AFF A150 systems

Install and maintain

NetApp
April 23, 2026

This PDF was generated from https://docs.netapp.com/us-en/ontap-systems/a150/install-setup.html on April 23, 2026. Always check docs.netapp.com for the latest.

Table of Contents

AFF A150 systems . . . 1
- Install and setup . . . 1
  - Start here: Choose your installation and setup experience . . . 1
  - Quick guide - AFF A150 . . . 1
  - Video steps - AFF A150 . . . 1
  - Detailed guide - AFF A150 . . . 2
- Maintain . . . 14
  - Maintain AFF A150 hardware . . . 14
  - Boot media . . . 15
  - Chassis . . . 35
  - Controller . . . 44
  - Replace a DIMM - AFF A150 . . . 62
  - Replace SSD Drive or HDD Drive - AFF A150 . . . 70
  - Replace the NVMEM battery - AFF A150 . . . 74
  - Hot-swap a power supply - AFF A150 . . . 80
  - Replace the real-time clock battery - AFF A150 . . . 82
- Key specifications for AFF A150 . . . 88

AFF A150 systems

Install and setup

Start here: Choose your installation and setup experience

For most configurations, you can choose from different content formats.

- Quick steps

  A printable PDF of step-by-step instructions with live links to additional content.

- Video steps

  Video step-by-step instructions.

- Detailed steps

  Online step-by-step instructions with live links to additional content.

If your system is in a MetroCluster IP configuration, see the Install MetroCluster IP Configuration instructions.

Warning: If your system is being installed with ONTAP 9.13.1P8 - 9.13.1P11, ONTAP 9.14.1P1 - 9.14.1P7 or ONTAP 9.15.1 - 9.15.1P2 and your system contains 10 or more internal solid-state drives, you must take additional steps to prepare the system for installation. See issue CONTAP-285173 - ADP does not leave a spare root partition on an AFF A150 with 10 or more internal drives.

Quick guide - AFF A150

Warning: If your system is being installed with ONTAP 9.13.1P8 - 9.13.1P11, ONTAP 9.14.1P1 - 9.14.1P7 or ONTAP 9.15.1 - 9.15.1P2 and your system contains 10 or more internal solid-state drives, you must take additional steps to prepare the system for installation. See issue CONTAP-285173 - ADP does not leave a spare root partition on an AFF A150 with 10 or more internal drives.

The Installation and Setup instructions give graphic instructions for a typical installation of your system, from racking and cabling, through initial system bring-up. Use this guide if you are familiar with installing NetApp systems.

Use the xref:./a150/AFF A150 System Installation and Setup Instructions

The ASA A150 uses the same installation procedure as the AFF A150 system.

Video steps - AFF A150

The following video shows how to install and cable your system.

Animation - Install and setup of an AFF A150

If you have a MetroCluster configuration, use the MetroCluster documentation.

Warning: If your system is being installed with ONTAP 9.13.1P8 - 9.13.1P11, ONTAP 9.14.1P1 - 9.14.1P7 or ONTAP 9.15.1 - 9.15.1P2 and your system contains 10 or more internal solid-state drives, you must take additional steps to prepare the system for installation. See issue CONTAP-285173 - ADP does not leave a spare root partition on an AFF A150 with 10 or more internal drives.

Detailed guide - AFF A150

This section gives detailed step-by-step instructions for installing an AFF A150 system.

If you have a MetroCluster configuration, use the MetroCluster documentation.

Warning: If your system is being installed with ONTAP 9.13.1P8 - 9.13.1P11, ONTAP 9.14.1P1 - 9.14.1P7 or ONTAP 9.15.1 - 9.15.1P2 and your system contains 10 or more internal solid-state drives, you must take additional steps to prepare the system for installation. See issue CONTAP-285173 - ADP does not leave a spare root partition on an AFF A150 with 10 or more internal drives.

Step 1: Prepare for installation

To install your AFF A150 system, you create an account on the NetApp Support Site, register your system, and obtain your license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

Before you begin

- Make sure you have access to NetApp Hardware Universe (HWU) for information about site requirements as well as additional information on your configured system.
- Make sure you have access to the Release Notes for your version of ONTAP for more information about this system.
- Contact your network administrator for information about connecting your system to the switches.
- Make sure you have the following items at your site:
  - Rack space for the storage system
  - Phillips #2 screwdriver
  - Additional networking cables to connect your system to your network switch and laptop or console with a Web browser
  - A laptop or console with an RJ-45 connection and access to a Web browser

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.



3. Set up your account:
   a. Log in to your existing account or create an account.
   b. Register your system.
4. Download and install Config Advisor on your laptop.
5. Inventory and make a note of the number and types of cables you received.

The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see NetApp Hardware Universe to locate the cable and identify its use.

Type of cable…	Part number and length	Connector type	For…
10 GbE cable (order dependent)	X6566B-05-R6 (112-00297), 0.5m X6566B-2-R6 (112-00299), 2m		Cluster interconnect network
10 GbE cable (order dependent)	Part number X6566B-2-R6 (112-00299), 2m or X6566B-3-R6 (112-00300), 3m X6566B-5-R6 (112-00301), 5m		Data
Optical network cables (order dependent)	X6553-R6 (112-00188), 2m X6536-R6 (112-00090), 5m X6554-R6(112-00189), 15m		FC host network
Cat 6, RJ-45 (order dependent)	Part numbers X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network and Ethernet data
Storage (order dependent)	Part number X66030A (112-00435), 0.5m X66031A (112-00436), 1m X66032A (112-00437), 2m X66033A (112-00438), 3m		Storage
Micro-USB console cable	Not applicable		Console connection during software setup on non-Windows or Mac laptop/console
Power cables	Not applicable		Powering up the system

6. Download and complete the Cluster Configuration Worksheet.

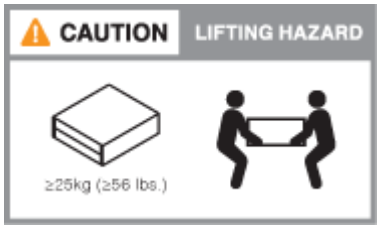
Step 2: Install the hardware

You install your system in a 4-post rack or NetApp system cabinet, as applicable.

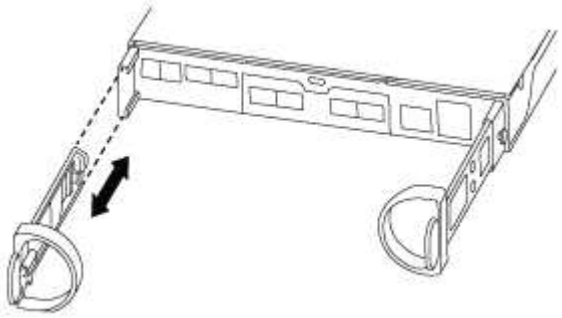
Steps

1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

You need to be aware of the safety concerns associated with the weight of the system.



3. Attach cable management devices (as shown).



4. Place the bezel on the front of the system.

Step 3: Cable controllers to network

You cable the controllers to your network by using either the two-node switchless cluster method or the switched cluster method.

About this task

The following table identifies the cable type with the call out number and cable color in the illustrations for both two-node switchless cluster network cabling and switched cluster network cabling.

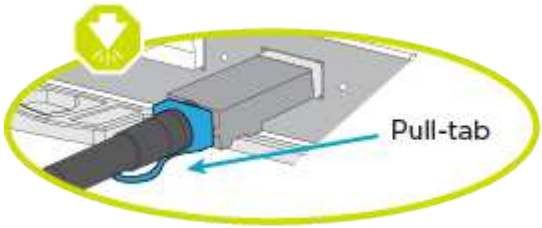
Cabling	Connection type
1	Cluster interconnect
2	Controllers to host data network switches
3	Controllers to management network switch

Option 1: Two-node switchless cluster

Cable your two-node switchless cluster.

About this task

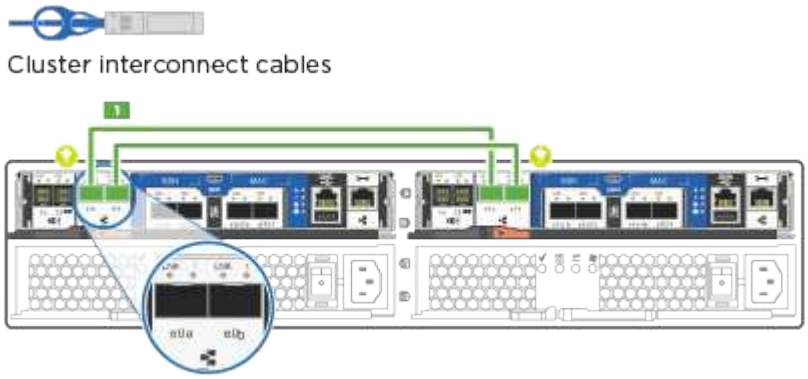
Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.



As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Cable the cluster interconnect ports e0a to e0a and e0b to e0b with the cluster interconnect cable.

