NetApp

Maintain

Install and maintain

NetApp
February 06, 2026

Table of Contents

Maintain

Overview of the maintenance procedures - AFF C80

Maintain the hardware of your AFF C80 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 C80 system has already been deployed as a storage node in the ONTAP environment.

System components

For the AFF C80 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 runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Drive	A drive is a device that provides the physical storage needed for data.
Fan	A fan cools the controller.
NVRAM	The NVRAM (Non-Volatile Random Access Memory) is a module that allows the controller to protect and save in-flight data if the system loses power. The system ID resides in the NVRAM module. When replaced, the controller assumes the new system ID from the replacement NVRAM module.

NV battery	The NV battery is responsible for providing power to the NVRAM module while data in-flight is being destaged to flash memory after a power loss.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.
System Management module	The System Management module provides the interface between the controller and a console or laptop for controller or system maintenance purposes. The System Management module contains the boot media and stores the system serial number (SSN).

Boot media - automated recovery

Boot media automated recovery workflow - AFF C80

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 C80 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 System Management 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 C80

Before replacing the boot media in your AFF C80 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming the cluster ports on the impaired controller are working properly, 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.

Review the following requirements.

- You must replace the failed boot media with a replacement boot media you received from NetApp.
- The cluster ports are used to communicate between the two controllers during the automated boot recovery process. Make sure that the cluster ports on the impaired controller are working properly.
- 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
    - /cfcard/kmip/certs/client.crt
    - /cfcard/kmip/certs/client.key
    - /cfcard/kmip/certs/CA.pem
- Understand the controller terminology used in this procedure:
    - 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 C80

Shut down the impaired controller in your AFF C80 storage system to prevent data loss and 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 C80

The boot media in your AFF C80 storage 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 in the

System Management module, and then reinstalling the System Management 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.

Replace the boot media, which is located inside the System Management module at the rear of the controller.

Before you begin

- You must have a replacement boot media.
- Have an anti-static mat available for the System Management module.

Steps

1. Verify that NVRAM destage has completed before proceeding. When the LED on the NV module is off, NVRAM is destaged.

   If the LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact NetApp Support for assistance.
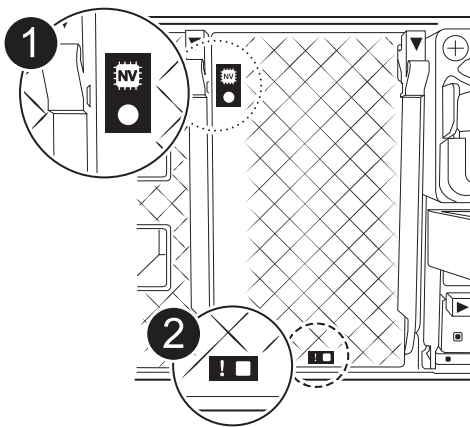
①	NVRAM status LED
②	NVRAM attention LED

2. Go to the rear of the chassis and properly ground yourself if you are not already grounded.
3. Disconnect power from the controller:
   - For AC power supplies, disconnect the power cords from the power supplies.
   - For DC power supplies, disconnect the power block from the power supplies.
4. Remove the System Management module:
   a. Remove any cables connected to the System Management module. Label the cables to identify their correct ports for reinstallation.
   b. Rotate the cable management arm down by pulling the buttons on both sides of the cable management arm.
   c. Depress the system management cam button.

      The cam lever moves away from the chassis.

d. Rotate the cam lever all the way down and remove the System Management module from the controller.

e. Place the System Management module on an anti-static mat with the boot media accessible.

5. Remove the boot media from the System Management module:
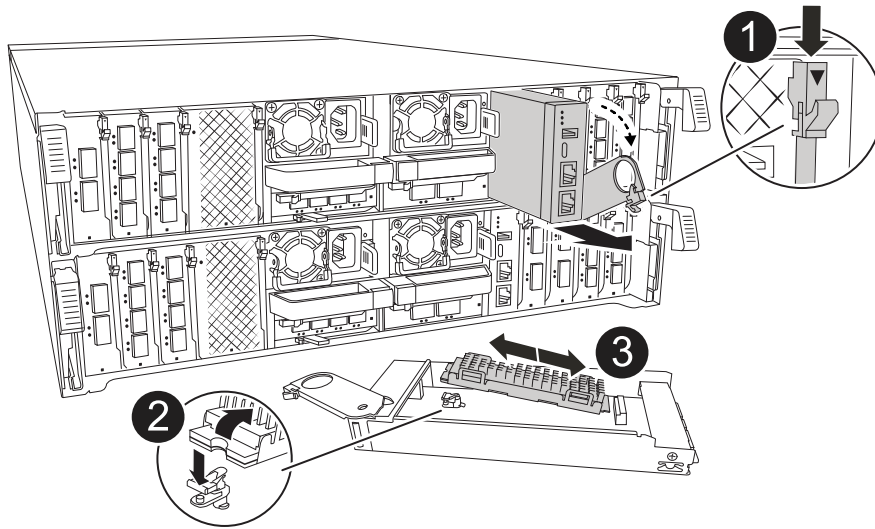
❶	System Management module cam latch
❷	Boot media locking button
❸	Boot media

   a. Press the blue locking button.

   b. Rotate the boot media up, slide it out of the socket, and set it aside.

6. Install the replacement boot media into the System Management 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 locking button.

   c. Push the locking button, rotate the boot media all the way down, and then release the locking button.

7. Reinstall the System Management module:

   a. Align the edges of the System Management module with the chassis opening.

   b. Gently slide the module into the chassis until the cam latch begins to engage.

   c. Rotate the cam latch all the way up to lock the module in place.

   d. Reconnect the cables to the System Management module using the labels you created during removal.

   e. Rotate the cable management arm up to the closed position.

8. Reconnect power to the controller:

   ◦ For AC power supplies, plug the power cords into the power supplies.

   ◦ For DC power supplies, reconnect the power block to the power supplies.

The controller automatically reboots when power is restored.

9. Interrupt the boot process by pressing `Ctrl-C` to stop 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 C80

After installing the new boot media device in your AFF C80 storage 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`

```
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 part to NetApp - AFF C80

If a component in your AFF C80 storage 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 C80

The manual recovery of the boot image involves using a USB drive to reinstall ONTAP onto the AFF C80 system's replacement boot media. You must download the appropriate ONTAP recovery image from the NetApp Support Site and copy it to a USB drive. This prepared USB drive is then used to perform the recovery and restore the system to operational status.

If your system is running in ONTAP 9.17.1 and later, use the automatic boot recovery procedure.

To get started, review the recovery requirements, shut down the controller, replace the boot media, use the USB drive to restore the image, and reapply encryption settings if necessary.

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 impaired 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 ONTAP boot menu.

⑦ **Return the failed part to NetApp**

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

Requirements for manual boot media recovery - AFF C80

Before replacing the boot media in your AFF C80 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 system is running in ONTAP 9.17.1 and later, use the automatic 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 support for manual boot media recovery - AFF C80

To ensure data security on your AFF C80 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 system is running in ONTAP 9.17.1 and later, use the automatic 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.

> (i) 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.

What's next?

After checking the encryption key support and status on the boot media, you need to shut down the controller.

Shut down the controller for manual boot media recovery - AFF C80

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

If your system is running in ONTAP 9.17.1 and later, use the automatic boot recovery procedure.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `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.

Option 2: Controller is in a MetroCluster

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.

- 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:

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

   ```
   storage failover modify -node local -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 section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…

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 -halt true`
	
	The *-halt true* parameter brings you to the LOADER prompt.

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 C80

The boot media in your AFF C80 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 system is running in ONTAP 9.17.1 and later, use the automatic boot recovery procedure.

Step 1: Replace the boot media

The boot media is located inside the System Management module and is accessed by removing the module from the system.

Steps

1. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.

2. Unplug the controller's PSUs.

   > ⓘ If your system has DC power, disconnect the power block from the PSUs.

   a. Remove any cables connected to the System Management module. Make sure to label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.

   b. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.

   c. Depress the system management cam button. The cam lever moves away from the chassis.

   d. Rotate the cam lever all the way down and remove the System Management module from the controller module.

   e. Place the System Management module on an anti-static mat, so that the boot media is accessible.

3. Remove the boot media from the management module:
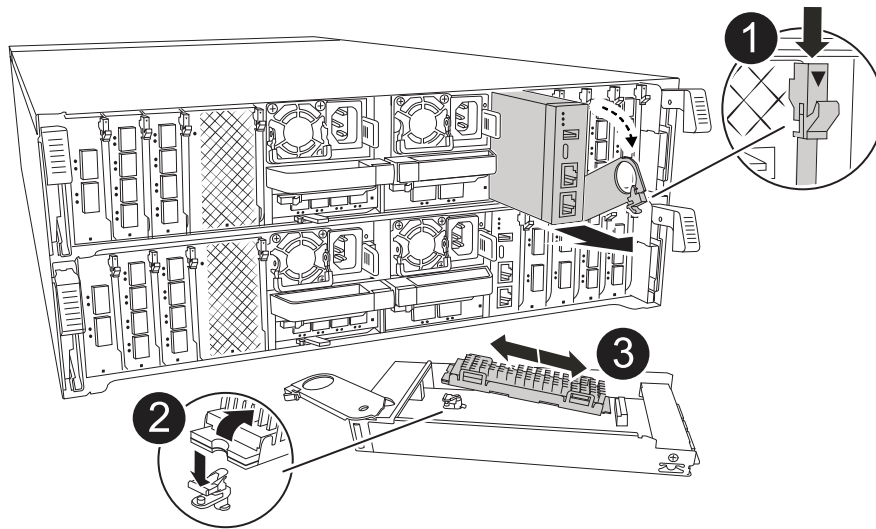
❶	System Management module cam latch
❷	Boot media locking button
❸	Boot media

    a. Press the blue locking button.

    b. Rotate the boot media up, slide it out of the socket, and set it aside.

4. Install the replacement boot media into the System Management 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 locking button.

    c. Push the locking button, rotate the boot media all the way down and then release the locking button.

5. Reinstall the System Management module:

    a. Rotate the cable management tray up to the closed position.

    b. Recable the System Management module.

Step 2: Transfer the boot image to the boot media

The replacement boot media that you installed is without an ONTAP image. You can transfer the ONTAP image to the replacement boot media by downloading the appropriate ONTAP service image from the NetApp Support Site to a USB flash drive and then to the replacement boot media.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.

- Download 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. Use the `version -v` command to display if your version of ONTAP supports NVE. If the command output displays `<10no- DARE>`, your version of ONTAP does not support NVE.

    ◦ If NVE is supported by your version of ONTAP, download the image with NetApp Volume Encryption,

as indicated in the download button.

  ◦ If NVE is not supported, 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 between the node management ports of
  the controllers (typically the e0M interfaces).

Steps

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.

   a. Download the service image from the Downloads link on the page, to your work space on your laptop.

   b. Unzip the service image.

      (i) 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.

      The USB flash drive should have the appropriate ONTAP image of what the impaired controller is
      running.

   c. Remove the USB flash drive from your laptop.

2. Insert the USB flash drive into the USB-A port on the System Management module.

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

3. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

   (i) If you have DC power supplies, reconnect the power block to the power supplies.

4. 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 C80

After installing the new boot media device in your AFF C80 system, you can boot the
recovery image manually from a USB drive to restore the configuration from the partner
node.

If your system is running in ONTAP 9.17.1 and later, use the automatic 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
```

> (i)   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 keys after manual boot recovery - AFF C80

Restore encryption on the replacement boot media in your AFF C80 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 system is running in ONTAP 9.17.1 and later, use the automatic boot recovery procedure.

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

Onboard Key Manager (OKM)

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

Before you begin

Ensure you have the following information available:

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

Steps

On the impaired controller:

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

ONTAP version	Select this option

ONTAP 9.8 or later	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>

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	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.
	Selection (1-19)?
	recover_onboard_keymanager</pre>

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-------------------------
012345678901234567890123456789012345678901234567890123456789012345678901
23
123456789012345678901234567890123456789012345678901234567890123456789012
34
234567890123456789012345678901234567890123456789012345678901234567890123
45
345678901234567890123456789012345678901234567890123456789012345678901234
56
456789012345678901234567890123456789012345678901234567890123456789012345
67
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
```

```
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
0123456789012345678901234567890123456789012345678901234567890 1
23
1234567890123456789012345678901234567890123456789012345678901 2
34
2345678901234567890123456789012345678901234567890123456789012 3
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.

