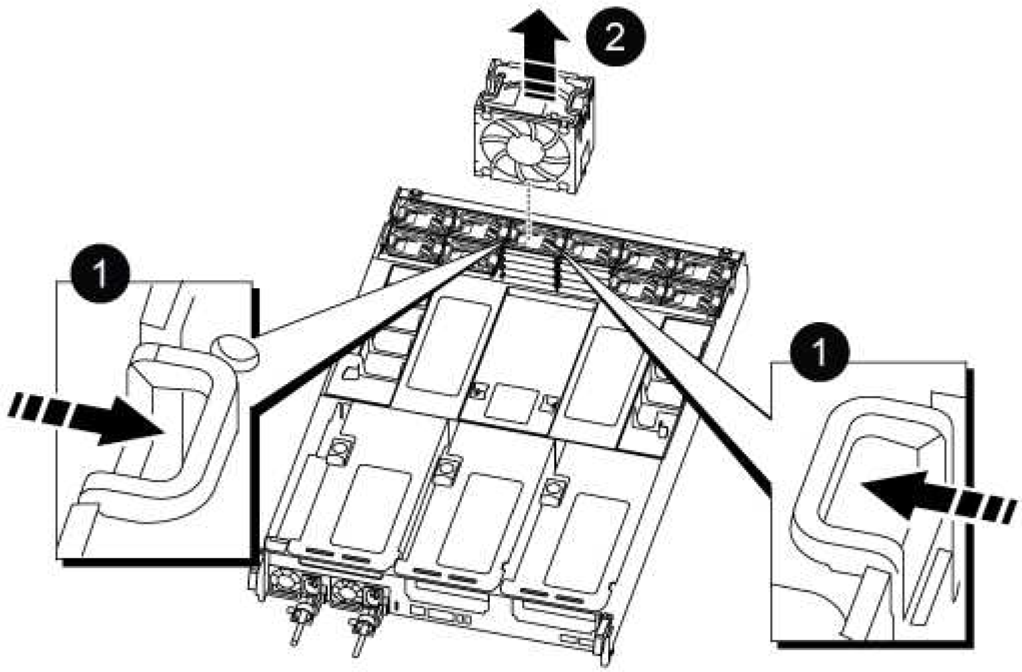

1	Fan locking tabs
2	Fan module

3. Align the edges of the replacement fan module with the opening in the controller module, and then slide the replacement fan module into the controller module until the locking latches click into place.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

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 system, as needed.

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

3. Complete the reinstallation of the controller module:

   a. 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. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them 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.

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

   `storage failover giveback -ofnode impaired_node_name`.

5. If automatic giveback was disabled, reenable it:

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

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

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

Step 5: 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 an NVDIMM - AFF C800

Replace the NVDIMM in your AFF C800 system 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.

Before you begin

- Make sure you have the replacement NVDIMM that you received from NetApp and is compatible with your AFF C800 system.
- Make sure all other components in the storage system are functioning properly; if not, contact NetApp Support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

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 `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

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

2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

   Video - Confirm drive seating

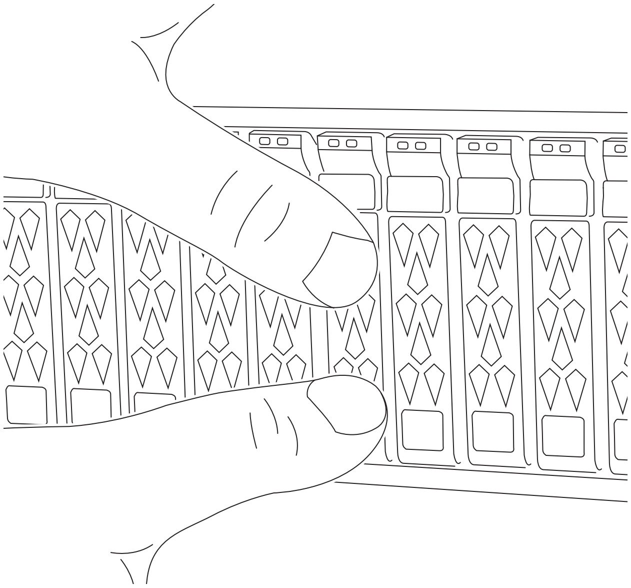

3. Check the controller drives based on the system status:

   a. On the healthy controller, check if any active RAID group is in a degraded state, failed state, or both:

```
storage aggregate show -raidstatus !*normal*
```

   - If the command returns `There are no entries matching your query.` continue to go to the next sub-step to check for missing drives.
   - If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message
'<message_name>'
```

   b. Check for missing drive issues for both the file system or spare drives:

```
event log show -severity * -node * -message-name *disk.missing*
```

   - If the command returns `There are no entries matching your query.` continue to go to the next step.
   - If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message
'<message_name>'
```

4. Remove the power cable retainers, then unplug the cables from the power supplies.

5. Loosen the hook and loop strap on the cable management device. Unplug the system cables and SFP/QSFP modules (if needed) from the controller module. Note each cable’s location.

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

6. Remove the cable management device from the controller module and set it aside.
7. 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.