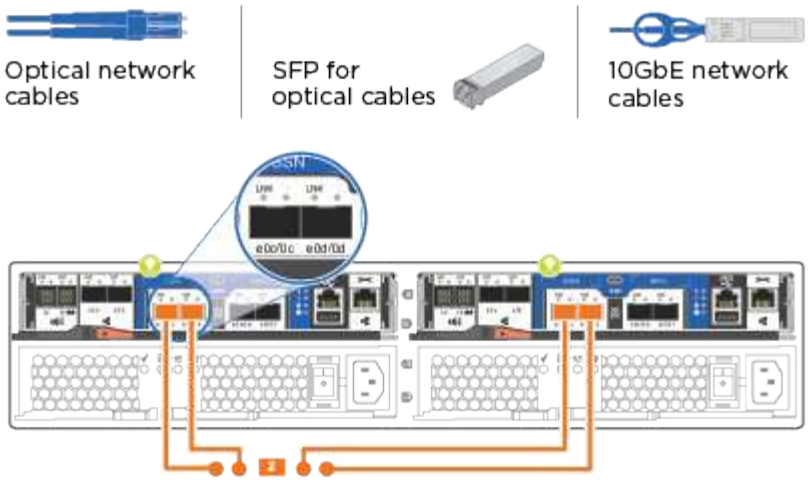


2. Cable the controllers to either a UTA2 data network or an Ethernet network:

UTA2 data network configurations

Use one of the following cable types to cable the UTA2 data ports to your host network.

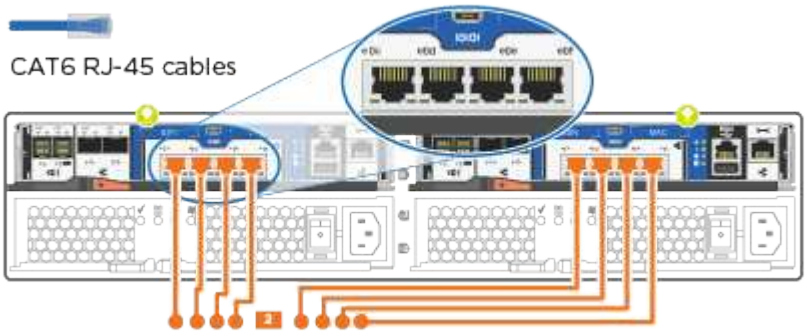
- For an FC host, use 0c and 0d **or** 0e and 0f.
- For an 10GbE system, use e0c and e0d **or** e0e and e0f.



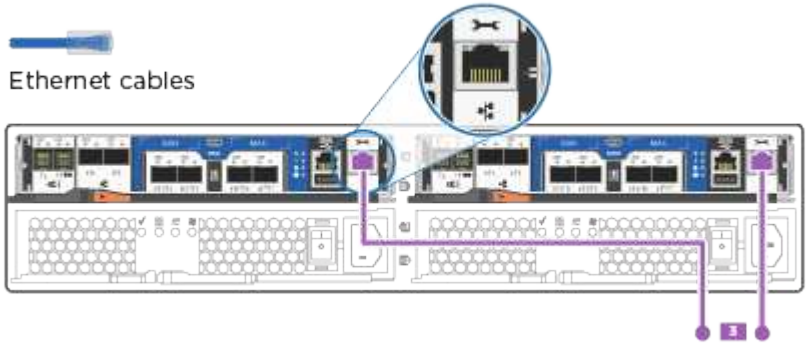
You can connect one port pair as CNA and one port pair as FC, or you can connect both port pairs as CNA or both port pairs as FC.

Ethernet network configurations

Use the Cat 6 RJ45 cable to cable the e0c through e0f ports to your host network. in the following illustration.



3. Cable the e0M ports to the management network switches with the RJ45 cables.



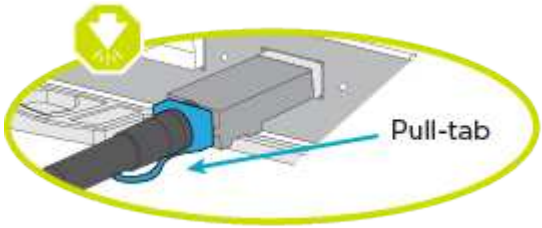
DO NOT plug in the power cords at this point.

Option 2: Switched cluster

Cable your switched cluster.

About this task

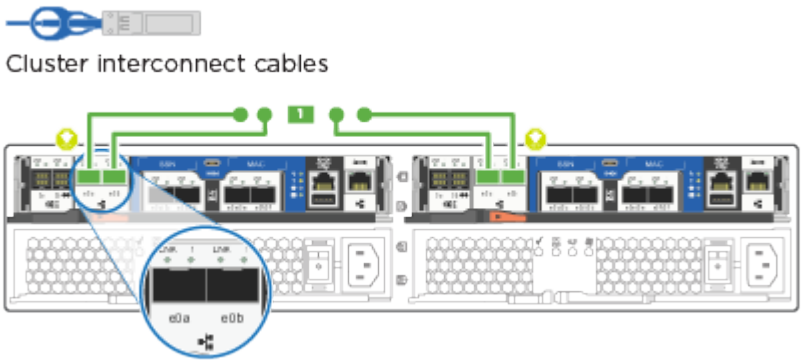
Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.



As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. For each controller module, cable e0a and e0b to the cluster interconnect switches with the cluster interconnect cable.

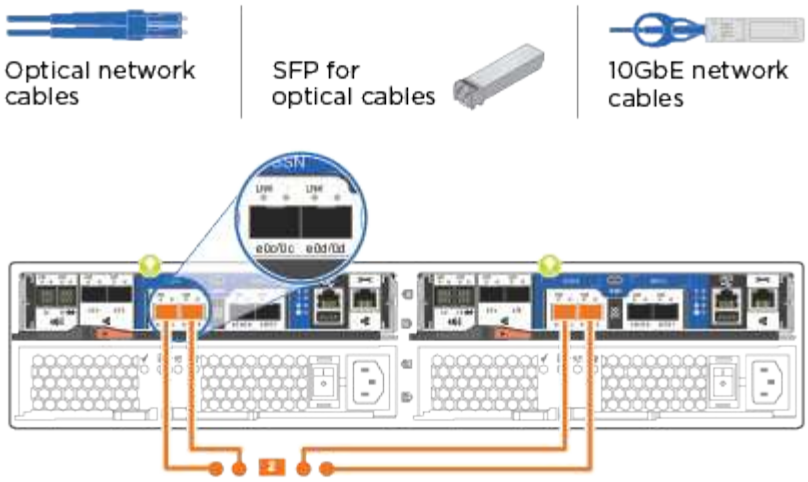


2. You can use either the UTA2 data network ports or the ethernet data network ports to connect the controllers to your host network:

UTA2 data network configurations

Use one of the following cable types to cable the UTA2 data ports to your host network.

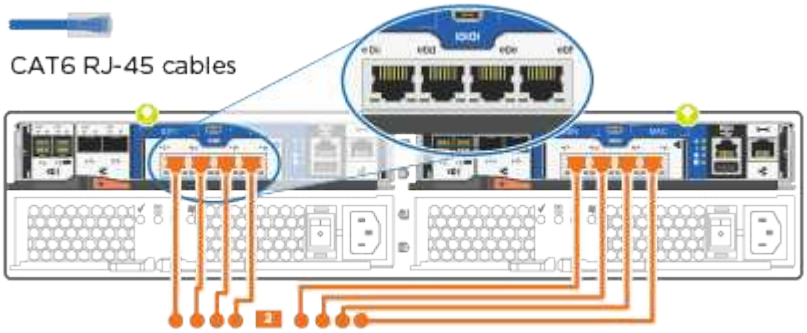
- For an FC host, use 0c and 0d **or** 0e and 0f.
- For an 10GbE system, use e0c and e0d **or** e0e and e0f.



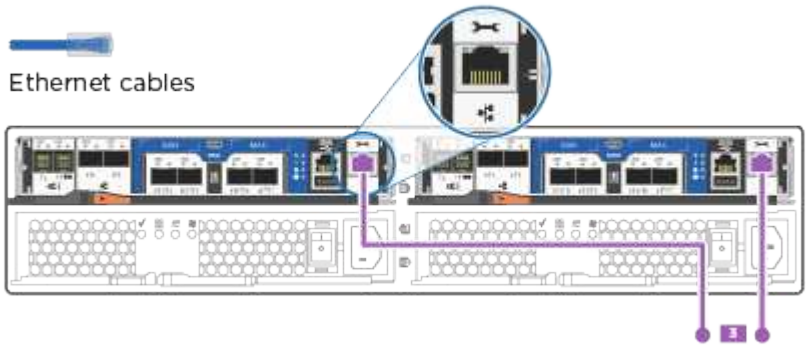
You can connect one port pair as CNA and one port pair as FC, or you can connect both port pairs as CNA or both port pairs as FC.

Ethernet network configurations

Use the Cat 6 RJ45 cable to cable the e0c through e0f ports to your host network.



3. Cable the e0M ports to the management network switches with the RJ45 cables.



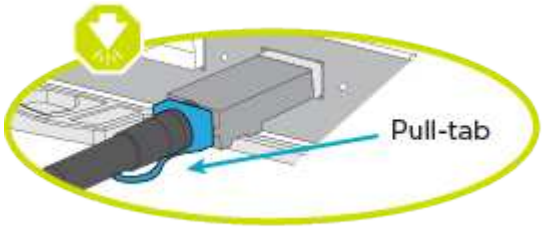
DO NOT plug in the power cords at this point.

Step 4: Cable controllers to drive shelves

Cable the controllers to your shelves using the onboard storage ports. NetApp recommends MP-HA cabling for systems with external storage.