```
*****************************************************************
***********
* 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 part to NetApp - AFF C80

If a component in your AFF C80 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Chassis

Chassis replacement workflow - AFF C80

Get started with replacing the chassis in your AFF C80 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 chassis replacement**

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

Requirements to replace the chassis - AFF C80

Before replacing the chassis in your AFF C80 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.

Shut down the controller to replace the chassis - AFF C80

Shut down the controller in your AFF C80 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
```

> (i) 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 you've shut down the controller, you need to replace the chassis.

Replace the chassis - AFF C80

Replace the chassis in your AFF C80 system when a hardware failure requires it. The replacement process involves removing the controllers and power supply units (PSUs), removing the drives, installing the replacement chassis, and reinstalling the chassis components.

Step 1: Remove the PSUs and cables

You need to remove all four power supply units (PSUs), two per controller, before removing the controller. Removing them lightens the overall weight of each controller.

Steps
1. Remove the four PSUs:
   a. If you are not already grounded, properly ground yourself.
   b. Unplug power cords from the controller module PSU.

      If your system has DC power, disconnect the power block from the PSUs.

   c. Remove the PSU from the controller by rotating the PSU handle up so that you can pull the PSU out, press the PSU 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.
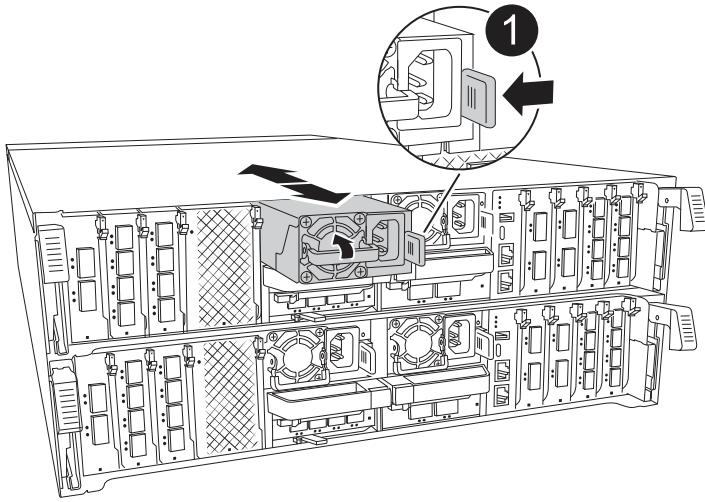
1	Terracotta PSU locking tab

    d. Repeat these steps for the remaining PSUs.

2. Remove the cables:

    a. Unplug the system cables and any SFP and QSFP modules (if needed) from the controller module, but leave them in the cable management device to keep them organized.

> ℹ️ Cables should have been labeled at the beginning of this procedure.

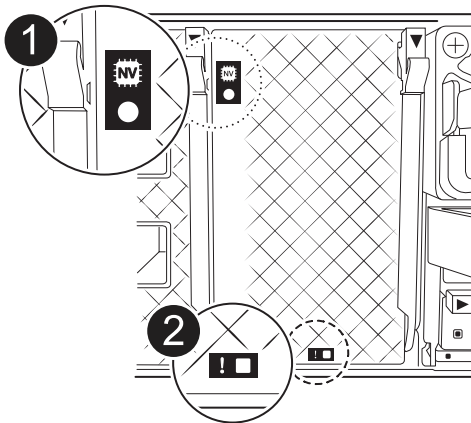   b. Remove the cable management device from the controller modules and set them aside.

Step 2: Remove the controller modules and drives

Remove the controllers from the chassis and then remove the drives from the chassis.

Steps

1. Check the that amber NVRAM status LED located in slot 4/5 on the back of each controller module is off. Look for the NV icon.

| 1	NVRAM status LED

❷	NVRAM attention LED

- ◦ If the NVRAM LED is off, go to the next step.
- ◦ If the NVRAM LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact NetApp Support Site for assistance.

2. Remove the controller modules:

a. Press down on both of the locking latches on the controller, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
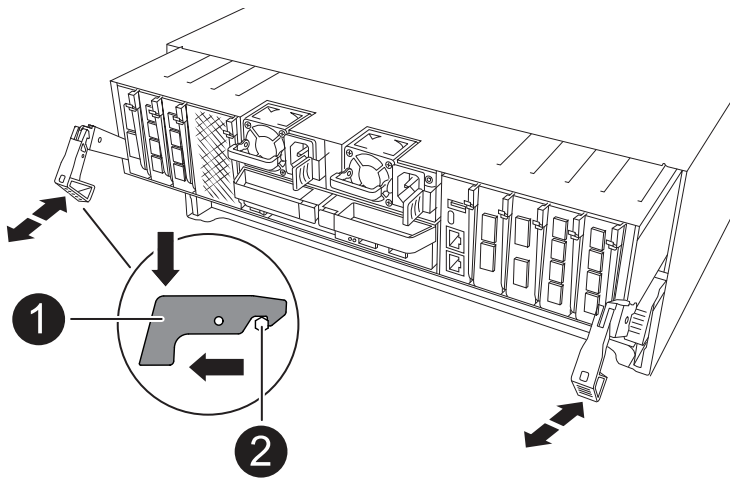
❶	Locking latch
❷	Locking pin

b. Slide the controller module out of the chassis by the locking latches, and place it on a flat, stable surface.

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

c. Repeat these steps for the second controller module.

3. Remove the drives:

a. Gently remove the bezel from the front of the system.

b. Press the release button at the top of the drive carrier face below the LEDs.

c. 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.

   ⓘ Drives are fragile. Always use two hands to support the drive weight when removing a drive to prevent damage to them.

d. Keep track of which drive bay each drive was in and set the drive aside on a static-free cart or table.

e. Repeat this step for the remaining drives in the chassis.

Step 3: Replace the chassis and install components

Remove the impaired chassis, install the replacement chassis, and reinstall all components.

Steps

1. Remove the impaired chassis:

   a. Remove the screws from the chassis mount points.

   b. Using two people or a lift, slide the impaired chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.

2. Install the replacement chassis:

   a. Using two people or a lift, 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.

   b. Slide the chassis all the way into the equipment rack or system cabinet.

   c. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the impaired chassis.

3. Beginning with the bottom controller module, install the controller modules in the replacement chassis:

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

   b. Rotate the locking latches upward into the locked position.

   c. If you have not already done so, reinstall the cable management device and recable the controller.

      If you removed the media converters (QSFPs or SFPs), remember to reinstall them.

      Make sure that the cables are connected referencing the cable labels.

4. Reinstall the drives into their corresponding drive bays in the front of the chassis.

5. Install all four of the PSUs:

   a. Using both hands, support and align the edges of the 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.

      (i) To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the PSU power cables to all four of the PSUs.

   a. Secure the power cable to the PSU using the power cable retainer.

      If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis and secure the power cable to the PSU with the thumbscrews.

   The controller modules begin to boot as soon as PSUs are installed and power is restored.

What's next?

After replacing the impaired chassis and reinstalling the components, complete the chassis replacement.

Complete chassis replacement - AFF C80

Reboot the controllers, verify system health, and return the failed part to NetApp to complete the final step in the AFF C80 chassis replacement procedure.

Step 1: Boot the controllers and perform giveback

Boot ONTAP on the controllers and perform controller giveback to return storage ownership.

Steps

1. Check the console output:

    a. If the controller boots to the LOADER prompt, reboot the controller with the `boot_ontap` command.

    b. If the console displays `waiting for giveback` after the reboot, log into the partner controller and check that the replaced controller is ready for giveback with the `storage failover show` command.

2. Perform the giveback:

    a. Connect the console cable to the partner controller.

    b. Give back the controller with the `storage failover giveback -fromnode local` command.

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 C80

Get started with replacing the controller in your AFF C80 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 controller replacement requirements**

To replace the controller module, you must meet certain requirements.

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**

Replacing the controller includes removing the impaired controller, moving the FRU components to the replacement controller module, and then installing 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 C80

Before replacing the controller in your AFF C80 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 C80 controller, you need to shut down the impaired controller.

Shut down the impaired controller - AFF C80

Shut down the controller in your AFF C80 storage system to prevent data loss and ensure system stability when replacing the controller.

Shut down the controller module using one of the following options.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `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.

Option 2: Controller is in a MetroCluster

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.

- 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:

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

   ```
   storage failover modify -node local -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 section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
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 -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 - AFF C80

Replace the controller in your AFF C80 system when a hardware failure requires it. This process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting the system.