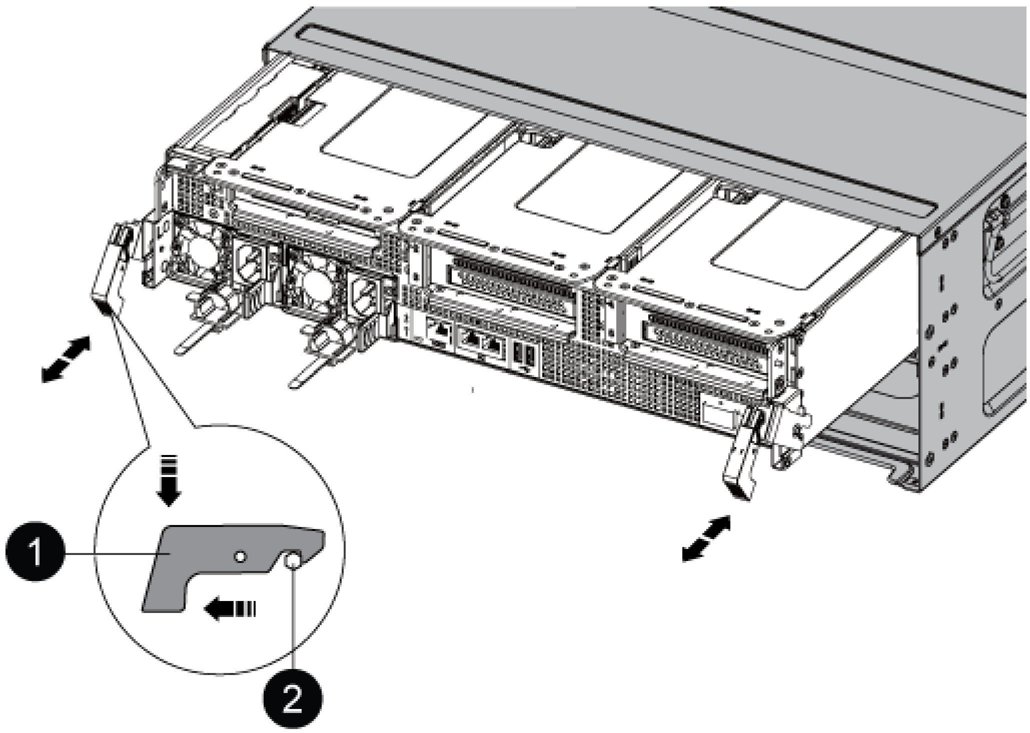

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a stable, flat surface.

   Support the bottom of the controller module while sliding it out of the chassis.

Step 3: Replace the NVDIMM

Locate the NVDIMM in the controller module using the NVDIMM map label on top of the air duct, and then replace it.

Steps

1. Access the NVDIMM by unlocking the locking latch on the appropriate riser, and then remove the riser.

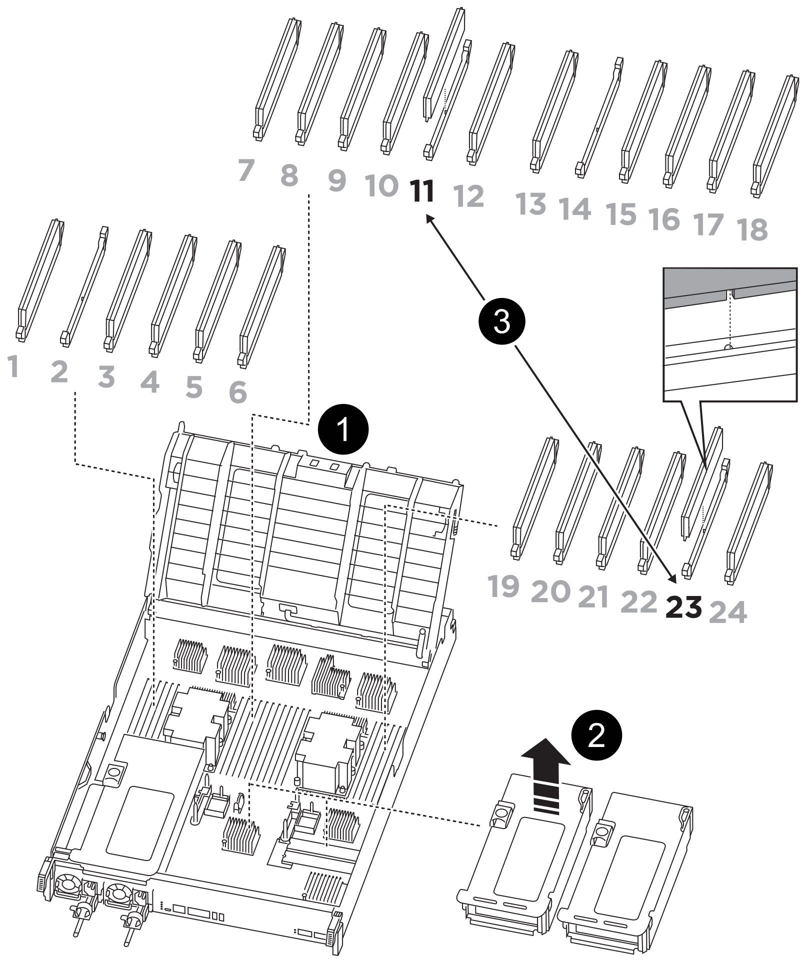

1	Air duct cover
2	Riser 2
3	NVDIMM in slots 11 and 23

2. Note the orientation of the NVDIMM in the socket so that you can insert the NVDIMM in the replacement controller module in the proper orientation.
3. 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.

4. 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.

5. Locate the slot where you are installing the NVDIMM.

6. 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.

7. 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.
8. Reinstall any risers that you removed from the controller module.
9. Close the air duct.

Step 4: Reinstall the controller module and booting the system

Reinstall the controller module and reboot it.

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 system, as needed.

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

3. Complete the reinstallation of the controller module:
   a. 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. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them 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.
4. Return the impaired controller to normal operation by giving back its storage:

   `storage failover giveback -ofnode` *`impaired_node_name`*.

5. If automatic giveback was disabled, reenable it:

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

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

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

Step 5: 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 - AFF C800

Replace the NV battery in your AFF C800 system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

Step 1: Shutdown the impaired controller

Shut down or take over the impaired controller.