About this task

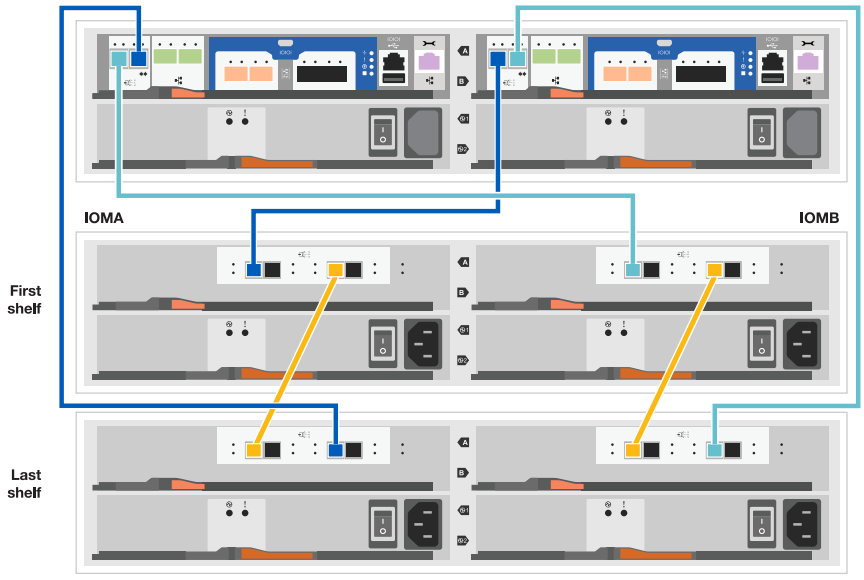
- If you have a SAS tape drive, you can use single-path cabling. If you have no external shelves, MP-HA cabling to internal drives is optional (not shown) if the SAS cables are ordered with the system.
- You must cable the shelf-to-shelf connections, and then cable both controllers to the drive shelves.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.



Steps

1. Cable the HA pair with external drive shelves.

   The following example shows cabling for DS224C drive shelves. The cabling is similar with other supported drive shelves.



2. Cable the shelf-to-shelf ports.
   - Port 3 on IOM A to port 1 on the IOM A on the shelf directly below.
   - Port 3 on IOM B to port 1 on the IOM B on the shelf directly below.

    mini-SAS HD to mini-SAS HD cables

3. Connect each node to IOM A in the stack.
   - Controller 1 port 0b to IOM A port 3 on last drive shelf in the stack.

- Controller 2 port 0a to IOM A port 1 on the first drive shelf in the stack.

 mini-SAS HD to mini-SAS HD cables

4. Connect each node to IOM B in the stack
   - Controller 1 port 0a to IOM B port 1 on first drive shelf in the stack.
   - Controller 2 port 0b to IOM B port 3 on the last drive shelf in the stack.

 mini-SAS HD to mini-SAS HD cables

For additional cabling information, see Install and cable shelves for a new system installation - shelves with IOM12/IOM12B modules.

Step 5: Complete system setup

You can complete the system setup and configuration using cluster discovery with only a connection to the switch and laptop, or by connecting directly to a controller in the system and then connecting to the management switch.

Option 1: If network discovery is enabled

If you have network discovery enabled on your laptop, you can complete system setup and configuration using automatic cluster discovery.

Steps

1. Use the following animation to set one or more drive shelf IDs

   Animation - Set drive shelf IDs

2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Turn on the power switches to both nodes.

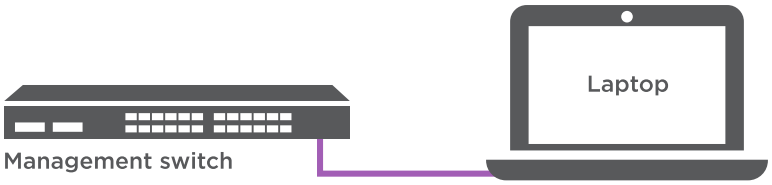


Initial booting may take up to eight minutes.

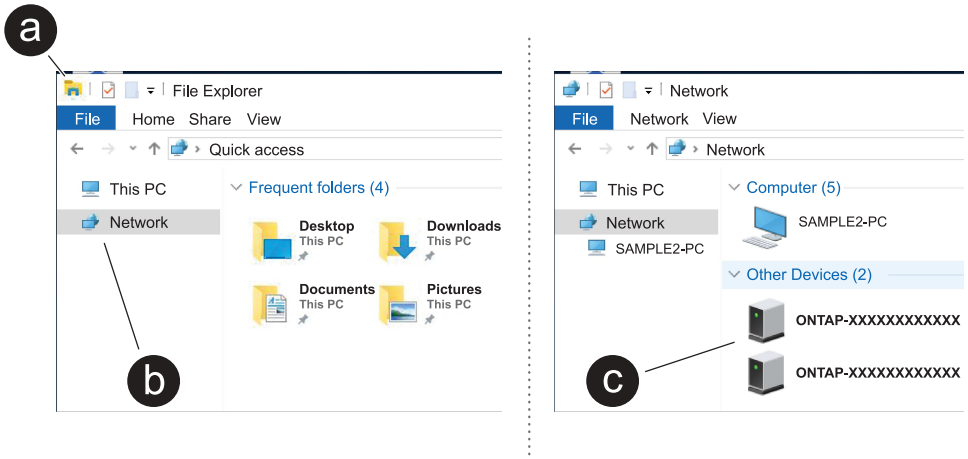
4. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

5. Connect your laptop to the Management switch.



6. Select an ONTAP icon listed to discover:



a. Open File Explorer.
b. Click **Network** in the left pane and right-click and select **refresh**.
c. Double-click either ONTAP icon and accept any certificates displayed on your screen.

XXXXX is the system serial number for the target node.

System Manager opens.

7. Configure the system using the data you collected in the ONTAP Configuration Guide.
8. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create and account.
   b. Register your system.
   c. Download Active IQ Config Advisor.
9. Verify the health of your system by running Config Advisor.
10. After you have completed the initial configuration, go to the ONTAP documentation site for information about configuring additional features in ONTAP.

Option 2: If network discovery is not enabled

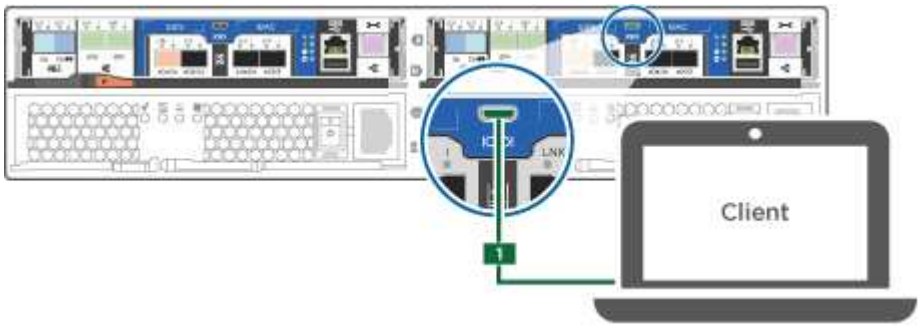
If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

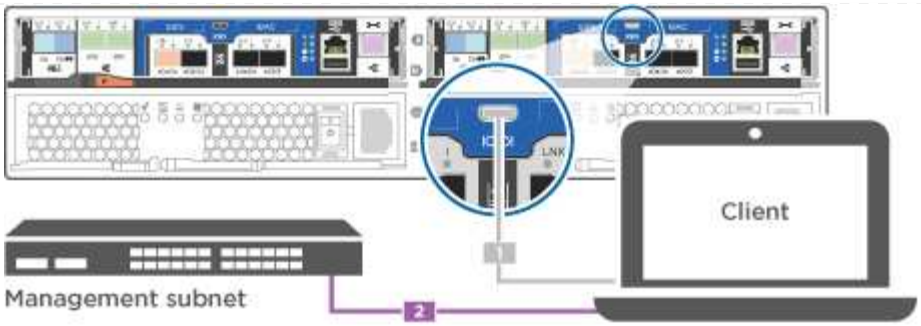
1. Cable and configure your laptop or console.
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for instructions on how to configure the console port.

   b. Connect the console cable to the laptop or console, and connect the console port on the controller using the console cable that came with your system.



   c. Connect the laptop or console to the switch on the management subnet.



d. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.

2. Use the following animation to set one or more drive shelf IDs:

   Animation - Set drive shelf IDs

3. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
4. Turn on the power switches to both nodes.



Initial booting may take up to eight minutes.

5. Assign an initial node management IP address to one of the nodes.

If the management network has DHCP…	Then…
Configured	Record the IP address assigned to the new controllers.
Not configured	a. Open a console session using PuTTY, a terminal server, or the equivalent for your environment. Check your laptop or console's online help if you do not know how to configure PuTTY. b. Enter the management IP address when prompted by the script.

6. Using System Manager on your laptop or console, configure your cluster.
   a. Point your browser to the node management IP address.

The format for the address is https://x.x.x.x.

b. Configure the system using the data you collected in the ONTAP Configuration Guide.

7. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create and account.
   b. Register your system.
   c. Download Active IQ Config Advisor.
8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to the ONTAP documentation site for information about configuring additional features in ONTAP.

Maintain

Maintain AFF A150 hardware

For the AFF A150 storage system, you can perform maintenance procedures on the following components.

Boot media

The boot media stores a primary and secondary set of boot image files that the system uses when it boots.

Chassis

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

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

NVEM Battery

A battery is included with a controller and preserves cached data if the AC power fails.

Power supply

A power supply provides a redundant power source in a controller shelf.

Real time clock battery

A real time clock battery preserves system date and time information if the power is off.

Boot media

Overview of boot media replacement - AFF A150

The boot media stores a primary and secondary set of system (boot image) files that the system uses when it boots. Depending on your network configuration, you can perform either a nondisruptive or disruptive replacement.