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>'
   ```
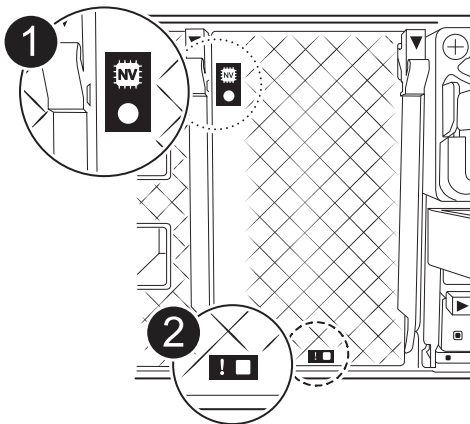
4. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

1	NVRAM status LED

❷	NVRAM attention LED

- ◦ If the NV LED is off, go to the next step.
- ◦ If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

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

6. Unplug the controller module power supply cables from the controller module power supplies (PSU).

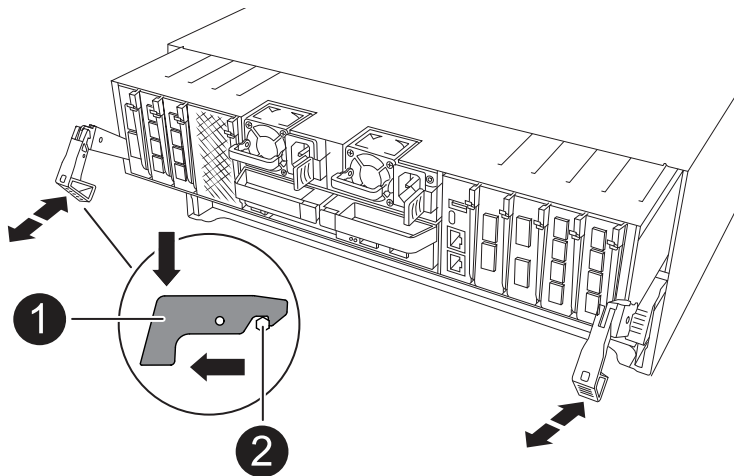> ⓘ If your system has DC power, disconnect the power block from the PSUs.

7. 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.

8. Remove the cable management device from the controller module.

9. 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.

❶	Locking latch
❷	Locking pin

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

    Make sure that you support the bottom of the controller module as you slide it out of the chassis.
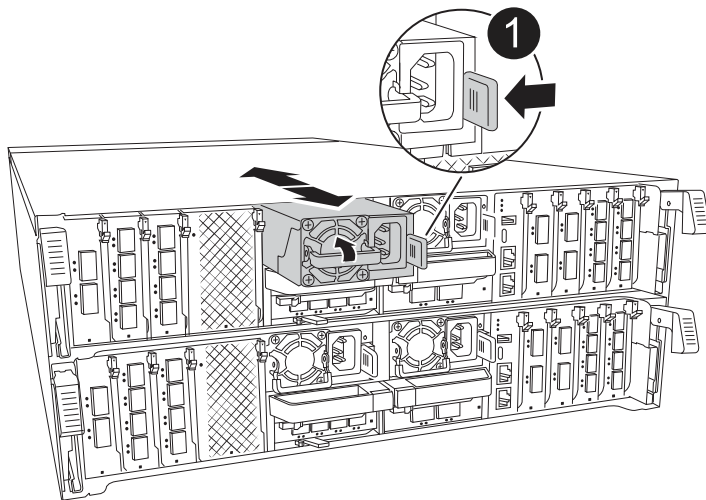
Step 2: Move the power supplies

Move the power supplies to the replacement controller.

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.

❶	Terracotta PSU locking tab
❷	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.
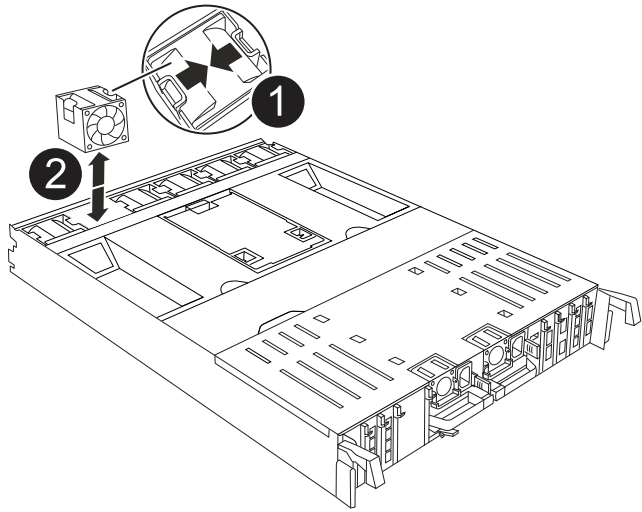
   > ⓘ To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

Step 3: Move the fans

Move the fans 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 it by aligning its edges with the opening in the controller module and sliding it in until the locking latches click into place.

3. Repeat these steps for the remaining fan modules.

Step 4: Move the NV battery

Move the NV battery to the replacement controller module.

Steps

1. Open the air duct cover in the middle of the controller module and locate the NV battery.
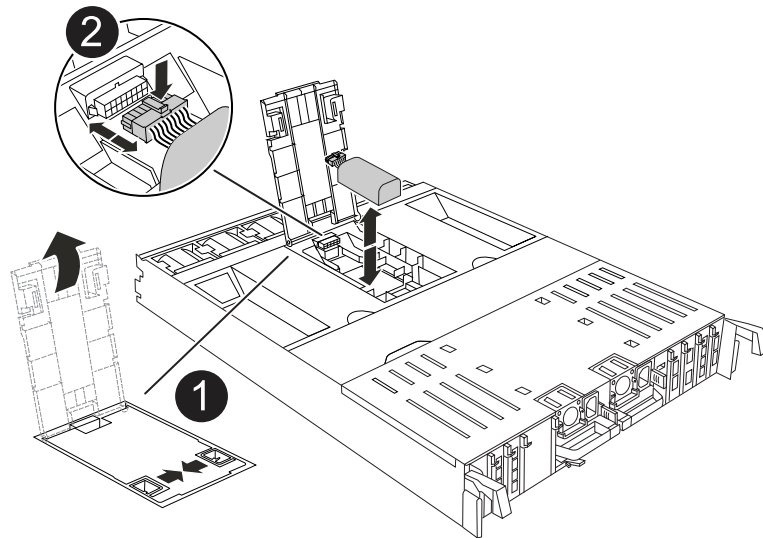
①	NV battery air duct

②	NV battery pack plug

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

2. Lift the battery up to access the battery plug.

3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.

4. Lift the battery out of the air duct and controller module.

5. Move the battery pack to the replacement controller module and then install it in the replacement controller module:

   a. Open the NV battery air duct in the replacement controller module.

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

   c. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.

   d. Close the NV battery air duct.

Step 5: Move system DIMMs

Move the DIMMs to the replacement controller module.

Steps

1. Open the controller air duct on the top of the controller.

   a. Insert your fingers in the recesses at the far ends of the air duct.

   b. Lift the air duct and rotate it upward as far as it will go.

2. Locate the system DIMMs on the motherboard, using the DIMM map on top of the air duct.

   The DIMM locations, by model, are listed in the following table:

Model	DIMM slot location
FAS70	3, 10, 19, 26
FAS90	3, 7, 10, 14, 19, 23, 26, 30
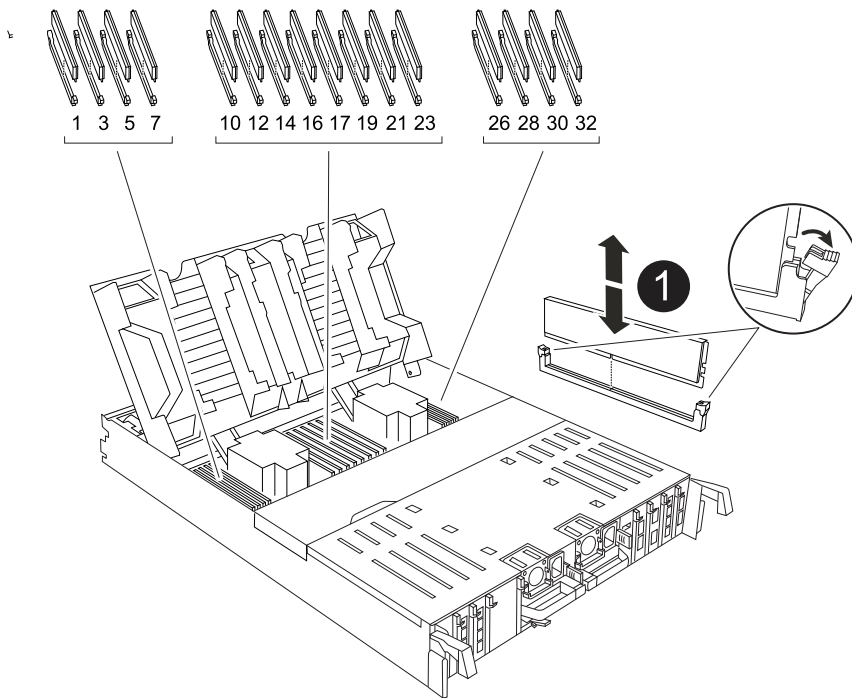
1	System DIMM

3. 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.

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.

> (i) Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

5. Locate the slot on the replacement controller module where you are installing the DIMM.

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.

> (i) 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. Repeat these steps for the remaining DIMMs.

9. Close the controller air duct.
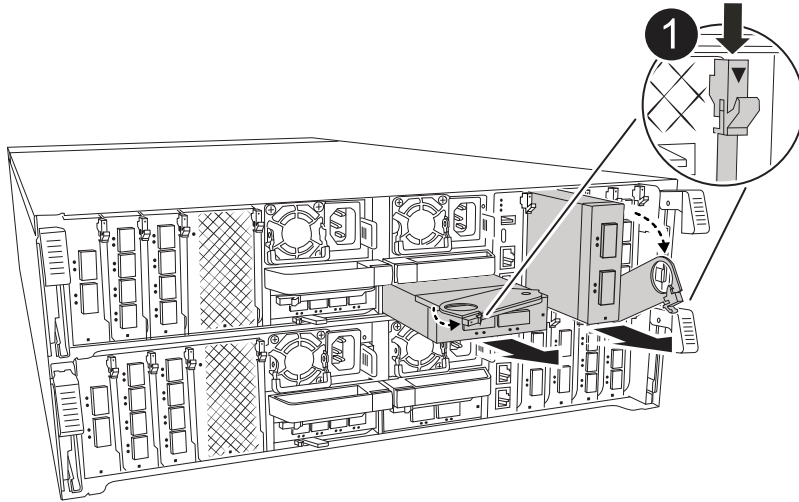
Step 6: Move the I/O modules

Move the I/O modules to the replacement controller module.

Steps

1. Unplug any cabling on the target I/O module.

Make sure that you label the cables so that you know where they came from.

2. Rotate the cable management arm down by pulling the buttons on the inside of the cable management arm and rotating it down.

1	I/O module cam lever

3. Remove the I/O modules from the controller module:

   a. Depress the target I/O module cam latch button.

   b. Rotate the cam latch down as far as it will go. For horizontal modules, rotate the cam away from the module as far as it will go.

   c. Remove the module from the controller module by hooking your finger into the cam lever opening and pulling the module out of the controller module.

      Make sure that you keep track of which slot the I/O module was in.

   d. Install the replacement I/O module into the replacement controller module by gently sliding the I/O module into the slot until the I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.

4. Repeat these steps to move the remaining I/O modules, except the modules in slots 6 and 7, to the replacement controller module.

   > (i) To move the I/O modules from slots 6 and 7, you must move the carrier containing these I/O modules from the impaired controller module to the replacement controller module.

5. Move the carrier containing the I/O modules in slots 6 and 7 to the replacement controller module:

   a. Press the button on the right-most handle on the carrier handle. ..Slide the carrier out of the impaired controller module insert it into the replacement controller module in the same position it was in the impaired controller module.

   b. Gently push the carrier all the way into the replacement controller module until it locks into place.

Step 7: Move the System Management module

Move the System Management module to the replacement controller module.

Steps

1. Remove the System Management module from the impaired controller module:
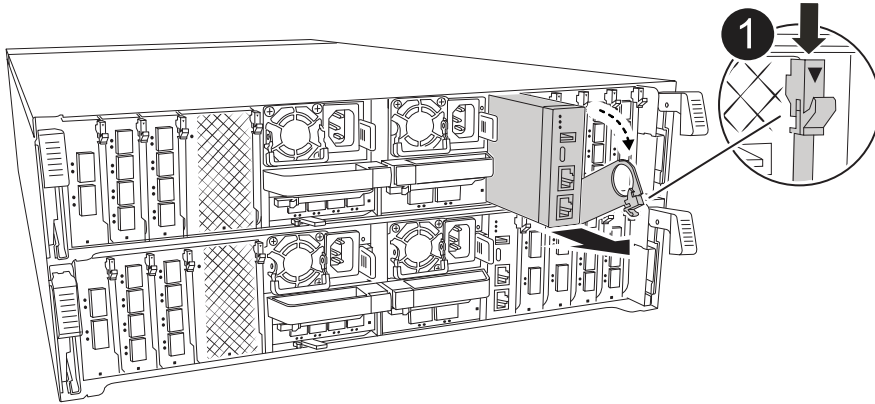
1	System Management module cam latch

   a. Depress the system management cam button.

   b. Rotate the cam lever all the way down.

   c. Loop your finger into the cam lever and pull the module straight out of the system.

2. Install the system management module into the replacement controller module in the same slot that it was in on the impaired controller module:

   a. Align the edges of the System Management module with the system opening and gently push it into the controller module.

   b. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.

Step 8: Move the NVRAM module

Move the NVRAM module to the replacement controller module.

Steps

1. Remove the NVRAM module from the impaired controller module:
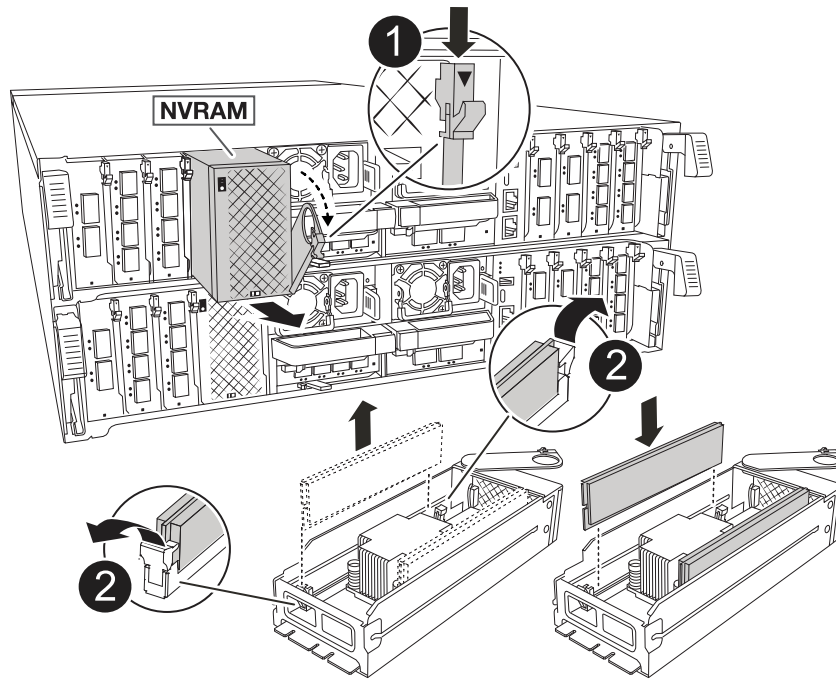
①	Cam locking button
②	DIMM locking tab

   a. Depress the cam latch button.

       The cam button moves away from the chassis.

   b. Rotate the cam latch as far as it will go.

   c. Remove the NVRAM module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

2. Install the NVRAM module into slot 4/5 in the replacement controller module:

   a. Align the module with the edges of the chassis opening in slot 4/5.

   b. Gently slide the module into the slot all the way, and then push the cam latch all the way up to lock the module in place.

Step 9: Install the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

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

3. Reinstall the cable management arm, if removed, but do not reconnect any cables to the replacement controller.

4. Plug the console cable into the console port of the replacement controller module and reconnect it to the laptop so that so that it receives console messages when it reboots.

5. 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.

    a. Rotate the locking latches upward into the locked position.

    b. Plug in the power supplies. The controller boots to the LOADER prompt as soon as power is restored.

6. From the LOADER prompt, enter `show date` to display the date and time on the replacement controller. Date and time are in GMT.

	Time displayed is local time not always GMT and is displayed in 24hr mode.

7. Set the current time in GMT with the `set time hh:mm:ss` command. You can get the current GMT from the partner node the `date -u` command.

8. Recable the storage system, as needed.

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

9. Plug the power cords into the power supplies.

	If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

What's next?

After you've replaced the impaired AFF C80 controller, you need to restore the system configuration.

Restore and verify the system configuration - AFF C80

Verify that the controller's HA configuration is active and functioning correctly in your AFF C80 storage system, and confirm that the system's adapters list all the paths to the disks.

Verify the HA configuration settings and disk list before returning the system to service.

Steps

1. Boot to maintenance mode:

```
boot_ontap maint
```

2. Enter `y` when you see *Continue with boot?*.

   If you see the *System ID mismatch* warning message, enter `y`.

3. Enter `sysconfig -v` and capture the display contents.

   > ⓘ   If you see *PERSONALITY MISMATCH* contact customer support.

4. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.

5. Verify that all components display the same `HA` state:

   ```
   ha-config show
   ```

   The HA state should be the same for all components.

6. 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
   ```

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

   - `ha`
   - `mcc` (not supported)
   - `mccip` (not supported in ASA systems)
   - `non-ha` (not supported)

7. Confirm that the setting has changed:

   ```
   ha-config show
   ```

8. Verify that the adapter lists the paths to all disks: `storage show disk -p`

   If you see any issues, check cabling and reseat cables.

9. Exit Maintenance mode: `halt`

What's next?

After you've restored and verified the system configuration for your AFF C80 system, you need to give back the controller.

Give back the controller - AFF C80

Return control of storage resources to the replacement controller so your AFF C80 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.

No encryption

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

Steps

1. From the LOADER prompt, enter `boot_ontap`.

2. Press <enter> when console messages stop.

    ◦ If you see the *login* prompt, go to the next step at the end of this section.

    ◦ If you see *Waiting for giveback*, press the <enter> key, log into the partner node, and then go to the next step at the end of this section.

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

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

5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

1. From the LOADER prompt, enter `boot_ontap maint`.

2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap menu` and select option 10.

3. Enter the OKM passphrase.

    > ℹ️ You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.

5. At the boot menu, enter option `1` for normal boot.

6. Press <enter> when *Waiting for giveback* is displayed.

7. Move the console cable to the partner node and login as `admin`.

8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates` *true*

    ◦ If you encounter errors, contact NetApp Support.

9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.

10. Synchronize and verify status of the keys:

    a. Move the console cable back to the replacement controller.

    b. Synchronize missing keys: `security key-manager onboard sync`

    > ℹ️ You are prompted for the cluster-wide passphrase of OKM for the cluster.

    c. Verify status of the keys: `security key-manager key query -restored` *false*

The output should show no results when when properly synchronized.

If the output shows results (the key IDs of keys that are not present in the system's internal key table), contact NetApp Support.

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

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

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

Complete controller replacement - AFF C80

To complete the controller replacement for your AFF C80 system, first restore the NetApp Storage Encryption configuration (if necessary). Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, return the failed part to NetApp.

Step 1: Verify LIFs and check cluster health

Before returning the replacement node to service, verify that the logical interfaces are on their home ports, check the cluster health, 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 logical interfaces are listed as false, revert them to their home ports:

   ```
   network interface revert -vserver * -lif *
   ```

2. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article.

3. If automatic giveback was disabled, reenable it:

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

Step 2: 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 C80

Replace a DIMM in your AFF C80 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 using one of the following options.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `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.

Option 2: Controller is in a MetroCluster

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.

- 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:

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

      `storage failover modify -node local -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:

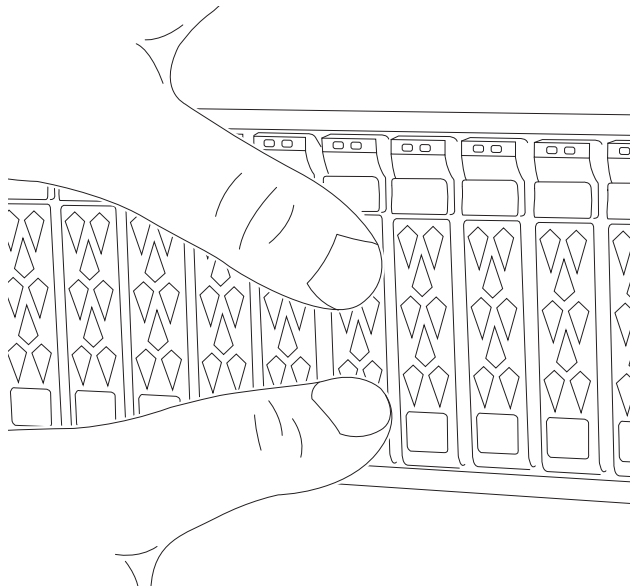| If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond y when prompted.