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 `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

   Video - Confirm drive seating

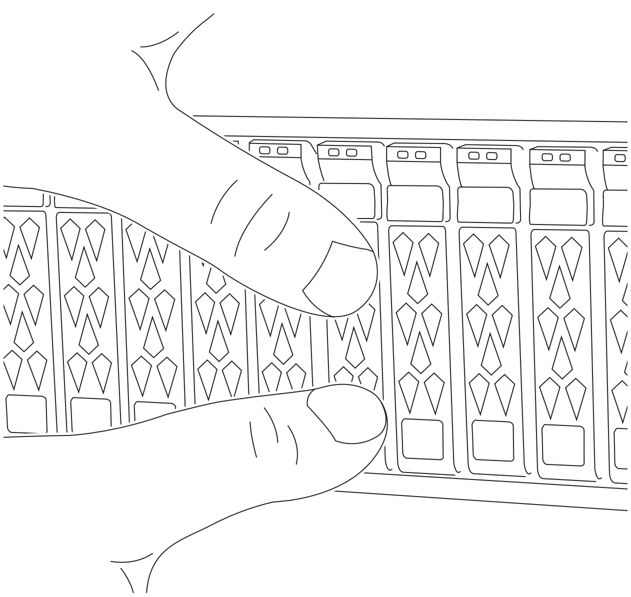

3. Check the controller drives based on the system status:
   a. On the healthy controller, check if any active RAID group is in a degraded state, failed state, or both:

```
storage aggregate show -raidstatus !*normal*
```

- If the command returns `There are no entries matching your query.` continue to go to the next sub-step to check for missing drives.
- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message
'<message_name>'
```

b. Check for missing drive issues for both the file system or spare drives:

```
event log show -severity * -node * -message-name *disk.missing*
```

- If the command returns `There are no entries matching your query.` continue to go to the next step.
- If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