You must have a USB flash drive, formatted to FAT32, with the appropriate amount of storage to hold the `image_xxx.tgz` file.

You also must copy the `image_xxx.tgz` file to the USB flash drive for later use in this procedure.

- The nondisruptive and disruptive methods for replacing a boot media both require you to restore the `var` file system:
  - For nondisruptive replacement, the HA pair must be connected to a network to restore the `var` file system.
  - For disruptive replacement, you do not need a network connection to restore the `var` file system, but the process requires two reboots.
- You must replace the failed component with a replacement FRU component you received from your provider.
- It is important that you apply the commands in these steps on the correct node:
  - The *impaired* node is the node on which you are performing maintenance.
  - The *healthy node* is the HA partner of the impaired node.

Check encryption key support and status - AFF A150

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

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

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

Steps

1. Check if your ONTAP version supports encryption:

   ```
   version -v
   ```

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

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

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

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

Step 2: Verify key manager status and back up configuration

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

Steps

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

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

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

   If no key manager is configured:

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

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

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

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

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

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

External Key Manager (EKM)

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

If all keys show `true` in the Restored column:

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

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

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

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

   If the command fails, contact NetApp Support.

b. Verify that all authentication keys are restored:

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

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

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

Onboard Key Manager (OKM)

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

If all keys show `true` in the Restored column:

a. Back up the OKM information:

   i. Switch to advanced privilege mode:

      ```
      set -priv advanced
      ```

      Enter `y` when prompted to continue.

   ii. Display the key management backup information:

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

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

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

   iv. Return to admin mode:

      ```
      set -priv admin
      ```

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

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

a. Synchronize the onboard key manager:

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

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

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

b. Verify all authentication keys are restored:

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

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

c. Back up the OKM information:

i. Switch to advanced privilege mode:

```
set -priv advanced
```

Enter `y` when prompted to continue.

ii. Display the key management backup information:

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

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

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

iv. Return to admin mode:

```
set -priv admin
```

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

Shut down the impaired controller - AFF A150

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

Complete the shutdown of the impaired controller by halting it to the LOADER prompt and capturing the boot environment variables. The LOADER prompt is the safe halted state from which you can replace the boot media.

Steps

a. Take the impaired controller to the LOADER prompt:

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

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

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

Option 2: Controller is in a MetroCluster

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

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

Take over and halt the impaired controller so that the healthy controller continues to serve data from the impaired controller's storage. To do this, you suppress automatic case creation in AutoSupport, disable automatic giveback, and bring the impaired controller to the LOADER prompt. The LOADER prompt is the safe halted state from which you can replace the FRU.

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

Steps

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

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

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

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

Replace the boot media - AFF A150

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

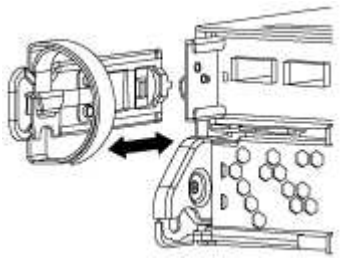
Step 1: Remove the controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

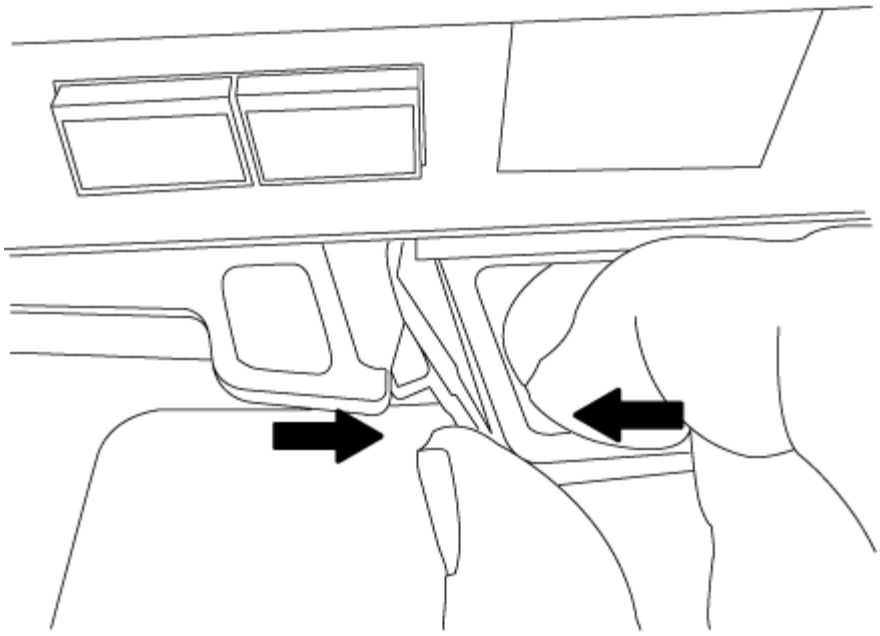
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

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

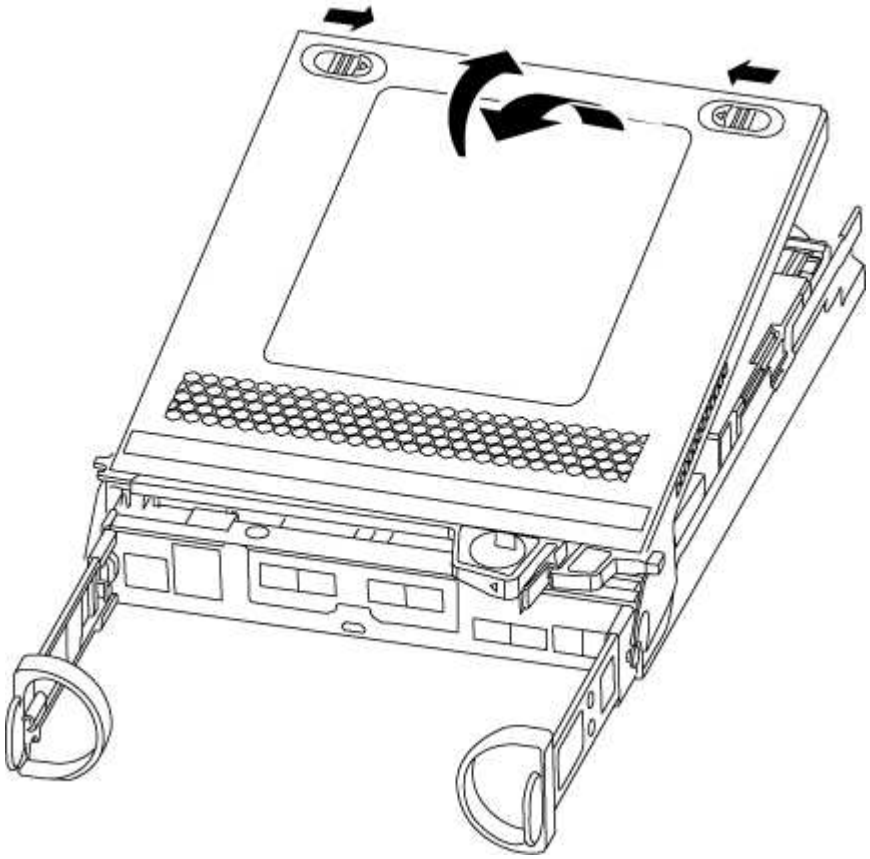
3. Remove and set aside the cable management devices from the left and right sides of the controller module.



4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.



5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

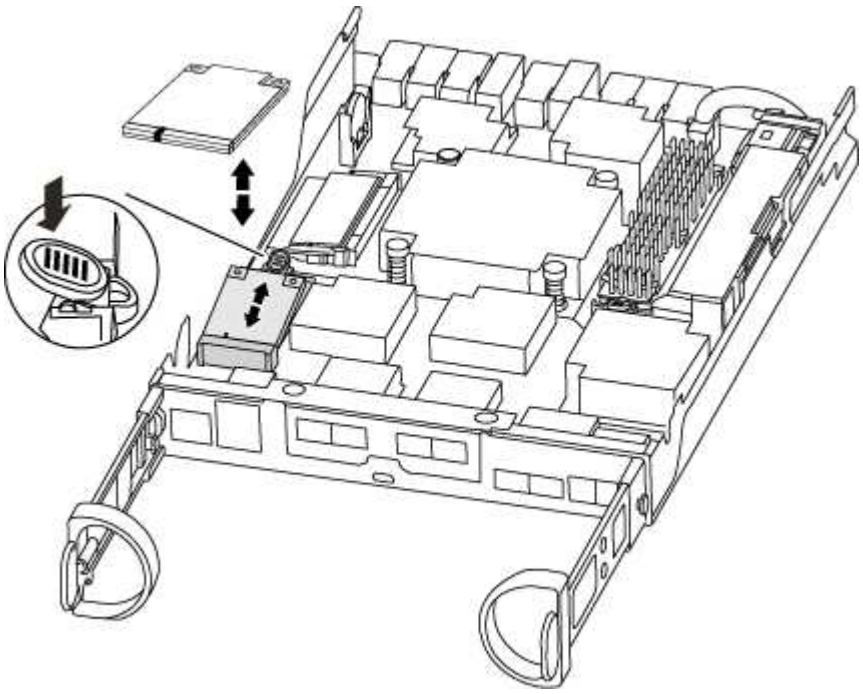


Step 2: Replace the boot media

You must locate the boot media in the controller and follow the directions to replace it.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the boot media using the following illustration or the FRU map on the controller module:



3. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

> Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

4. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
5. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.
6. Push the boot media down to engage the locking button on the boot media housing.
7. Close the controller module cover.