If the impaired controller is displaying…	Then…
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`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove 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>'
```
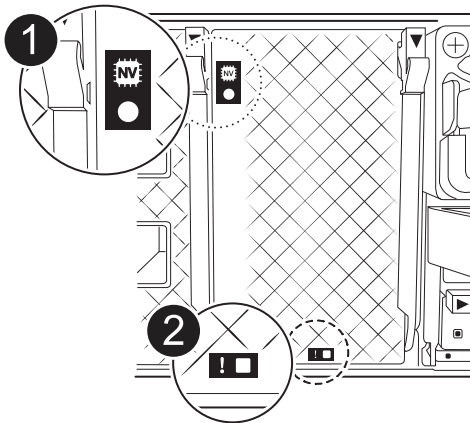
4. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

1	NVRAM status LED
2	NVRAM attention LED

-   If the NV LED is off, go to the next step.
-   If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

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

6. Unplug the controller module power supply cables from the controller module power supplies (PSU).
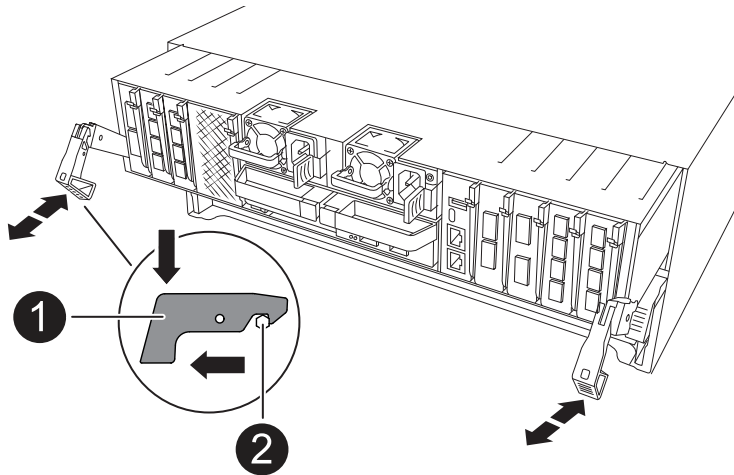
> ⓘ If your system has DC power, disconnect the power block from the PSUs.

7. 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.

8. Remove the cable management device from the controller module.

9. 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

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

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

Step 3: Replace the DIMM

To replace the DIMM, locate them inside the controller and follow the specific sequence of steps.

Steps

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

2. Open the controller air duct on the top of the controller.

   a. Insert your fingers in the recesses at the far ends of the air duct.

   b. Lift the air duct and rotate it upward as far as it will go.

3. Locate the DIMMs on your controller module and identify the target DIMM.

   Use the FRU map on the controller airduct to locate the DIMM slot.

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.

> **ℹ** Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.
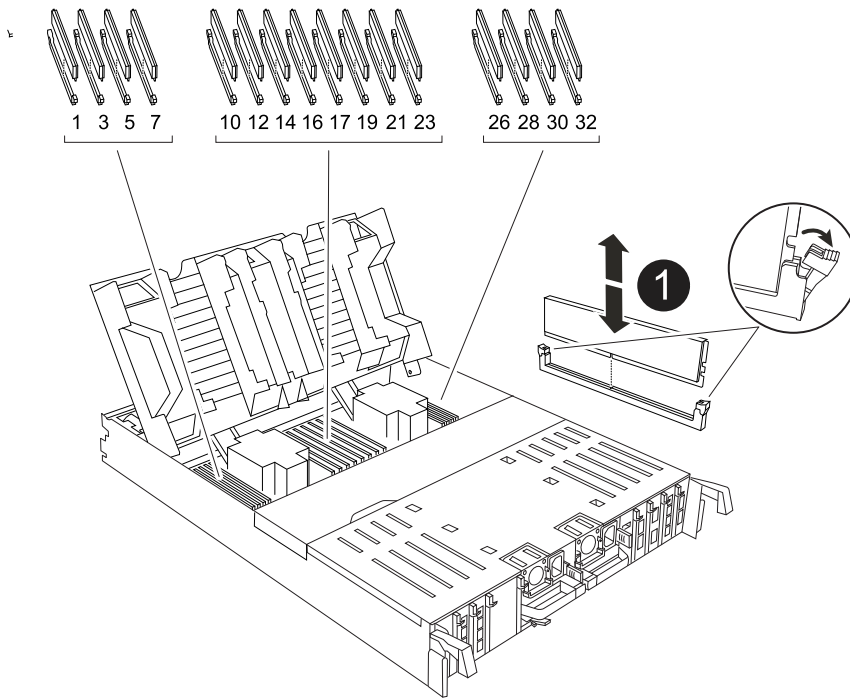
❶	DIMM and DIMM ejector tabs

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

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

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

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

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

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

8. Close the controller air duct.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

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

3. Recable the storage system, as needed.

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

   Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. 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 into the locked position.

5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

   If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

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

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

7. If automatic giveback was disabled, reenable it:

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

8. 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 - AFF C80

Replace a drive in your AFF C80 system when a drive fails or requires an upgrade. This 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 SSD drive nondisruptively while I/O is in progress.

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 disk 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.

  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 disk depends on how the disk 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 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 several disk drives, you must wait 70 seconds between the removal of each failed disk drive and the insertion of the replacement disk drive to allow the storage system to recognize the existence of each new disk.

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

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.

    > (i) 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.

Replace a fan module - AFF C80

Replace a fan module in your AFF C80 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 controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `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.

Option 2: Controller is in a MetroCluster

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.

- 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:

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

      ```
      storage failover modify -node local -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:

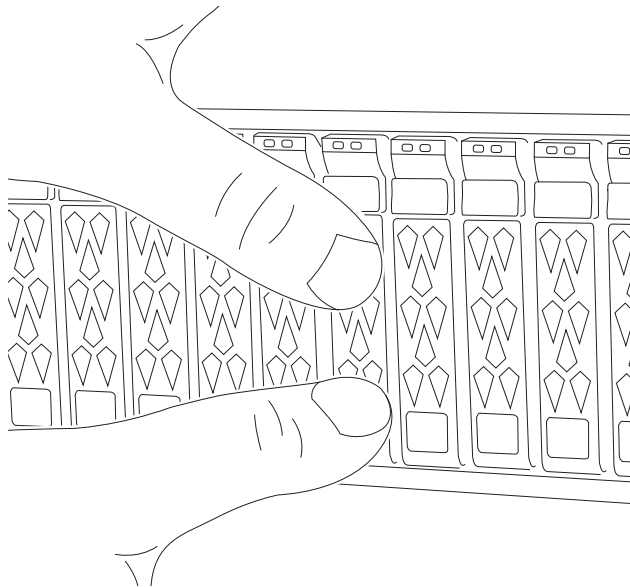| If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond y when prompted.

If the impaired controller is displaying…	Then…
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`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove 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>'
```
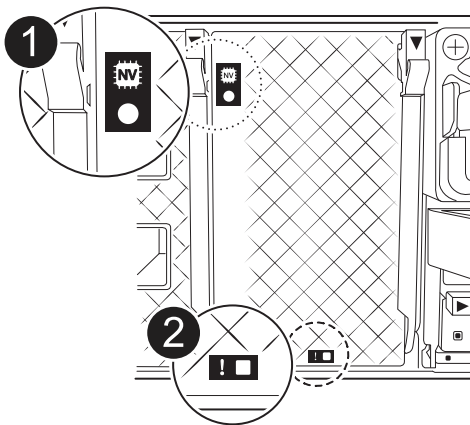
b.   Check for missing drive issues for both the file system or spare drives:

```
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.   Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

1	NVRAM status LED
2	NVRAM attention LED

- ◦ If the NV LED is off, go to the next step.
- ◦ If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

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

6. Unplug the controller module power supply cables from the controller module power supplies (PSU).

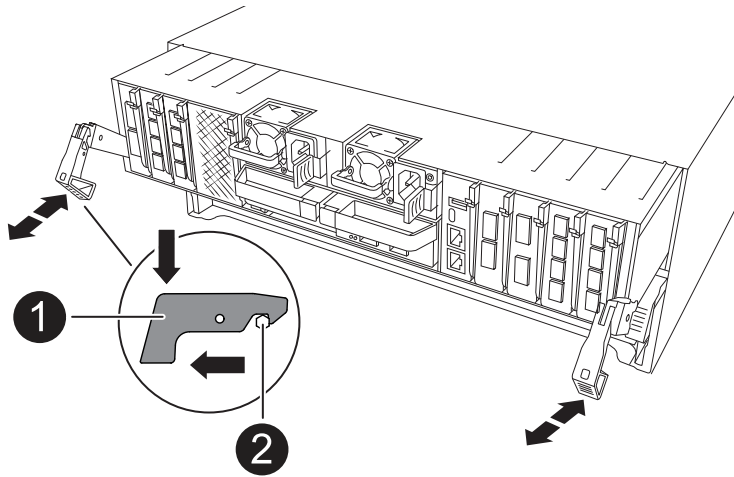> ⓘ   If your system has DC power, disconnect the power block from the PSUs.

7. 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.

8. Remove the cable management device from the controller module.

9. 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.

❶	Locking latch
❷	Locking pin

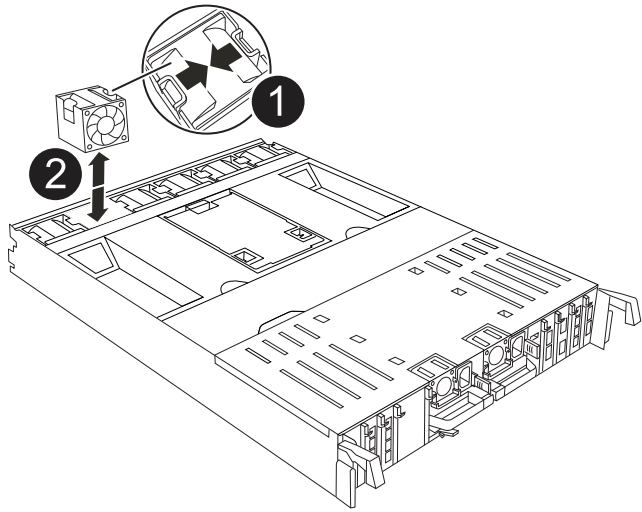10. Slide the controller module out of the chassis and place it on a flat, stable surface.

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

Step 3: Replace the fan

To replace a fan, remove 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.

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.

❶	Fan locking tabs
❷	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. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

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

3. Recable the storage system, as needed.

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

   Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. 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 into the locked position.

5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

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

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

7. If automatic giveback was disabled, reenable it:

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

8. 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 NVRAM - AFF C80

Replace the NVRAM in your AFF C80 system when the non-volatile memory becomes faulty or requires an upgrade. The replacement process involves shutting down the impaired controller, completely removing power by pulling the controller out a few inches, replacing the NVRAM module or the NVRAM DIMM, reassigning the disks, and returning the failed part to NetApp.

Before you begin
• Make sure you have the replacement part available. You must replace the failed component with a replacement component you received from NetApp.
• 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 using one of the following options.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `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.

Option 2: Controller is in a MetroCluster

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.

- 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:

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

      ```
      storage failover modify -node local -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 section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

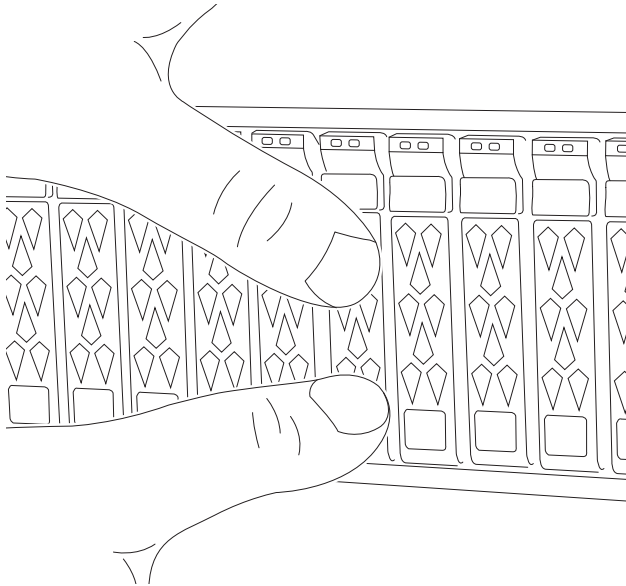| If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace the NVRAM module or NVRAM DIMM

Replace the NVRAM module or NVRAM DIMMs using the appropriate following option.

Option 1: Replace the NVRAM module

To replace the NVRAM module, locate it in slot 4/5 in the chassis and follow the specific sequence of 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. Go to the rear of the chassis.

5. Unplug the controller's PSUs.

> **ⓘ** If your system has DC power, disconnect the power block from the PSUs.

6. Press down on both of the locking latches on the controller, rotate both latches downward at the same time and pull the controller out about 3 or 4 inches.

7. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.

8. Remove the target NVRAM module from the chassis:

   a. Depress the cam latch button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch as far as it will go.

   c. Remove the impaired NVRAM module from the chassis by hooking your finger into the cam lever opening and pulling the module out of the chassis.
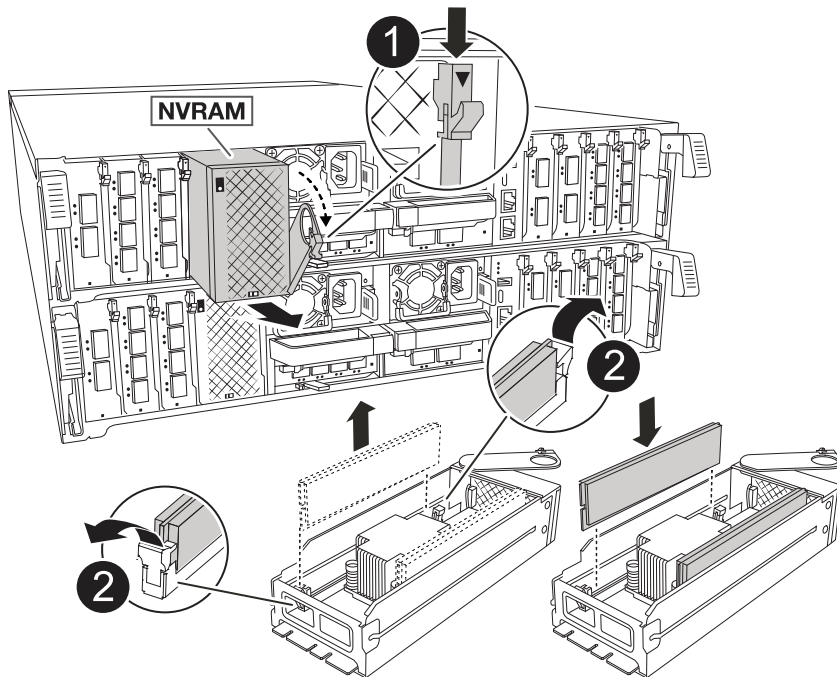
①	Cam locking button
②	DIMM locking tabs

9. Set the NVRAM module on a stable surface.

10. Remove the DIMMs, one at a time, from the impaired NVRAM module and install them in the replacement NVRAM module.

11. Install the replacement NVRAM module into the chassis:

    a. Align the module with the edges of the chassis opening in slot 4/5.

    b. Gently slide the module into the slot all the way, and then push the cam latch all the way up to lock the module in place.

12. Reinstall the controller module. The controller reboots as soon as it is fully inserted.

    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 into the locked position.

13. Plug the power cords into the power supplies.

    > ⓘ If you have DC power supplies, reconnect the power block to the power supplies.
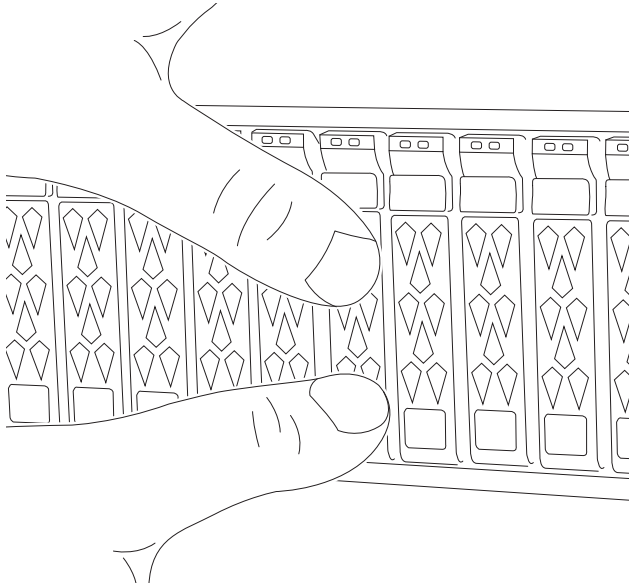
14. Rotate the cable management tray up to the closed position.

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

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

17. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Option 2: Replace the NVRAM DIMM

To replace NVRAM DIMMs in the NVRAM module, first must remove the NVRAM module and then replace the target DIMM.

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. Go to the rear of the chassis.

5. Unplug the controller's PSUs.

> ℹ️   If your system has DC power, disconnect the power block from the PSUs.

6. Press down on both of the locking latches on the controller, rotate both latches downward at the same time and pull the controller out about 3 or 4 inches.

7. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.

8. Remove the target NVRAM module from the chassis:

   a. Depress the cam button.

      The cam button moves away from the chassis.

   b. Rotate the cam latch as far as it will go.

   c. Remove the NVRAM module from the chassis by hooking your finger into the cam lever opening and pulling the module out of the chassis.
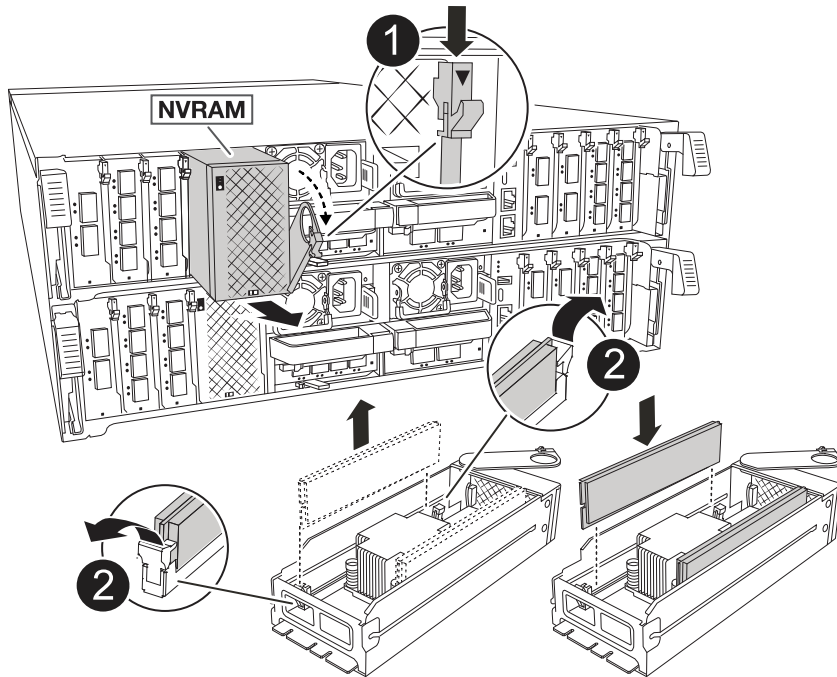
❶	Cam locking button
❷	DIMM locking tabs

9. Set the NVRAM module on a stable surface.

10. Locate the DIMM to be replaced inside the NVRAM module.

> ℹ️   Consult the FRU map label on the side of the NVRAM module to determine the locations of DIMM slots 1 and 2.

11. Remove the DIMM by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.

12. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.

13. Install the NVRAM module into the chassis:

    a. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.

14. Reinstall the controller module. The controller reboots as soon as it is fully inserted.

    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.

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

    b. Rotate the locking latches upward into the locked position.

15. Plug the power cords into the power supplies.

    (i) If you have DC power supplies, reconnect the power block to the power supplies.

16. Rotate the cable management tray up to the closed position.

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

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

19. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 3: Reassign disks

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

(!) Disk reassignment is only needed when replacing the NVRAM module and does not apply to NVRAM DIMM replacement.

Steps

1. If the 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 controller, boot the controller and enter *y* when prompted to override the system ID due to a system ID mismatch.

3. Wait until the Waiting for giveback message is displayed on the console of the controller with the replacement module 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, node 2 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. Give back the controller:

    a. From the healthy controller, give back the replaced controller's storage:
       ```
       storage failover giveback -ofnode replacement_node_name
       ```

       The 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.

       For more information, see manual giveback commands to override the veto.

    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.

5. Verify that the disks were assigned correctly:

   ```
   storage disk show -ownership
   ```

   The disks belonging to the controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 151759706:

```
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  -         151759706  151759706  -
151759706 Pool0
1.0.1  aggr0_1  node1 node1             151759706  151759706  -
151759706 Pool0
.
.
.
```

6. 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.

7. 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 controller is the current owner of the disks on the disaster site.

     See Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration for more information.

8. 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.
```

9. Verify that the expected volumes are present for each controller:

```
vol show -node node-name
```

10. If storage encryption is enabled, you must restore functionality.

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

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

12. If automatic giveback was disabled, reenable it:

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

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

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

Step 4: Return the failed part to NetApp

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

Replace the NV battery - AFF C80

Replace the NV battery in your AFF C80 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.

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 using one of the following options.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `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.

Option 2: Controller is in a MetroCluster

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.

- 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:

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

   ```
   storage failover modify -node local -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:

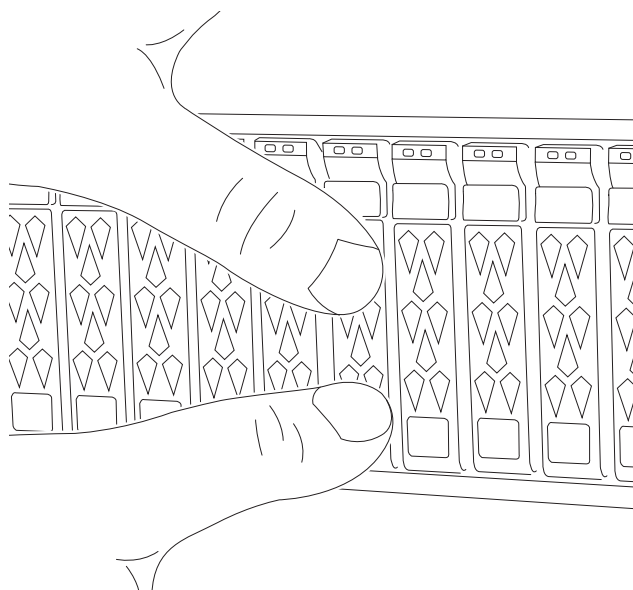| If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…

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 -halt true`
	
	The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove 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 .

- 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. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

❶	NVRAM status LED
❷	NVRAM attention LED

- If the NV LED is off, go to the next step.

- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

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

6. Unplug the controller module power supply cables from the controller module power supplies (PSU).

> ⓘ   If your system has DC power, disconnect the power block from the PSUs.

7. 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.

8. Remove the cable management device from the controller module.

9. 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.

❶	Locking latch
❷	Locking pin

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

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

Step 3: Replace the NV battery

Remove the failed NV battery from the controller module and install the replacement NV battery.

Steps

1. Open the air duct cover and locate the NV battery.
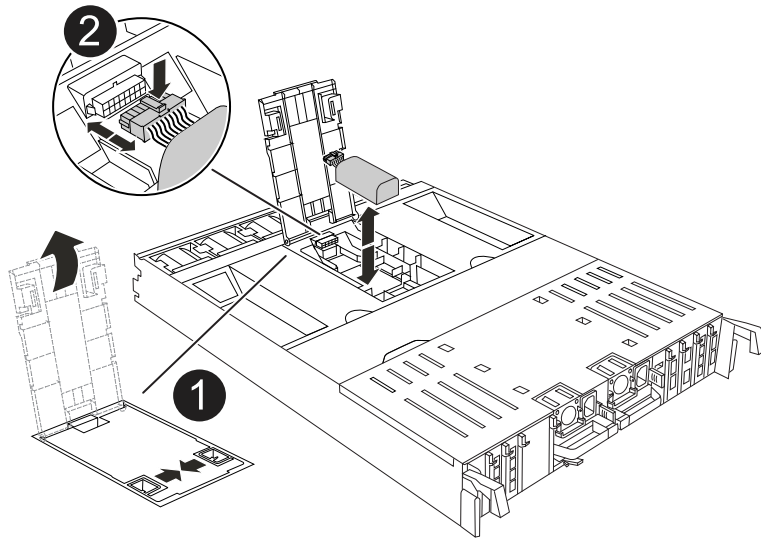
❶	NV battery air duct cover
❷	NV battery plug

2. Lift the battery up to access the battery plug.

3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.

4. Lift the battery out of the air duct and controller module, and then set it aside.

5. Remove the replacement battery from its package.

6. Install the replacement battery pack into the controller:

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

    b. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.

7. Close the NV air duct cover.

   Make sure that the plug locks into the socket.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

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

3. Recable the storage system, as needed.

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

    Make sure that the console cable is connected to the repaired controller module so that it receives console messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. 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 into the locked position.

5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

    If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

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

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

7. If automatic giveback was disabled, reenable it:

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

8. 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.

I/O module

Overview of add and replace an I/O module - AFF C80

The AFF C80 system offers flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your AFF C80 storage system with the same type of I/O module, or with a different kind of I/O module. You can also add an I/O module into a system with empty slots.

- Add an I/O module

  Adding additional modules can improve redundancy, helping to ensure that the system remains operational even if one module fails.

- Hot swap an I/O module

  Hot swapping an I/O module allows you to replace a failed module without shutting down the system, minimizing downtime and maintaining system availability.

- Replace an I/O module

  Replacing a failing I/O module can restore the system to its optimal operating state.

I/O slot numbering

The I/O slots on AFF C80 controllers are numbered 1 through 11, as shown in the following illustration.
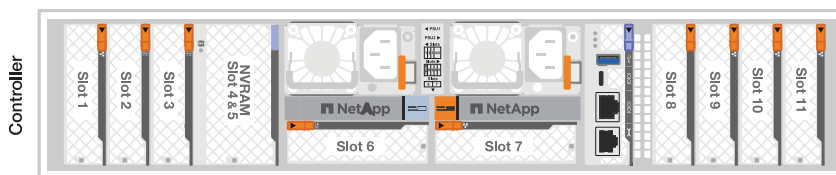
Add an I/O module - AFF C80

Add an I/O module to your AFF C80 system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your system when there are empty slots available or when all slots are fully populated.

About this task

If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led` *on* command.

A storage system has two location LEDs, one on each controller. Location LEDs remain illuminated for 30 minutes.

You can turn them off by entering the `system location-led` *off* command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Shut down the impaired controller module

Shut down or take over the impaired controller module using one of the following options.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `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.

Option 2: MetroCluster configuration

> (i) 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 command:
`system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command 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.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Halt or take over 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`.

Step 2: Add the new I/O module

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- Make sure that all other components are functioning properly.
- Make sure you have the replacement component you received from NetApp.

Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

Steps

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

2. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.

3. Remove the target slot blanking module from the carrier:

    a. Depress the cam latch on the blanking module in the target slot.

    b. Rotate the cam latch away from the module as far as it will go.

    c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

4. Install the I/O module:

    a. Align the I/O module with the edges of the enclosure slot opening.

    b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.

5. Cable the I/O module to the designated device.

    > ⓘ Make sure that any unused I/O slots have blanks installed to prevent possible thermal issues.

6. Rotate the cable management tray up to the closed position.

7. From the LOADER prompt, reboot the node:

    ```
    bye
    ```

    > ⓘ This reinitializes the I/O module and other components and reboots the node.

8. Give back the controller from the partner controller:

    ```
    storage failover giveback -ofnode target_node_name
    ```

9. Repeat these steps for controller B.

10. From the healthy node, restore automatic giveback if you disabled it:

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

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

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

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

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

2. Unplug any cabling on the target I/O module.

3. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.

4. Remove the target I/O module from the chassis:

   a. Depress the cam latch button.

   b. Rotate the cam latch away from the module as far as it will go.

   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

   Make sure that you keep track of which slot the I/O module was in.

5. Install the I/O module into the target slot in the enclosure:

   a. Align the module with the edges of the enclosure slot opening.

   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.

6. Cable the I/O module to the designated device.

7. Repeat the remove and install steps to replace additional modules for the controller.

8. Rotate the cable management tray up to the closed position.

9. Reboot the controller from the LOADER prompt:_bye_

   This reinitializes the PCIe cards and other components and reboots the node.

   > ⓘ If you encounter an issue during reboot, see BURT 1494308 - Environment shutdown might be triggered during I/O module replacement

10. Give back the controller from the partner controller:

    `storage failover giveback -ofnode target_node_name`

11. Enable automatic giveback if it was disabled:

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

12. Do one of the following:

    ◦ If you removed a NIC I/O module and installed a new NIC I/O module, use the following network command for each port:

      ```
      storage port modify -node *<node name> -port *<port name> -mode network
      ```

    ◦ If you removed a NIC I/O module and installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.

13. Repeat these steps for controller B.

Hot swap an I/O module - AFF C80

You can hot swap an Ethernet I/O module in your AFF C80 storage system if a module fails and your storage system meets all ONTAP version requirements.

To hot swap an I/O module, make sure your storage system is running ONTAP 9.18.1 GA or later, prepare your storage system and I/O module, hot-swap the failed module, bring the replacement module online, restore the storage system to normal operation, and return the failed module to NetApp.

About this task

- You do not need to perform a manual takeover before replacing the failed I/O module.

- Apply commands to the correct controller and I/O slot during the hot-swap:

    ◦ The *impaired controller* is the controller where you are replacing the I/O module.

    ◦ The *healthy controller* is the HA partner of the impaired controller.

- You can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

  The storage system includes three location LEDs: one on the operator display panel and one on each controller. The LEDs remain illuminated for 30 minutes.

  You can turn them off by entering the `system location-led off` command. If you are unsure if the LEDs are on or off, you can check their state by entering the `system location-led show` command.

Step 1: Ensure the storage system meets the procedure requirements

To use this procedure, your storage system must be running ONTAP 9.18.1 GA or later, and your storage system must meet all requirements.

> ⓘ If your storage system is not running ONTAP 9.18.1 GA or later, you cannot use this procedure, you must use the replace an I/O module procedure.

- You are hot swapping an Ethernet I/O module in any slot having any combination of ports used for cluster, HA, and client with an equivalent I/O module. You cannot change the I/O module type.

  Ethernet I/O modules with ports used for storage or MetroCluster are not hot-swappable.

- Your storage system (switchless or switched cluster configuration) can have any number of nodes

supported for your storage system.

- All nodes in the cluster must be running the same ONTAP version (ONTAP 9.18.1GA or later) or running different patch levels of the same ONTAP version.

  If nodes in your cluster are running different ONTAP versions, this is considered a mixed-version cluster and hot-swapping an I/O module is not supported.

- The controllers in your storage system can be in either of the following states:

  - Both controllers can be up and running I/O (serving data).

  - Either controller can be in a takeover state if the takeover was caused by the failed I/O module and the nodes are otherwise functioning properly.

    In certain situations, ONTAP can automatically perform a takeover of either controller due to the failed I/O module. For example, if the failed I/O module contained all of the cluster ports (all of the cluster links on that controller go down) ONTAP automatically performs a takeover.

- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Step 2: Prepare the storage system and I/O module slot

Prepare the storage system and I/O module slot so that it is safe to remove the failed I/O module:

Steps
1. Properly ground yourself.

2. Label the cables to identify where they came from, and then unplug all cables from the target I/O module.

   > ⓘ The I/O module should be failed (ports should be in the link down state); however, if the links are still up and they contain the last functioning cluster port, unplugging the cables triggers an automatic takeover.
   >
   > Wait five minutes after unplugging the cables to ensure any takeovers or LIF failovers complete before continuing with this procedure.

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

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

   For example, the following AutoSupport message suppresses automatic case creation for two hours:

   ```
   node2::> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

4. Disable automatic giveback if the partner node has been taken over:

If…	Then…
If either controller took over its partner automatically	Disable automatic giveback: 1. Enter the following command from the console of the controller that took over its partner: `storage failover modify -node local -auto -giveback false` 2. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
Both controllers are up and running I/O (serving data)	Go to the next step.