```
system node autosupport invoke -node * -type all -message
'<message_name>'
```

4. Remove the power cable retainers, then unplug the cables from the power supplies.
5. Loosen the hook and loop strap on the cable management device. Unplug the system cables and SFP/QSFP modules (if needed) from the controller module. Note each cable’s location.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
6. Remove the cable management device from the controller module and set it aside.
7. 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.

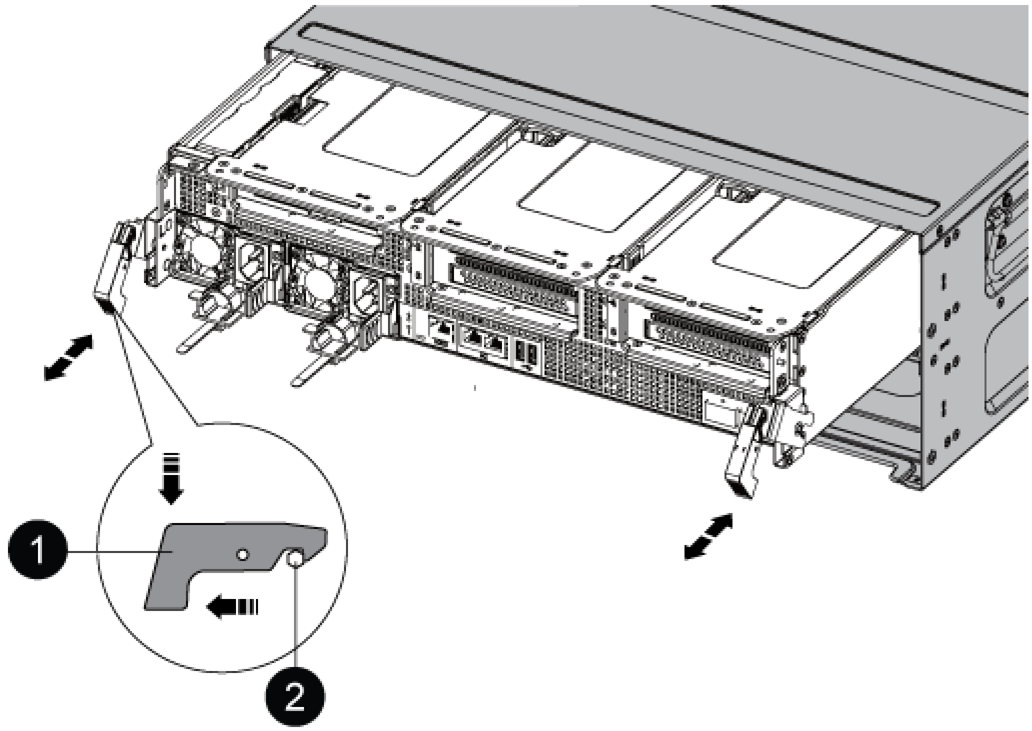

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a stable, flat surface.

   Support the bottom of the controller module while sliding it out of the chassis.

Step 3: Replace the NVDIMM battery

Replace the NVDIMM battery by removing the failed battery from the controller module and installing the replacement battery into the controller module.

Steps

1. Open the air duct cover and locate the NVDIMM battery in the riser.

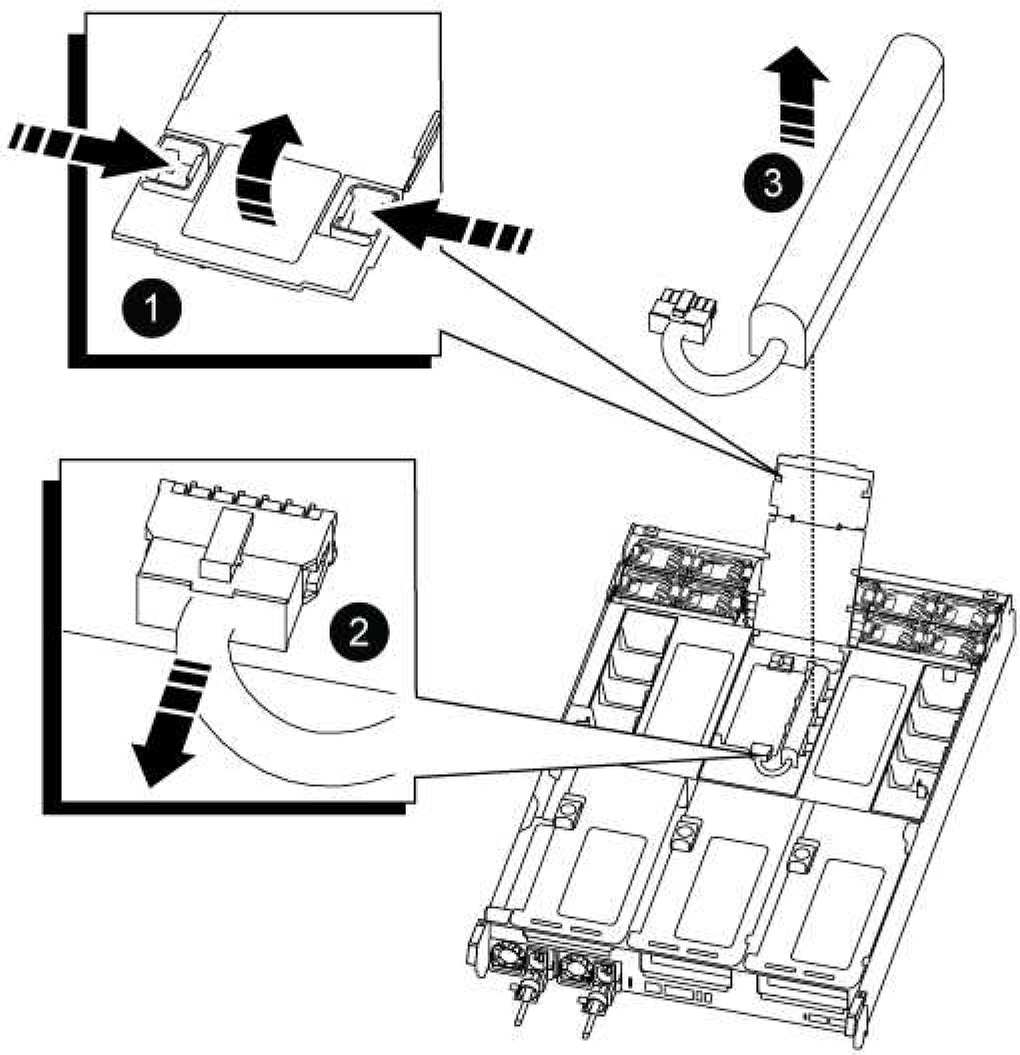

1	Air duct riser
2	NVDIMM battery plug
3	NVDIMM battery pack

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

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 lift the battery out of the air duct and controller module, and then set it aside.
3. Remove the replacement battery from its package.
4. Install the replacement battery pack in the NVDIMM air duct:
   a. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.

   b. Plug the battery plug into the riser socket and make sure that the plug locks into place.

5. Close the NVDIMM air duct.

   Make sure that the plug locks into the socket.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

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 system, as needed.

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

3. Complete the reinstallation of the controller module:

   a. 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. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them 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.

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

   `storage failover giveback -ofnode` *`impaired_node_name`*.

5. If automatic giveback was disabled, reenable it:

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

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

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

Step 5: 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 card - AFF C800

Replace or add an I/O module in your AFF C800 system when the module fails, requires an upgrade to support higher performance, or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

Before you begin

- You must have the NetApp new or replacement part available.
- Make sure all other components in the storage system are functioning properly; if not, contact technical support.
- 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

Shut down or take over the impaired controller.

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 `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

   Video - Confirm drive seating

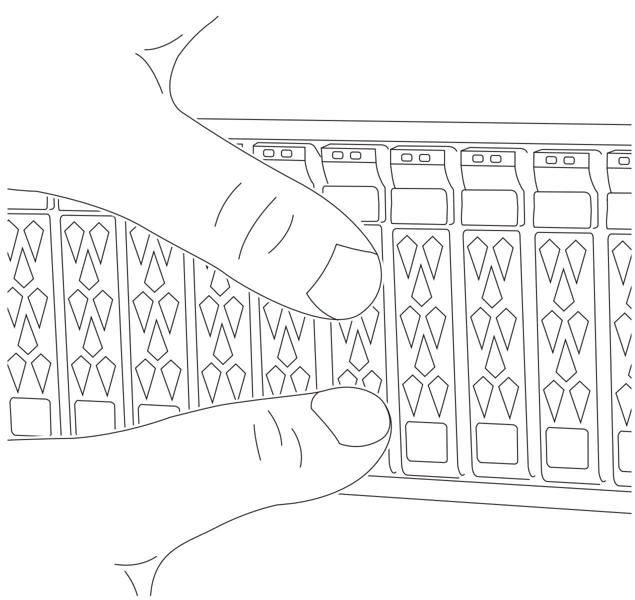

3. Check the controller drives based on the system status:
   a. On the healthy controller, check if any active RAID group is in a degraded state, failed state, or both:

   ```
   storage aggregate show -raidstatus !*normal*
   ```

      ▪ If the command returns `There are no entries matching your query.` continue to go to the next sub-step to check for missing drives.
      ▪ If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

      ```
      system node autosupport invoke -node * -type all -message
      '<message_name>'
      ```

   b. Check for missing drive issues for both the file system or spare drives:

   ```
   event log show -severity * -node * -message-name *disk.missing*
   ```

      ▪ If the command returns `There are no entries matching your query.` continue to go to the next step.
      ▪ If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

      ```
      system node autosupport invoke -node * -type all -message
      '<message_name>'
      ```

4. Remove the power cable retainers, then unplug the cables from the power supplies.
5. Loosen the hook and loop strap on the cable management device. Unplug the system cables and SFP/QSFP modules (if needed) from the controller module. Note each cable’s location.

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

6. Remove the cable management device from the controller module and set it aside.
7. 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.

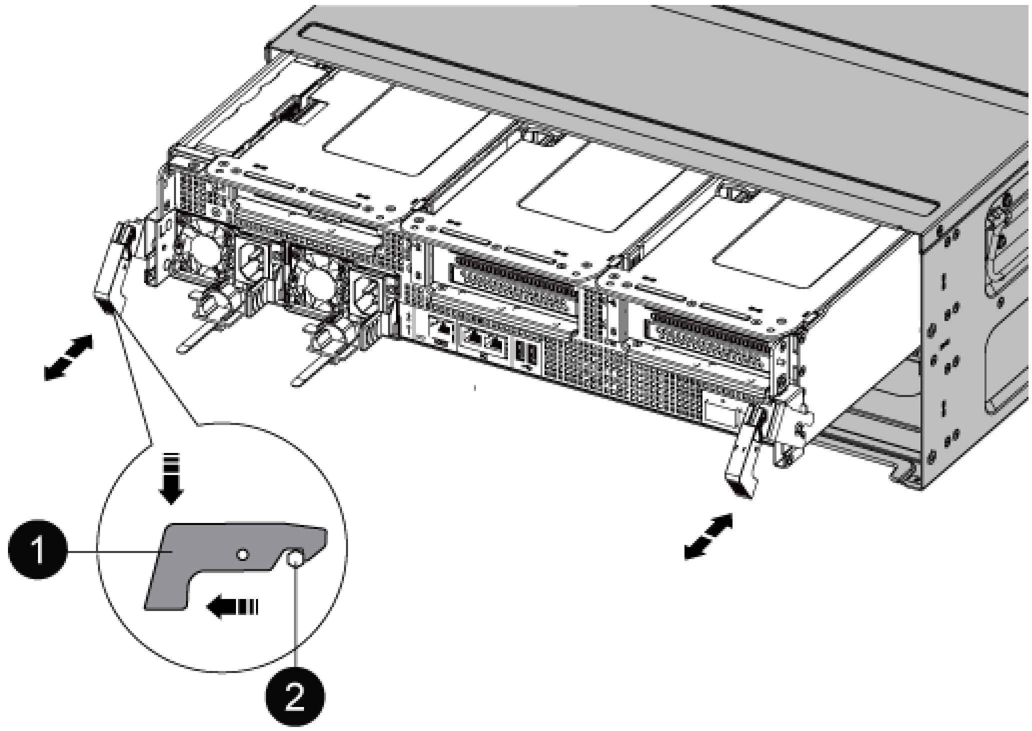

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a stable, flat surface.

   Support the bottom of the controller module while sliding it out of the chassis.

Step 3: Replace the PCIe card

Replace a PCIe card by removing the cabling and any QSFPs and SFPs from the ports on the PCIe cards in the target riser, removing the riser from the controller module, removing and replacing the PCIe card, reinstalling the riser and any QSFPs and SFPs onto the ports, and recable the ports.

Steps

1. Determine if the card you are replacing is from Riser 1 or if it is from Riser 2 or 3.
   - If you are replacing the 100GbE PCIe card in Riser 1, use Steps 2 - 3 and Steps 6 - 7.
   - If you are replacing a PCIe card from Riser 2 or 3, use Steps 4 through 7.
2. Remove Riser 1 from the controller module:
   a. Remove the QSFP modules that might be in the PCIe card.
   b. Rotate the riser locking latch on the left side of the riser up and toward the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

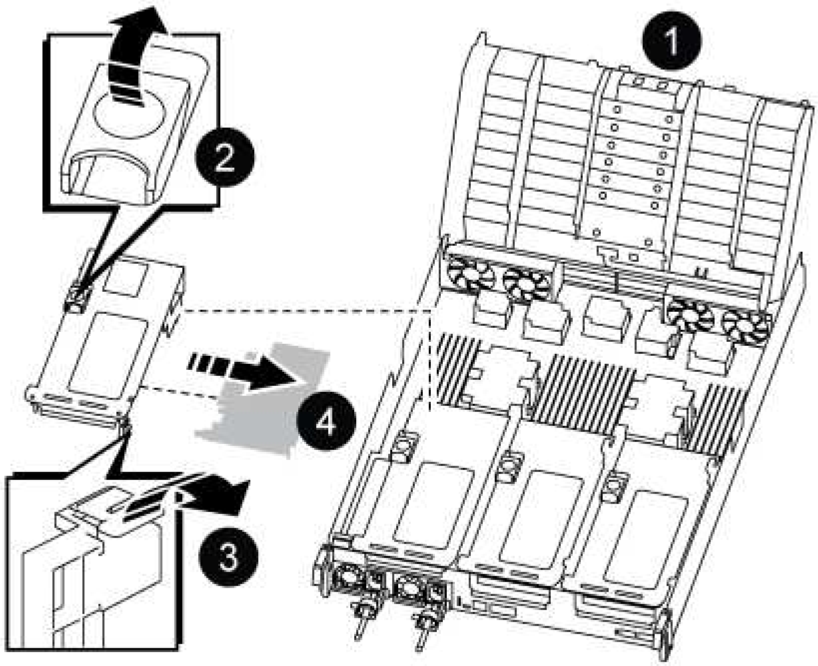

1	Air duct
2	Riser locking latch
3	Card locking bracket
4	Riser 1 (left riser) with 100GbE PCIe card in slot 1.

3. Remove the PCIe card from Riser 1:
   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. Remove the PCIe card from the riser.
4. Remove the PCIe riser from the 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 the fan modules.

      The riser raises up slightly from the controller module.
   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

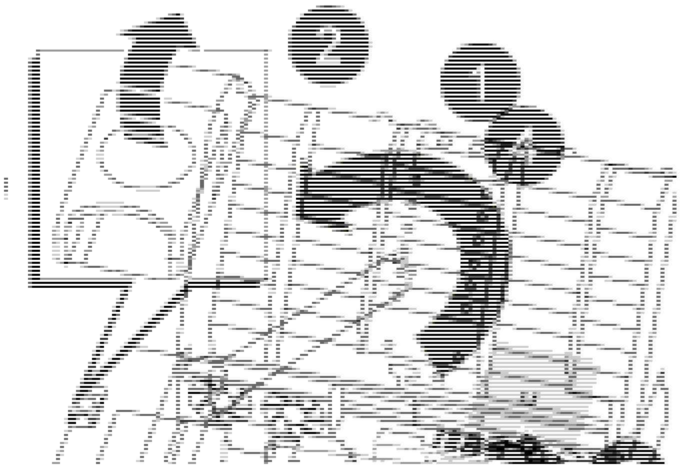

1	Air duct
2	Riser 2 (middle riser) or 3 (right riser) locking latch
3	Card locking bracket
4	Side panel on riser 2 or 3
5	PCIe cards in riser 2 or 3

5. Remove the PCIe card from the riser:
    a. Turn the riser so that you can access the PCIe cards.