Step 3: Transfer the boot image to the boot media

You can install the system image to the replacement boot media using a USB flash drive with the image installed on it. However, you must restore the var file system during this procedure.

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the var file system.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
2. Reinstall the cable management device and recable the system, as needed.

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

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

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

4. Push the controller module all the way into the system, making sure that the cam handle clears the USB flash drive, firmly push the cam handle to finish seating the controller module, push the cam handle to the closed position, and then tighten the thumbscrew.

   The controller begins to boot as soon as it is completely installed into the chassis.

5. Interrupt the boot process to stop at the LOADER prompt by pressing Ctrl-C when you see Starting AUTOBOOT press Ctrl-C to abort….

   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.

6. For systems with one controller in the chassis, reconnect the power and turn on the power supplies.

   The system begins to boot and stops at the LOADER prompt.

7. Set your network connection type at the LOADER prompt:

   - If you are configuring DHCP: `ifconfig e0a -auto`

     The target port you configure is the target port you use to communicate with the impaired controller from the healthy controller during var file system restore with a network connection. You can also use the e0M port in this command.

   - If you are configuring manual connections: `ifconfig e0a -addr=filer_addr -mask=netmask -gw=gateway-dns=dns_addr-domain=dns_domain`
     - filer_addr is the IP address of the storage system.
     - netmask is the network mask of the management network that is connected to the HA partner.
     - gateway is the gateway for the network.
     - dns_addr is the IP address of a name server on your network.
     - dns_domain is the Domain Name System (DNS) domain name.

       If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL. You need only the server's host name.

Other parameters might be necessary for your interface. You can enter `help ifconfig` at the firmware prompt for details.

Boot the recovery image - AFF A150

You must boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.

Steps

1. From the LOADER prompt, boot the recovery image from the USB flash drive: `boot_recovery`

   The image is downloaded from the USB flash drive.

2. When prompted, either enter the name of the image or accept the default image displayed inside the brackets on your screen.
3. Restore the var file system:

If your system has…	Then…
A network connection	a. Press `y` when prompted to restore the backup configuration. b. Set the healthy controller to advanced privilege level: `set -privilege advanced` c. Run the restore backup command: `system node restore-backup -node local -target-address` *`impaired_node_IP_address`* d. Return the controller to admin level: `set -privilege admin` e. Press `y` when prompted to use the restored configuration. f. Press `y` when prompted to reboot the controller.
No network connection	a. Press `n` when prompted to restore the backup configuration. b. Reboot the system when prompted by the system. c. Select the **Update flash from backup config** (sync flash) option from the displayed menu. If you are prompted to continue with the update, press `y`.

4. Ensure that the environmental variables are set as expected:
   a. Take the controller to the LOADER prompt.
   b. Check the environment variable settings with the `printenv` command.
   c. If an environment variable is not set as expected, modify it with the `setenv` *`environment-variable-name changed-value`* command.
   d. Save your changes using the `savenv` command.
5. The next depends on your system configuration:

- If your system has onboard keymanager, NSE or NVE configured, go to Restore OKM, NSE, and NVE as needed
- If your system does not have onboard keymanager, NSE or NVE configured, complete the steps in this section.

6. From the LOADER prompt, enter the `boot_ontap` command.

If you see…	Then…
The login prompt	Go to the next Step.
Waiting for giveback…	a. Log into the partner controller. b. Confirm the target controller is ready for giveback with the `storage failover show` command.

7. Connect the console cable to the partner controller.
8. Give back the controller using the `storage failover giveback -fromnode local` command.
9. At the cluster prompt, check the logical interfaces with the `net int -is-home false` command.

   If any interfaces are listed as "false", revert those interfaces back to their home port using the `net int revert` command.

10. Move the console cable to the repaired controller and run the `version -v` command to check the ONTAP versions.
11. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.

Restore encryption - AFF A150

Restore encryption on the replacement boot media.

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** 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

ONTAP version	Select this option

ONTAP 9.7 and earlier	Select the hidden option `recover_onboard_keymanager`
	
	Show example boot menu
	
	`Please choose one of the following:`
	
	`(1) Normal Boot.`
	`(2) Boot without /etc/rc.`
	`(3) Change password.`
	`(4) Clean configuration and initialize all disks.`
	`(5) Maintenance mode boot.`
	`(6) Update flash from backup config.`
	`(7) Install new software first.`
	`(8) Reboot node.`
	`(9) Configure Advanced Drive Partitioning.`
	`Selection (1-19)?`
	`recover_onboard_keymanager`

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

   Show example prompt

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

4. Enter the cluster-wide passphrase twice.

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

   Show example prompt

   ```
   Enter the passphrase for onboard key management:

   Enter the passphrase again to confirm:
   ```

5. Enter the backup information:

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

Show example prompt

```
Enter the backup data:

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

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

---------------------------END
BACKUP---------------------------
```

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

   The recovery process completes and displays the following message:

   ```
   Successfully recovered keymanager secrets.
   ```

Show example prompt

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

Successfully recovered keymanager secrets.

*******************************************************************
********************
* Select option "(1) Normal Boot." to complete recovery
process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
********************
```

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

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

Show example prompt

```
*******************************************************************
******************
* Select option "(1) Normal Boot." to complete the recovery
process.
*
*******************************************************************
******************


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

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

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

   On the partner controller:

8. Giveback the impaired controller:

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

   On the impaired controller:

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

   `security key-manager onboard sync`

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

Show example prompt

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

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

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

11. Verify that all keys are synced:

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

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

    On the partner controller:

12. Giveback the impaired controller:

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

13. Restore automatic giveback if you disabled it:

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

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

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

External Key Manager (EKM)

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

Before you begin

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

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

Steps

On the impaired controller:

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

   Show example boot menu

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

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

   Show example prompt

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

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

Show example

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

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

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

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

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

The recovery process completes and displays the following message:

```
Successfully recovered keymanager secrets.
```

Show example

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

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

Show example prompt

```
*****************************************************************
***********
* Select option "(1) Normal Boot." to complete the recovery
process.
*
*****************************************************************
***********

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

6. Restore automatic giveback if you disabled it:

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

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

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

Return the failed part to NetApp - AFF A150

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

Chassis

Overview of chassis replacement - AFF A150

To replace the chassis, you must move the power supplies, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

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

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is written with the assumption that you are moving all drives and controller module or modules to the new chassis, and that the chassis is a new component from NetApp.
- This procedure is disruptive. For a two-controller cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - AFF A150

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Most configurations

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

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessibility for each controller.
- Make sure you have the necessary tools and equipment for the replacement.

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 in to the SP/BMC of each controller using the IP address identified in the previous step:
   - If you log in from the BMC over SSH, connect using the SP/BMC address (for example, `ssh admin@<SP/BMC_address>`), then enter the command `system console` and authenticate.
   - If you are using a local console or laptop connected directly to the controller, log in using the same cluster administrator credentials.

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

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

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

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

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

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

Option 2: Controller is in a MetroCluster configuration

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

Take over and halt the impaired controller so that the healthy controller continues to serve data from the impaired controller's storage. To do this, you suppress automatic case creation in AutoSupport, disable automatic giveback, and bring the impaired controller to the LOADER prompt. The LOADER prompt is the safe halted state from which you can replace the FRU.

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

Steps

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

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

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

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

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` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Move and replace hardware - AFF A150

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

Step 1: Move a power supply

Moving out a power supply when replacing a chassis involves turning off, disconnecting, and removing the power supply from the old chassis and installing and connecting it on the replacement chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
3. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.
4. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

5. Repeat the preceding steps for any remaining power supplies.
6. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

7. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
8. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.

   Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

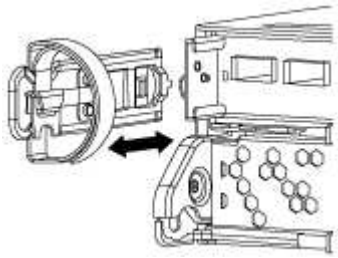
Step 2: Remove the controller module

Remove the controller module or modules from the old chassis.

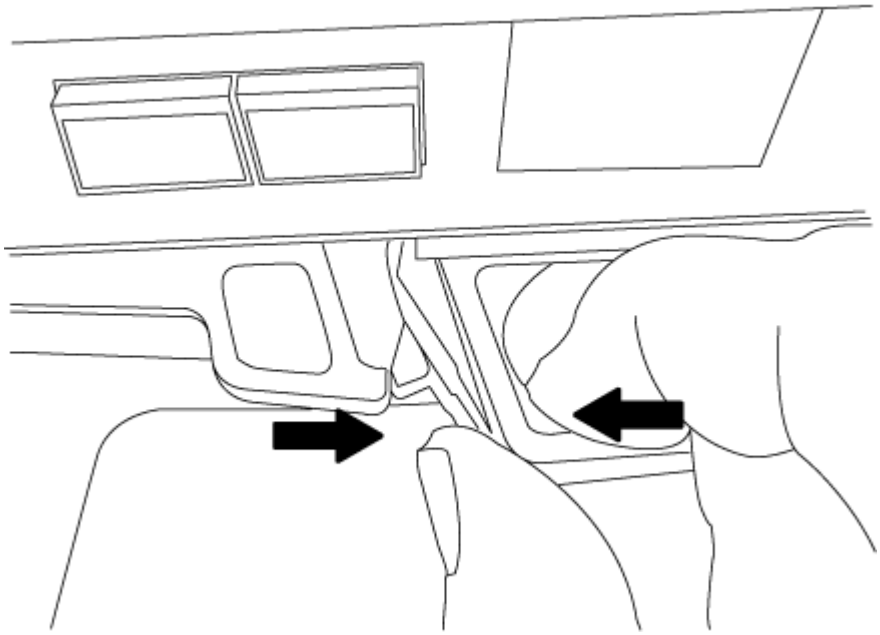
1. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

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

2. Remove and set aside the cable management devices from the left and right sides of the controller module.



3. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.



4. Set the controller module aside in a safe place, and repeat these steps if you have another controller module in the chassis.

Step 3: Move drives to the new chassis

You need to move the drives from each bay opening in the old chassis to the same bay opening in the new chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:

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

b. Pull the cam handle to its fully open position to unseat the drive from the midplane, and then gently slide the drive out of the chassis.

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

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

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

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

   The cam handle engages and begins to rotate upward.

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

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

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

Step 4: Replace a chassis from within the equipment rack or system cabinet

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

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

Step 5: Install the controller

After you install the controller module and any other components into the new chassis, boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

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

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

2. Recable the console to the controller module, and then reconnect the management port.
3. Repeat the preceding steps if there is a second controller to install in the new chassis.
4. Complete the installation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Repeat the preceding steps for the second controller module in the new chassis.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reinstall the blanking panel and then go to the next step.

5. Connect the power supplies to different power sources, and then turn them on.
6. Boot each controller to Maintenance mode:
   a. As each controller starts the booting, press `Ctrl-C` to interrupt the boot process when you see the message `Press Ctrl-C for Boot Menu`.

      If you miss the prompt and the controller modules boot to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then repeat this step.

   b. From the boot menu, select the option for Maintenance mode.

Restore and verify the configuration - AFF A150

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

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

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

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

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis *HA-state*`

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

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

   b. Confirm that the setting has changed: `ha-config show`
3. If you have not already done so, recable the rest of your system.
4. Reboot the system.

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

This task only applies to two-node MetroCluster configurations.

Steps

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