5. Prepare the failed I/O module for removal by removing it from service and powering it off:

   a. Enter the following command:

   ```
   system controller slot module remove -node impaired_node_name -slot
   slot_number
   ```

   b. Enter `y` when you see the prompt *Do you want to continue?*

   For example, the following command prepares the failed module in slot 7 on node 2 (the impaired controller) for removal, and displays a message that it is safe to remove:

   ```
   node2::> system controller slot module remove -node node2 -slot 7

   Warning: IO_2X_100GBE_NVDA_NIC module in slot 7 of node node2 will be
   powered off for removal.

   Do you want to continue? {y|n}: y

   The module has been successfully removed from service and powered
   off. It can now be safely removed.
   ```

6. Verify the failed I/O module is powered off:

   ```
   system controller slot module show
   ```

   The output should show `powered-off` in the `status` column for the failed module and its slot number.

Step 3: Replace the failed I/O module

Replace the failed I/O module with an equivalent I/O module.

Steps

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

2. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.

3. Remove the I/O module from the controller module:

> ⓘ The following illustration shows removing a horizontal and vertical I/O module. Typically, you will only remove one I/O module.
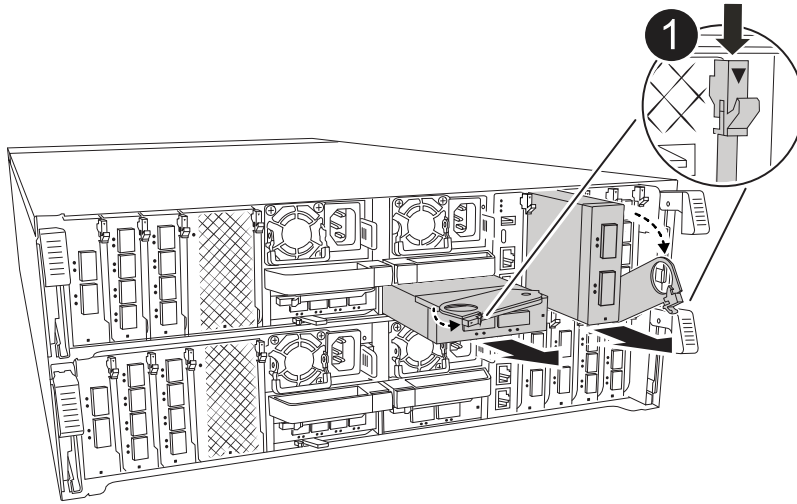
❶	Cam locking button

   a. Depress the cam latch button.

   b. Rotate the cam latch away from the module as far as it will go.

   c. Remove the module from the controller module by hooking your finger into the cam lever opening and pulling the module out of the controller module.

   Keep track of which slot the I/O module was in.

4. Set the I/O module aside.

5. Install the replacement I/O module into the target slot:

   a. Align the I/O module with the edges of the slot.

   b. Gently slide the module into the slot all the way into the controller module, and then rotate the cam latch all the way up to lock the module in place.

6. Cable the I/O module.

7. Rotate the cable management tray into the locked position.

Step 4: Bring the replacement I/O module online

Bring the replacement I/O module online, verify the I/O module ports initialized successfully, verify the slot is powered on, and then verify the I/O module is online and recognized.

About this task

After the I/O module is replaced and the ports are returned to a healthy state, LIFs are reverted to the replaced

I/O module.

Steps

1. Bring the replacement I/O module online:

   a. Enter the following command:

      ```
      system controller slot module insert -node impaired_node_name -slot
      slot_number
      ```

   b. Enter y when you see the prompt, *Do you want to continue?*

      The output should confirm the I/O module was successfully brought online (powered on, initialized, and placed into service).

      For example, the following command brings slot 7 on node 2 (the impaired controller) online, and displays a message that the process was successful:

      ```
      node2::> system controller slot module insert -node node2 -slot 7

      Warning: IO_2X_100GBE_NVDA_NIC module in slot 7 of node node2 will be
      powered on and initialized.

      Do you want to continue? {y|n}: `y`

      The module has been successfully powered on, initialized and placed
      into service.
      ```

2. Verify that each port on the I/O module successfully initialized:

   a. Enter the following command from the console of the impaired controller:

      ```
      event log show -event *hotplug.init*
      ```

      (i)   It might take several minutes for any required firmware updates and port initialization.

      The output should show one or more hotplug.init.success EMS events and `hotplug.init.success:` in the `Event` column, indicating each port on the I/O module initialized successfully.

      For example, the following output shows initialization succeeded for I/O ports e7b and e7a:

```
node2::> event log show -event *hotplug.init*

Time                      Node              Severity      Event

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

7/11/2025 16:04:06  node2       NOTICE        hotplug.init.success:
Initialization of ports "e7b" in slot 7 succeeded

7/11/2025 16:04:06  node2       NOTICE        hotplug.init.success:
Initialization of ports "e7a" in slot 7 succeeded

2 entries were displayed.
```

    b. If the port initialization fails, review the EMS log for the next steps to take.

3. Verify the I/O module slot is powered on and ready for operation:

```
system controller slot module show
```

The output should show the slot status as *powered-on* and therefore ready for operation of the I/O module.

4. Verify that the I/O module is online and recognized.

Enter the command from the console of the impaired controller:

```
system controller config show -node local -slot slot_number
```

If the I/O module was successfully brought online and is recognized, the output shows I/O module information, including port information for the slot.

For example, you should see output similar to the following for an I/O module in slot 7:

```
node2::> system controller config show -node local -slot 7

Node: node2
Sub- Device/
Slot slot Information
---- ---- ----------------------------
   7      - Dual 40G/100G Ethernet Controller CX6-DX
                  e7a MAC Address: d0:39:ea:59:69:74 (auto-100g_cr4-fd-
up)
                        QSFP Vendor:        CISCO-BIZLINK
                        QSFP Part Number:   L45593-D218-D10
                        QSFP Serial Number: LCC2807GJFM-B
                  e7b MAC Address: d0:39:ea:59:69:75 (auto-100g_cr4-fd-
up)
                        QSFP Vendor:        CISCO-BIZLINK
                        QSFP Part Number:   L45593-D218-D10
                        QSFP Serial Number: LCC2809G26F-A
                  Device Type:        CX6-DX PSID(NAP0000000027)
                  Firmware Version:   22.44.1700
                  Part Number:        111-05341
                  Hardware Revision:  20
                  Serial Number:      032403001370
```

Step 5: Restore the storage system to normal operation

Restore your storage system to normal operation by giving back storage to the controller that was taken over (as needed), restoring automatic giveback (as needed), verifying LIFs are on their home ports, and reenabling AutoSupport automatic case creation.

Steps

1. As needed for the version of ONTAP your storage system is running and the state of the controllers, give back storage and restore automatic giveback on the controller that was taken over:

If…	Then…

If either controller took over its partner automatically	1. Return the controller that was taken over to normal operation by giving back its storage:
	
	`storage failover giveback -ofnode controller that was taken over_name`
	
	2. Restore automatic giveback from the console of the controller that was taken over:
	
	`storage failover modify -node local -auto -giveback true`

If…	Then…
Both controllers are up and running I/O (serving data)	Go to the next step.

2. Verify that the logical interfaces are reporting to their home node 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 *`

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

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

Step 6: Return the failed part to NetApp

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

Replace an I/O module - AFF C80

Replace an I/O module in your AFF C80 system when the module fails or 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. You can use this procedure with all versions of ONTAP supported by your storage system.

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

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `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.

Option 2: Controller is in a MetroCluster

> (i) 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 command:
   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command 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.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Halt or take over 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`.

Step 2: Replace a failed I/O module

To replace an I/O module, locate it within the controller module and follow the specific sequence of steps.

Steps

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

2. Label the cables to identify where they came from, and then unplug all cables from the target I/O module.

3. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.

4. Remove the I/O module from the controller module:

   > ⓘ The following illustration shows removing a horizontal and vertical I/O module. Typically, you will only remove one I/O module.
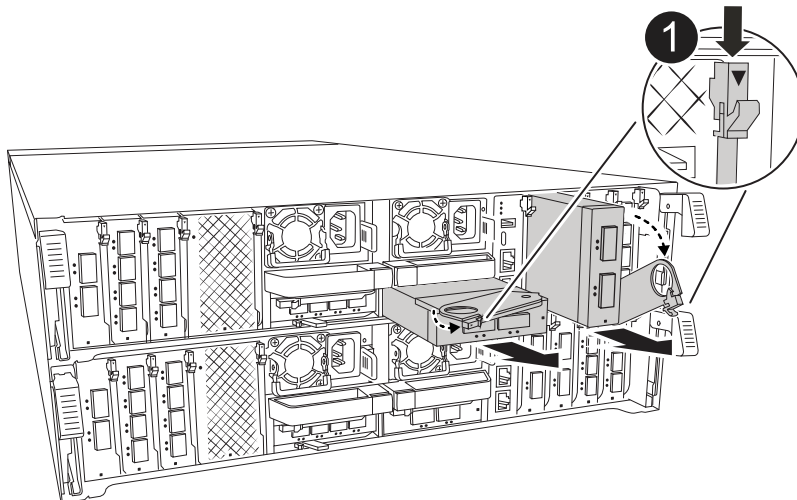
1	Cam locking button

   a. Depress the cam latch button.

   b. Rotate the cam latch away from the module as far as it will go.

   c. Remove the module from the controller module by hooking your finger into the cam lever opening and pulling the module out of the controller module.

Keep track of which slot the I/O module was in.

5. Set the I/O module aside.

6. Install the replacement I/O module into the target slot:

   a. Align the I/O module with the edges of the slot.

   b. Gently slide the module into the slot all the way into the controller module, and then rotate the cam latch all the way up to lock the module in place.

7. Cable the I/O module.

8. Rotate the cable management tray into the locked position.

Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller.

Steps

1. Reboot the controller from the LOADER prompt:

   ```
   bye
   ```

   > **(i)** Rebooting the impaired controller also reinitializes the I/O modules and other components.

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

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

3. Restore automatic giveback from the console of the healthy controller:

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

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

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

Step 4: Return the failed part to NetApp

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

Replace a power supply - AFF C80

Replace an AC or DC power supply unit (PSU) in your AFF C80 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.

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.

> ⓘ   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.

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.

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.
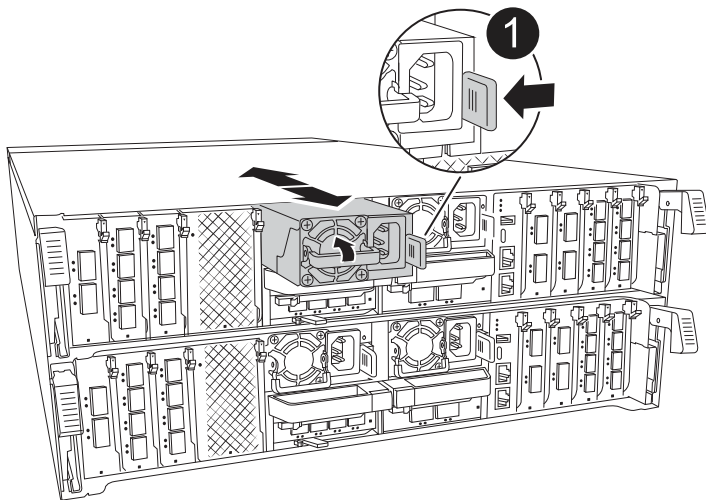
❶	Terracotta PSU locking tab

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 PSU cabling:

    a. Reconnect the power cable to the PSU.