    b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
    c. Swing the side panel off the riser.
    d. Remove the PCIe card from the riser.
6. Install the PCIe card into the same slot in the riser:
    a. Align the card with the card socket in the riser, and then slide it squarely into the socket in the riser.

        Make sure that the card is completely and squarely seated into the riser socket.

    b. For Riser 2 or 3, close the side panel.
    c. Swing the locking latch into place until it clicks into the locked position.
7. Install the riser into the controller module:
    a. Align the lip of the riser with the underside of the controller module sheet metal.
    b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
    c. Swing the locking latch down and click it into the locked position.

        When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

d. Reinsert any SFP modules that were removed from the PCIe cards.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

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 system, as needed.

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

3. Complete the reinstallation of the controller module:

   a. 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. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them 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.

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

   `storage failover giveback -ofnode` *`impaired_node_name`*.

5. If automatic giveback was disabled, reenable it:

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

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

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

Step 5: 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 - AFF C800

Replace an AC or DC power supply unit (PSU) in your AFF C800 system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cable, replacing the faulty PSU, and then reconnecting it to the power source.

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 it to the power source.

The power supplies are redundant and hot-swappable. You do not have to shut down the controller to replace a PSU.

About this task

This procedure is written for replacing one PSU at a time.

> It is a best practice to replace the PSU within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded PSU until the PSU is replaced.

> Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.

Use the appropriate procedure for your type of PSU: AC or DC.

Option 1: Hot-swap an AC PSU

To replace an AC PSU, complete the following steps.

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
   b. Unplug the power cable from the power source.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

> ! The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

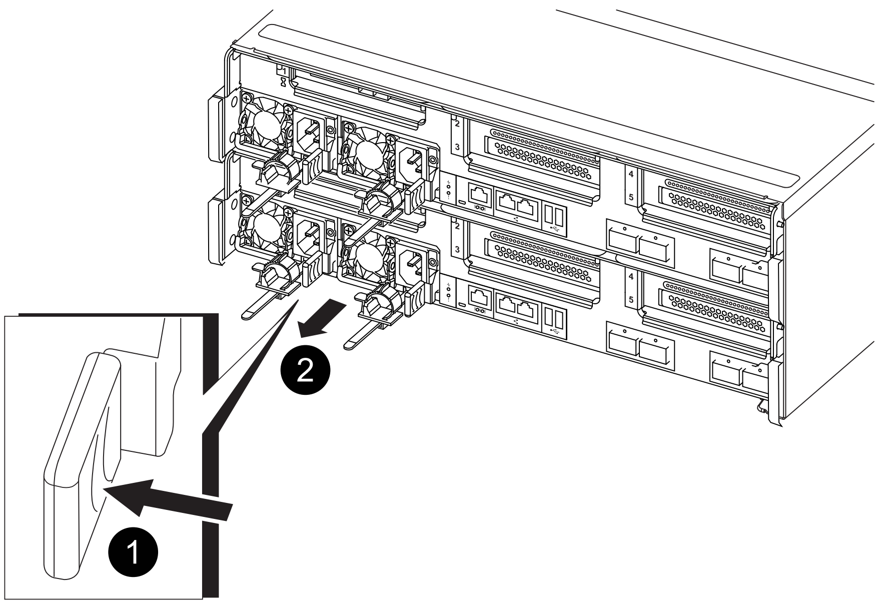

1	Blue PSU locking tab
2	Power supply

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU 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 PSU into the system.

6. Reconnect the PSU cabling:
   a. Reconnect the power cable to the PSU and the power source.
   b. Secure the power cable to the PSU using the power cable retainer.

   Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return and Replacements](https://mysupport.netapp.com/site/info/rma) page for further information.

Option 2: Hot-swap a DC PSU

To replace a DC PSU, complete the following steps.

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC cable connector using the thumb screws on the plug.
   b. Unplug the cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

   > ! The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.

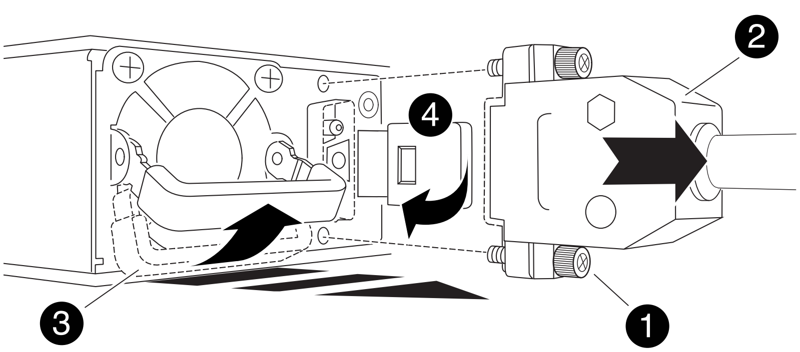

1	Thumb screws
2	D-SUB DC power PSU cable connector
3	Power supply handle

4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU 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 PSU into the system.
6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

      Once power is restored to the PSU, the status LED should be green.
7. 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 - AFF C800

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your AFF C800 system to ensure that services and applications relying on accurate time synchronization remain operational.

Before you begin

- Understand that you can use this procedure with all versions of ONTAP supported by your system.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller.

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 `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the chassis when you replace the controller module or replace a component inside the controller module.

Steps

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

2. Ensure that all drives in the chassis are firmly seated against the midplane by using your thumbs to push each drive until you feel a positive stop.

   Video - Confirm drive seating

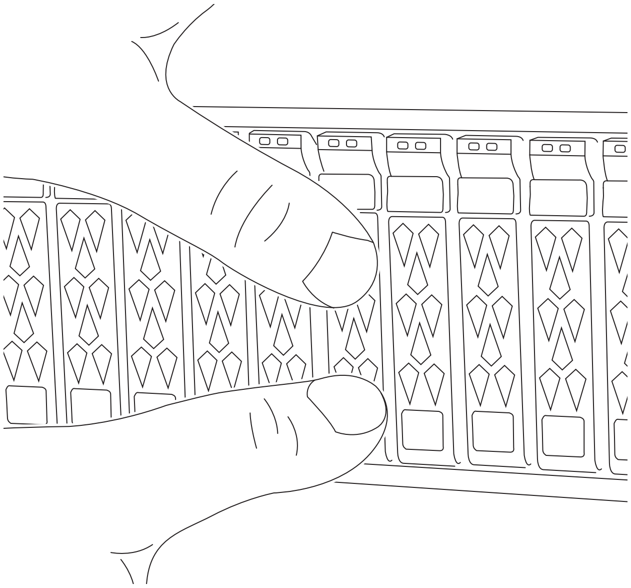

3. Check the controller drives based on the system status:

   a. On the healthy controller, check if any active RAID group is in a degraded state, failed state, or both:

   ```
   storage aggregate show -raidstatus !*normal*
   ```

      - If the command returns `There are no entries matching your query.` continue to go to the next sub-step to check for missing drives.
      - If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

      ```
      system node autosupport invoke -node * -type all -message
      '<message_name>'
      ```

   b. Check for missing drive issues for both the file system or spare drives:

   ```
   event log show -severity * -node * -message-name *disk.missing*
   ```

      - If the command returns `There are no entries matching your query.` continue to go to the next step.
      - If the command returns any other results, collect the AutoSupport data from both controllers and contact NetApp Support for further assistance.

      ```
      system node autosupport invoke -node * -type all -message
      '<message_name>'
      ```

4. Remove the power cable retainers, then unplug the cables from the power supplies.

5. Loosen the hook and loop strap on the cable management device. Unplug the system cables and SFP/QSFP modules (if needed) from the controller module. Note each cable’s location.

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

6. Remove the cable management device from the controller module and set it aside.