```
cluster_B::>  metrocluster node show

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

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

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

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

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

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

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

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: 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

Overview of controller module replacement - AFF A150

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

- All drive shelves must be working properly.
- If your system is in an HA pair, the healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

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

- This procedure includes steps for automatically or manually reassigning drives to the *replacement* controller, depending on your system’s configuration.

  You should perform the drive reassignment as directed in the procedure.

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

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

Shut down the impaired controller - AFF A150

Take over and halt the impaired controller so that the healthy controller continues to serve data from the impaired controller’s storage. To do this, you suppress automatic case creation in AutoSupport, disable automatic giveback, and bring the impaired controller to

the LOADER prompt. The LOADER prompt is the safe halted state from which you can replace the FRU.

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.

Replace the controller module hardware - AFF A150

To replace the controller module hardware, you must remove the impaired controller,

move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

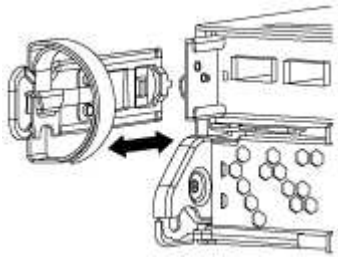
Step 1: Remove controller module

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

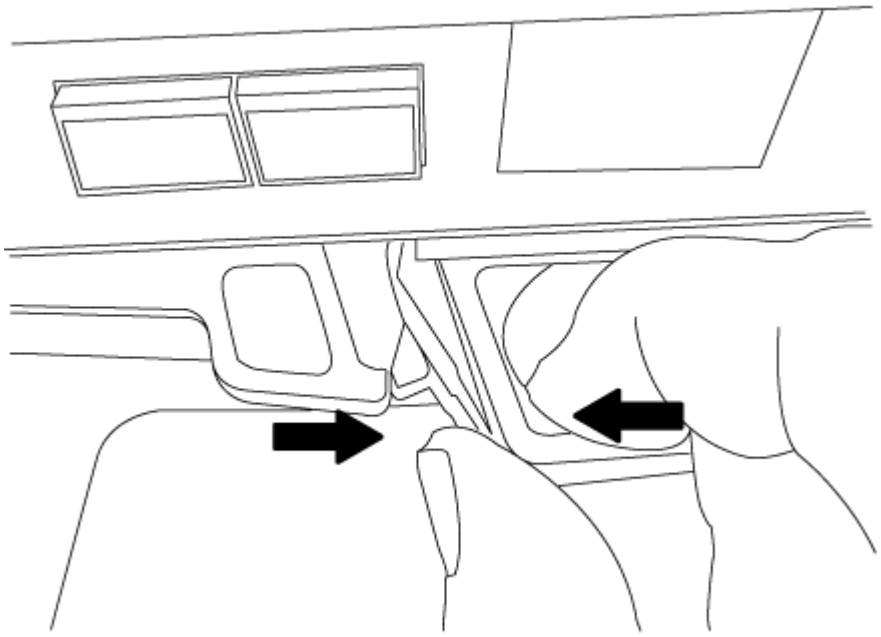
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

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

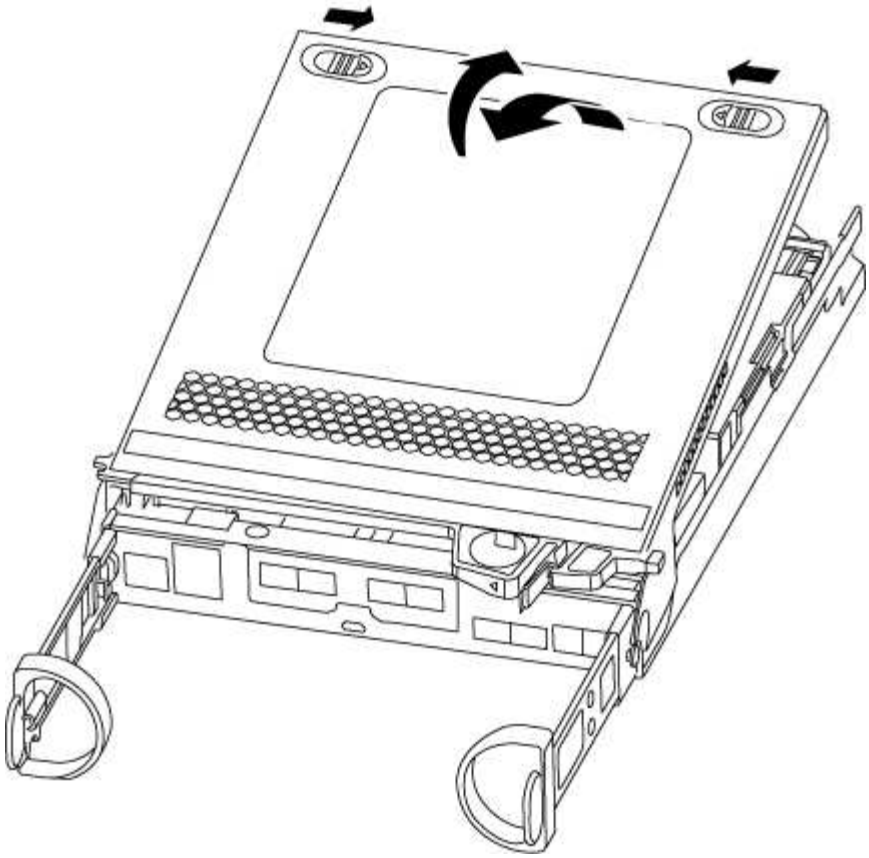
3. Remove and set aside the cable management devices from the left and right sides of the controller module.



4. If you left the SFP modules in the system after removing the cables, move them to the new controller module.
5. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.



6. Turn the controller module over and place it on a flat, stable surface.
7. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.



Step 2: Move the NVMEM battery

To move the NVMEM battery from the old controller module to the new controller module, you must perform a specific sequence of steps.

1. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

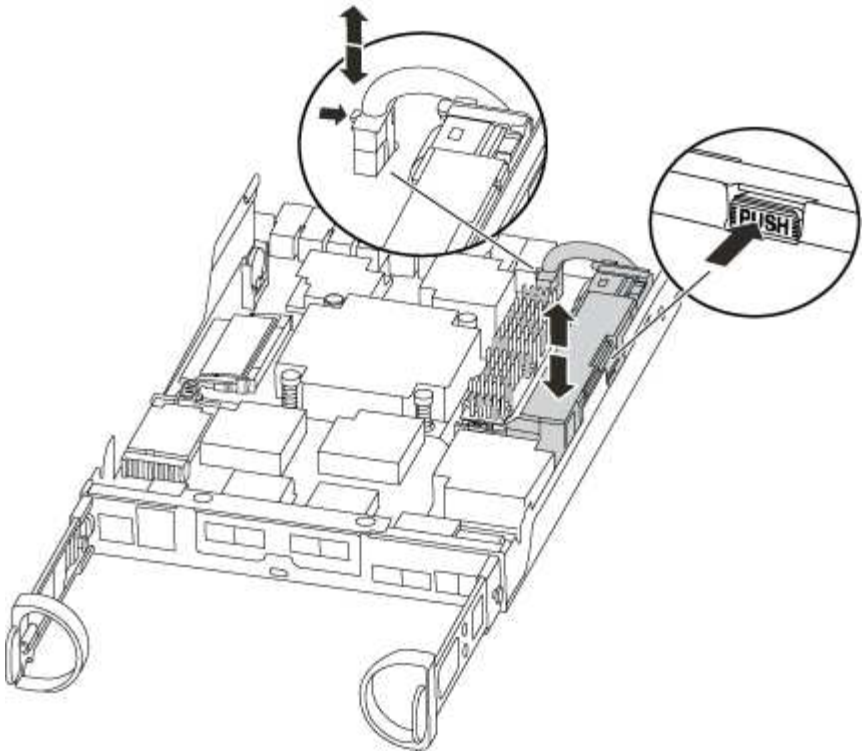


The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

2. Locate the NVMEM battery in the controller module.

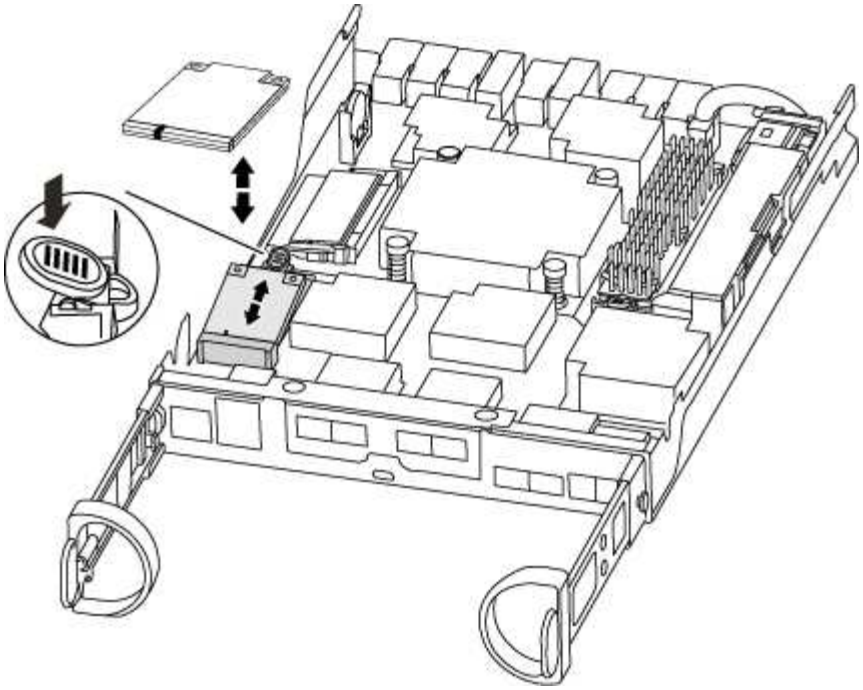


3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. Move the battery to the replacement controller module.
6. Loop the battery cable around the cable channel on the side of the battery holder.
7. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
8. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.

Step 3: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller module and insert it in the new controller module.

1. Locate the boot media using the following illustration or the FRU map on the controller module:



2. Press the blue button on the boot media housing to release the boot media from its housing, and then gently pull it straight out of the boot media socket.

   Do not twist or pull the boot media straight up, because this could damage the socket or the boot media.

3. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

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

5. Push the boot media down to engage the locking button on the boot media housing.

Step 4: Move the DIMMs

To move the DIMMs, you must follow the directions to locate and move them from the old controller module into the replacement controller module.

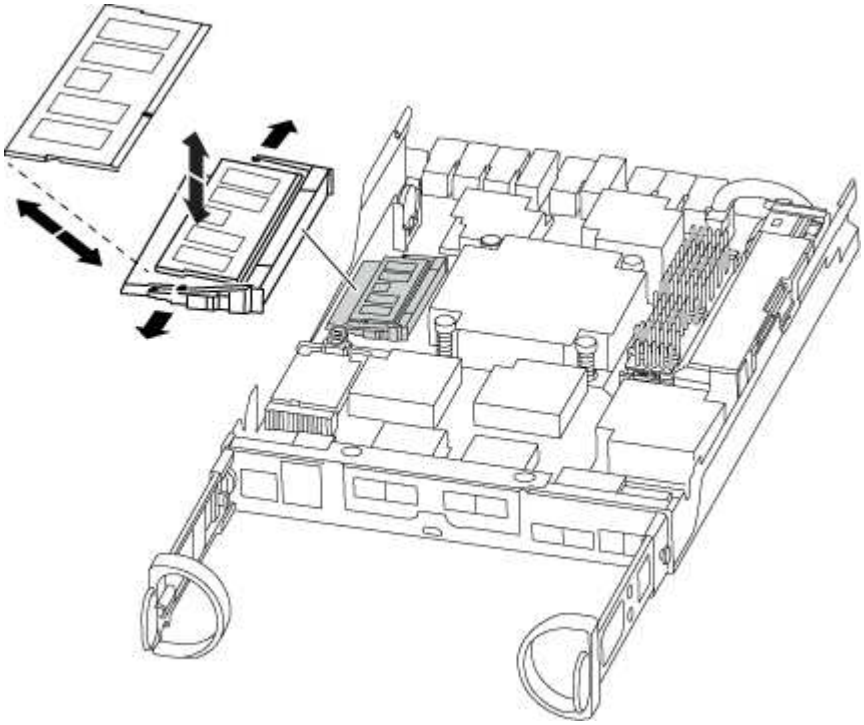
You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

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

   The number and placement of system DIMMs depends on the model of your system.

The following illustration shows the location of system DIMMs:



4. Repeat these steps to remove additional DIMMs as needed.
5. Verify that the NVMEM battery is not plugged into the new controller module.
6. Locate the slot where you are installing the DIMM.
7. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

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

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

8. Repeat these steps for the remaining DIMMs.
9. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

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

Step 5: Move a caching module, if present

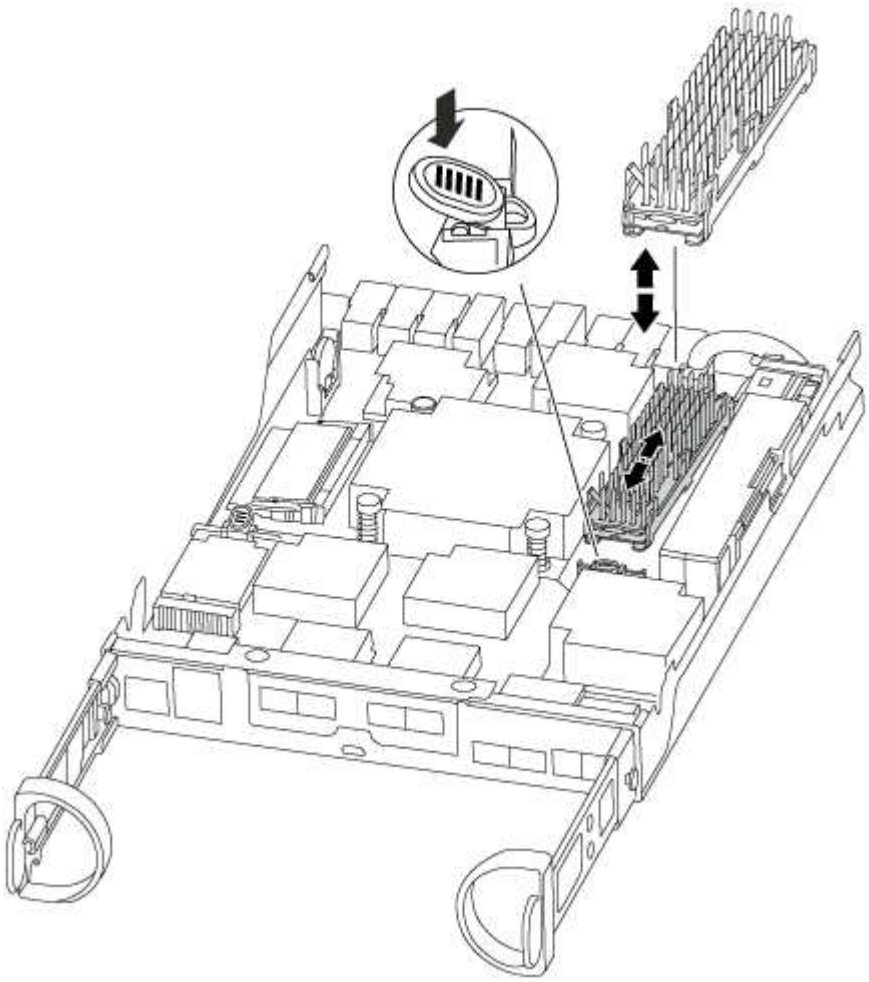
If your AFF A220 or FAS2700 system has a caching module, you need to move the caching module from the old controller module to the replacement controller module. The caching module is referred to as the “M.2 PCIe card” on the controller module label.

You must have the new controller module ready so that you can move the caching module directly from the old controller module to the corresponding slot in the new one. All other components in the storage system must be functioning properly; if not, you must contact technical support.

1. Locate the caching module at the rear of the controller module and remove it.

a. Press the release tab.

b. Remove the heatsink.



2. Gently pull the caching module straight out of the housing.
3. Move the caching module to the new controller module, and then align the edges of the caching module with the socket housing and gently push it into the socket.
4. Verify that the caching module is seated squarely and completely in the socket.

   If necessary, remove the caching module and reseat it into the socket.

5. Reseat and push the heatsink down to engage the locking button on the caching module housing.
6. Close the controller module cover, as needed.

Step 6: Install the controller

After you install the components from the old controller module into the new controller module, you must install the new controller module into the system chassis and boot the operating system.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

The system might update system firmware when it boots. Do not abort this process. The procedure requires you to interrupt the boot process, which you can typically do at any time after prompted to do so. However, if the system updates the system firmware when it boots, you must wait until after the update is complete before interrupting the boot process.