    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 page for further information.

Option 2: Hot-swap a DC PSU

To replace a DC PSU, complete the following steps.

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.
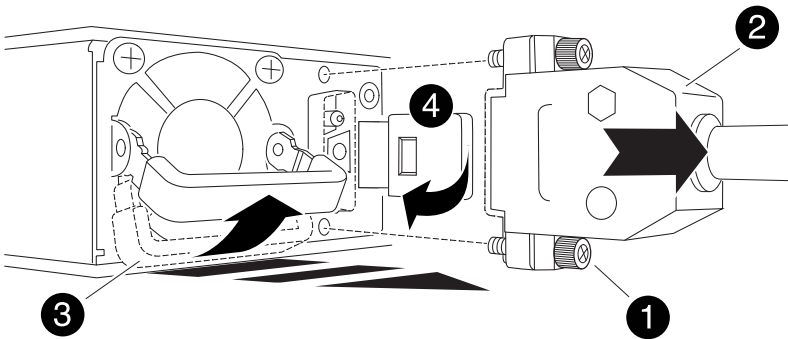
❶	Thumb screws
❷	D-SUB DC power PSU cable connector
❸	Power supply handle
❹	Blue PSU locking tab

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 C80

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your AFF C80 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 using one of the following options.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `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.

Option 2: Controller is in a MetroCluster

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.

- 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:

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

      ```
      storage failover modify -node local -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 section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
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`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove 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 .

- 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. Check the amber NVRAM that status LED located in slot 4/5 on the back of the impaired controller module is off. Look for the NV icon.

1	NVRAM status LED
2	NVRAM attention LED

- If the NV LED is off, go to the next step.
- If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

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

6. Unplug the controller module power supply cables from the controller module power supplies (PSU).

> ⓘ If your system has DC power, disconnect the power block from the PSUs.

7. 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.

8. Remove the cable management device from the controller module.

9. 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.

❶	Locking latch
❷	Locking pin

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

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

Step 3: Replace the RTC battery

Remove the failed RTC battery and install the replacement RTC battery.

You must use an approved RTC battery.

Steps

1. Open the controller air duct on the top of the controller.

   a. Insert your fingers in the recesses at the far ends of the air duct.

   b. Lift the air duct and rotate it upward as far as it will go.

2. Locate the RTC battery under the air duct.
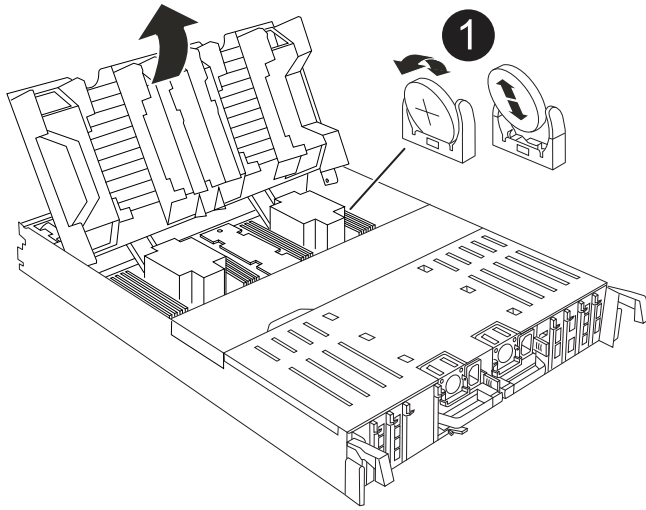
1	RTC battery and housing

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.

Step 4: Reinstall the controller module

Reinstall the controller module and reboot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

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

3. Recable the storage system, as needed.

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

   Make sure that the console cable is connected to the repaired controller module so that it receives console

messages when it reboots. The repaired controller receives power from the healthy controller and begins to reboot as soon as it is seated completely in the chassis.

4. 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 into the locked position.

5. Plug the power cords into the power supplies. The controller reboots as soon as power is restored.

   If you have DC power supplies, reconnect the power block to the power supplies after the controller module is fully seated in the chassis.

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

   `storage failover giveback -ofnode` *impaired_node_name*.

7. If automatic giveback was disabled, reenable it:

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

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

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

Step 5: Reset the time and date on the controller

After you replace the RTC battery, insert the controller, and power on for the first BIOS reset, you will see the following error messages:

`RTC date/time error. Reset date/time to default`

`RTC power failure error`

These messages are expected and you can continue with this procedure.

Steps

1. Check the date and time on the healthy controller with the `cluster date show` command.
   If your system stops at the boot menu, select the option for `Reboot node` and respond *y* when prompted, then boot to LOADER by pressing *Ctrl-C*

   a. At the LOADER prompt on the target controller, check the time and date with the `cluster date show` command.

   b. If necessary, modify the date with the `set date mm/dd/yyyy` command.

   c. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.

2. Confirm the date and time on the target controller.

3. At the LOADER prompt, enter *bye* to reinitialize the PCIe cards and other components and let the controller

reboot.

Step 6: Return the failed part to NetApp

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

Replace system management module - AFF C80

Replace the System Management module in your AFF C80 system when it becomes defective or its firmware is corrupted. The replacement process involves shutting down the controller, replacing the failed System Management module, rebooting the controller, updating the license keys, and returning the failed part to NetApp.

Replace the impaired system management module.

Steps

1. 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.
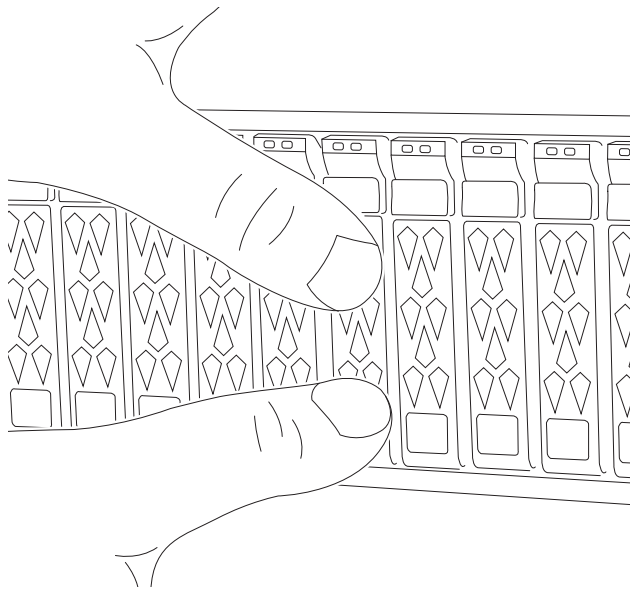
2. Make sure NVRAM destage has completed before proceeding. When the LED on the NV module is off, NVRAM is destaged. If the LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.
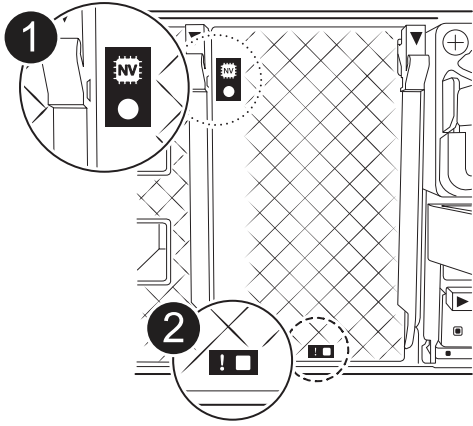
❶	NVRAM status LED
❷	NVRAM attention LED

- ◦ If the NV LED is off, go to the next step.
- ◦ If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.

3. Go to the rear of the chassis. If you are not already grounded, properly ground yourself.

4. Unplug the controller's PSUs.

> ℹ️ If your system has DC power, disconnect the power block from the PSUs.

5. Press down on both of the locking latches on the controller, rotate both latches downward at the same time and pull the controller out about 3 or 4 inches.

6. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.

7. Remove any cables connected to the System Management module. Make sure that label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
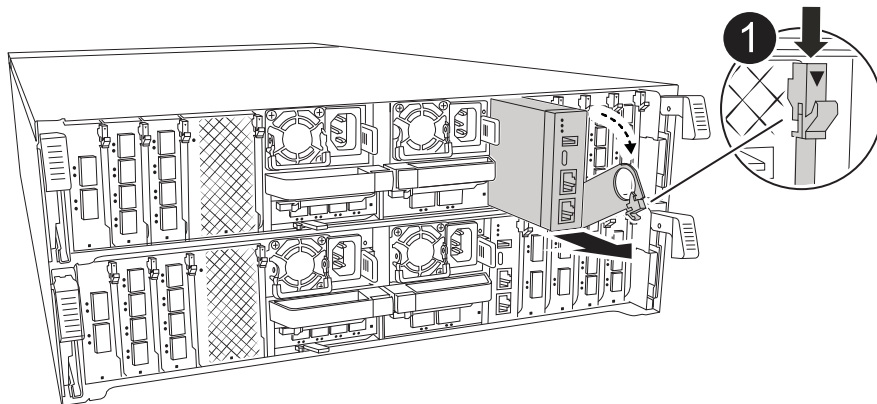
❶	System Management module cam latch

8. Remove the System Management module:

    a. Depress the system management cam button. The cam lever moves away from the chassis.

    b. Rotate the cam lever all the way down.

    c. Loop your finger into the cam lever and pull the module straight out of the system.

    d. Place the System Management module on an anti-static mat, so that the boot media is accessible.

9. Move the boot media to the replacement System Management module:
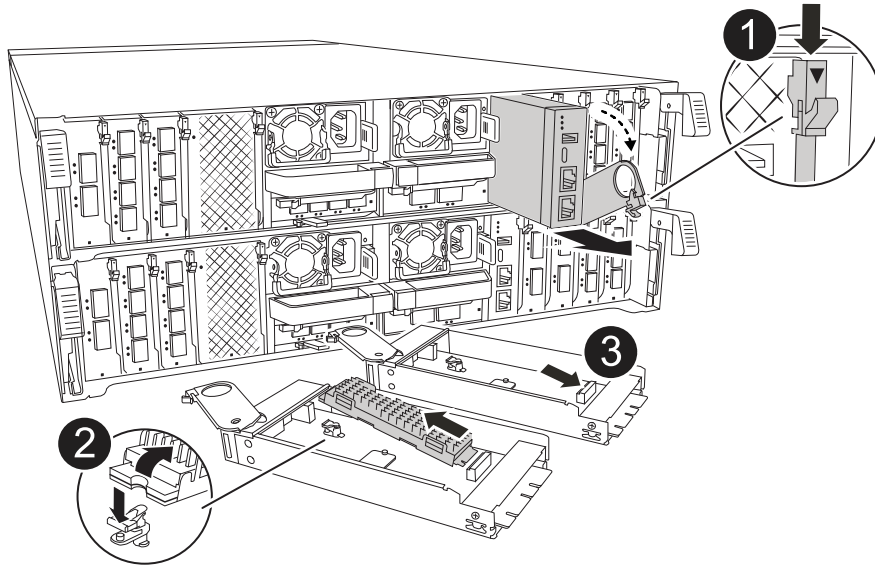
❶	System Management module cam latch
❷	Boot media locking button
❸	Boot media

    a. Press the blue locking button. The boot media rotates slightly upward.

    b. Rotate the boot media up, slide it out of the socket.

    c. Install the boot media in the replacement System Management module:

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

        ii. Rotate the boot media down toward until it engages the locking button. Depress the blue locking if necessary.

10. Install the system management module:

    a. Align the edges of the replacement System Management module with the system opening and gently push it into the controller module.

    b. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.

11. Recable the System Management module.

12. Reinstall the controller module. The controller reboots as soon as it is fully inserted.

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.

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

b. Rotate the locking latches upward into the locked position.

13. Plug the power cords into the power supplies.

(i) If you have DC power supplies, reconnect the power block to the power supplies.

14. Rotate the cable management tray up to the closed position.