7. 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.

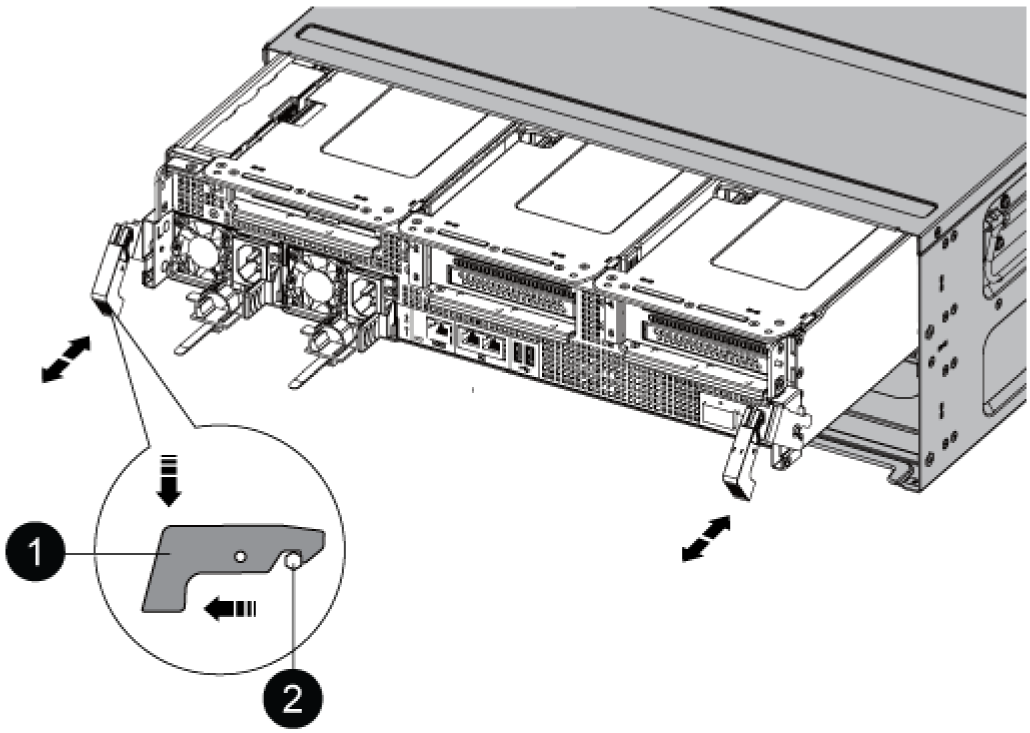

1	Locking latch
2	Locking pin

8. Slide the controller module out of the chassis and place it on a stable, flat surface.

   Support the bottom of the controller module while sliding it out of the chassis.

Step 3: Replace the RTC battery

Replace the RTC battery.

The procedure for replacing the RTC battery differs depending on whether your controller is a Original or VER2 model. Use the tabs below to select the appropriate instructions for your controller model.

About this task

The battery is located under Riser 2 (the middle riser) on Original controllers and near the DIMMs on VER2 controllers.

Original controller

Steps

1. Remove PCIe riser 2 (middle riser) from the 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 the fan modules.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, shift it toward the fans so that the sheet metal lip on the riser clears the edge of the controller module, lift the riser out of the controller module, and then place it on a stable, flat surface.

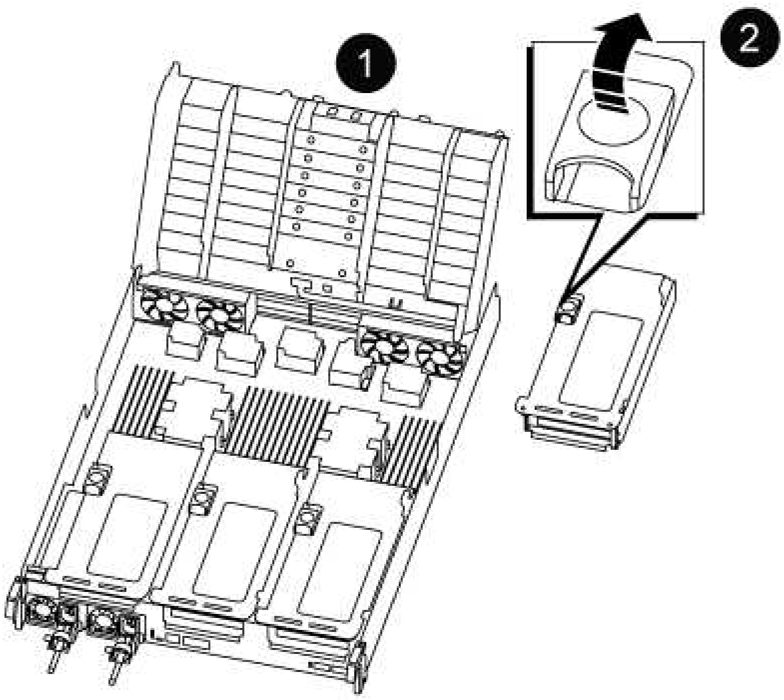

1	Air duct
2	Riser 2 (middle riser) locking latch

2. Locate the RTC battery under Riser 2.

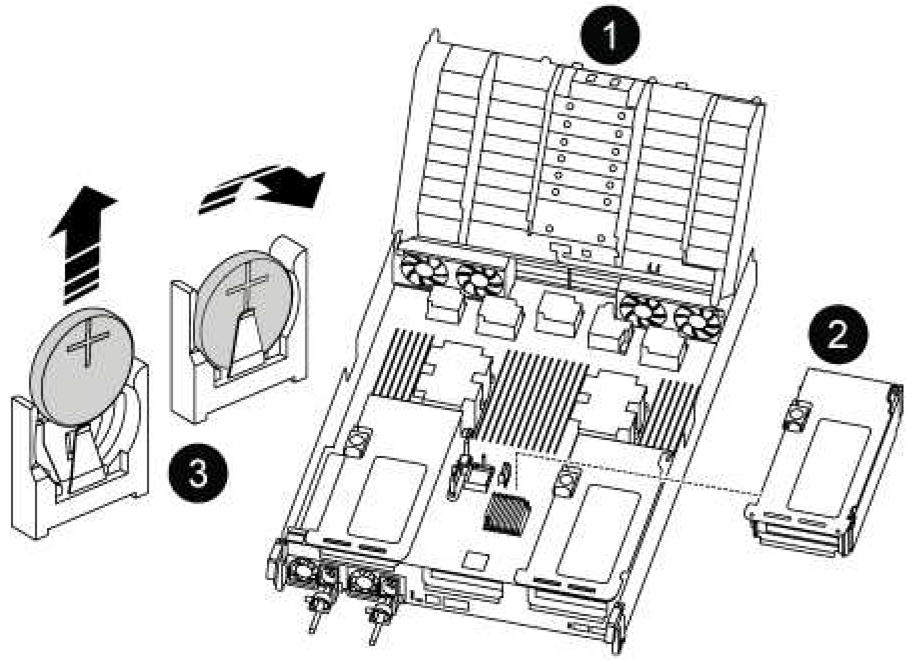

1	Air duct
2	Riser 2
3	RTC battery and housing

3. 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.

4. Remove the replacement battery from the antistatic shipping bag.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
7. Install the riser into the controller module:
   a. Align the lip of the riser with the underside of the controller module sheet metal.
   b. Guide the riser along the pins in the controller module, and then lower the riser into the controller module.
   c. Swing the locking latch down and click it into the locked position.

      When locked, the locking latch is flush with the top of the riser and the riser sits squarely in the controller module.

d. Reinsert any SFP modules that were removed from the PCIe cards.

VER2 controller

Steps

1. Locate the RTC battery near the DIMMs.

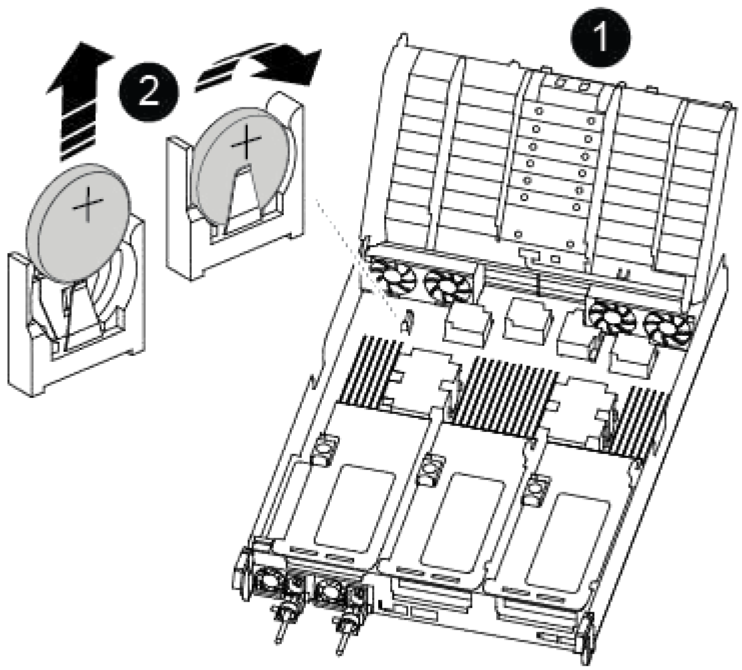

1	Air duct
2	RTC battery and housing

2. 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.

3. Remove the replacement battery from the antistatic shipping bag.
4. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
5. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

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 system, as needed.

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

3. Complete the reinstallation of the controller module:

   a. 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. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them 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.

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

   `storage failover giveback -ofnode` *`impaired_node_name`*.

5. If automatic giveback was disabled, reenable it:

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

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

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

Step 5: 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.