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

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

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

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

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Interrupt the boot process **only** after determining the correct timing: You must look for an Automatic firmware update console message. If the update message appears, do not press `Ctrl-C` to interrupt the boot process until after you see a message confirming that the update is complete. Only press `Ctrl-C` when you see the message `Press Ctrl-C for Boot Menu`. If the firmware update is aborted, the boot process exits to the LOADER prompt. You must run the update_flash command and then exit LOADER and boot to Maintenance mode by pressing `Ctrl-C` when you see Starting AUTOBOOT press Ctrl-C to abort. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. e. Select the option to boot to Maintenance mode from the displayed menu.

If your system is in…	Then perform these steps…
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process. e. Interrupt the boot process **only** after determining the correct timing: You must look for an Automatic firmware update console message. If the update message appears, do not press `Ctrl-C` to interrupt the boot process until after you see a message confirming that the update is complete. Only press `Ctrl-C` after you see the `Press Ctrl-C for Boot Menu` message. If the firmware update is aborted, the boot process exits to the LOADER prompt. You must run the update_flash command and then exit LOADER and boot to Maintenance mode by pressing `Ctrl-C` when you see Starting AUTOBOOT press Ctrl-C to abort. If you miss the prompt and the controller module boots to ONTAP, enter `halt`, and then at the LOADER prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then boot to Maintenance mode. f. From the boot menu, select the option for Maintenance mode.

Important: During the boot process, you might see the following prompts:

- A prompt warning of a system ID mismatch and asking to override the system ID.
- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down. You can safely respond `y` to these prompts.

Restore and verify the system configuration - AFF A150

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

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

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

About this task

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

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

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

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

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

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

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

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

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:
   a. Set the HA state for the chassis: `ha-config modify chassis` *`HA-state`*

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

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

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

Recable the system and reassign disks - AFF A150

To complete the replacement procedure and restore your system to full operation, you must recable the storage, confirm disk reassignment, restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller. You must complete a series of tasks before restoring your system to full operation.

Step 1: Recable the system

Verify the controller module's storage and network connections by using Active IQ Config Advisor.

Steps

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

Step 2: Reassign disks

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

You must use the correct procedure for your configuration:

Controller redundancy	Then use this procedure…
HA pair	Option 1: Verify the system ID change on an HA system
Stand-alone	Option 2: Manually reassign the system ID on a stand-alone system in ONTAP

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

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

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

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

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

```
node1> `storage failover show`
                                    Takeover
Node                Partner         Possible     State Description
------------      ------------      --------
-------------------------------------
node1               node2           false        System ID changed on
partner (Old:
                                                  151759755, New:
151759706), In takeover
node2               node1           -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

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

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

The *replacement* controller takes back its storage and completes booting.

If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

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

Find the High-Availability Configuration content for your version of ONTAP 9

b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

The output from the `storage failover show` command should not include the System ID changed on partner message.

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

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

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

Option 2: Manually reassign the system ID on a stand-alone system in ONTAP

In a stand-alone system, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems that are in a stand-alone configuration.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by pressing Ctrl-C, and then select the option to boot to Maintenance mode from the displayed menu.
2. You must enter `Y` when prompted to override the system ID due to a system ID mismatch.
3. View the system IDs: `disk show -a`

4. You should make a note of the old system ID, which is displayed as part of the disk owner column.

   The following example shows the old system ID of 118073209:

```
*> disk show -a
Local System ID: 118065481

  DISK      OWNER                    POOL   SERIAL NUMBER  HOME
--------    -------------            -----  -------------  -------------
disk_name    system-1  (118073209)   Pool0  J8XJE9LC       system-1
(118073209)
disk_name    system-1  (118073209)   Pool0  J8Y478RC       system-1
(118073209)
.
.
.
```

5. Reassign disk ownership by using the system ID information obtained from the disk show command: `disk reassign -s old system ID disk reassign -s 118073209`
6. Verify that the disks were assigned correctly: `disk show -a`

   The disks belonging to the replacement node should show the new system ID. The following example now show the disks owned by system-1 the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

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

7. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
8. Boot the node: `boot_ontap`

Complete system restoration - AFF A150

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

Step 1: Install licenses for the replacement controller in ONTAP

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

About this task

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

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

Before you begin

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

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

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

Steps

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

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and register the serial number

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

Steps

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

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

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

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

This task only applies to two-node MetroCluster configurations.

Steps

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

```
cluster_B::>  metrocluster node show

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

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

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

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

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

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

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

6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

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

Replace a DIMM - AFF A150

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

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

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

Step 1: Shut down the impaired controller

Take over and halt the impaired controller so that the healthy controller continues to serve data from the impaired controller's storage. To do this, you suppress automatic case creation in AutoSupport, disable automatic giveback, and bring the impaired controller to the LOADER prompt. The LOADER prompt is the safe halted state from which you can replace the FRU.

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.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

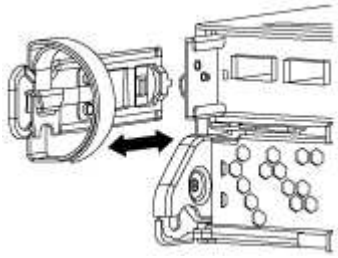
To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

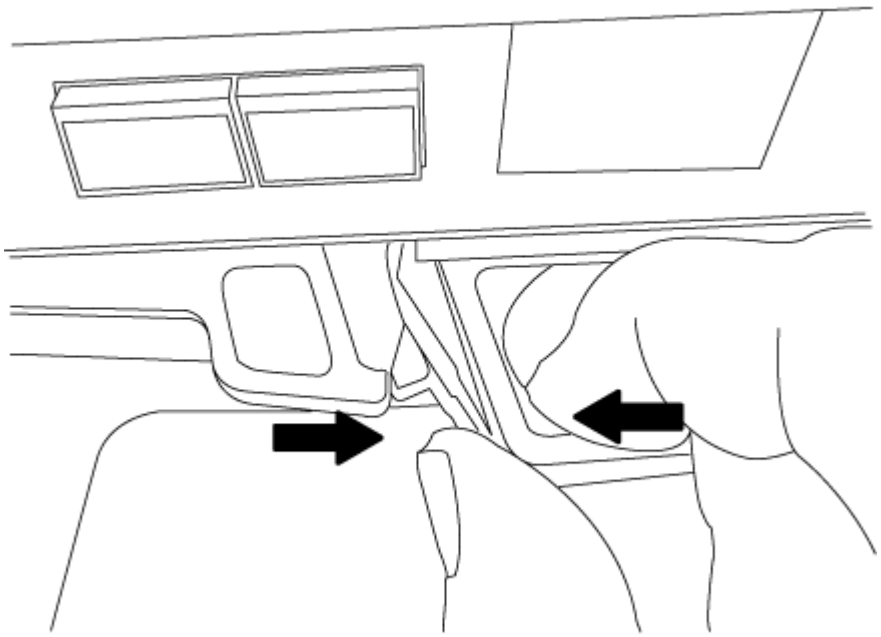
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

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

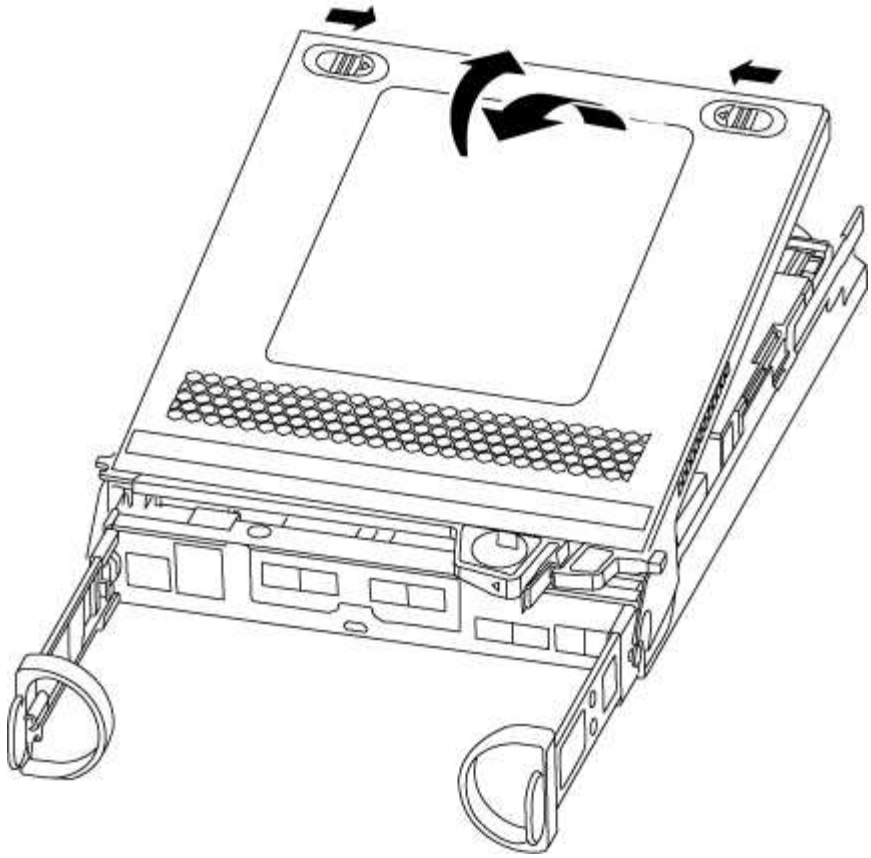
3. Remove and set aside the cable management devices from the left and right sides of the controller module.



4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.



5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.



Step 3: Replace the DIMMs

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

If you are replacing a DIMM, you need to remove it after you have unplugged the NVMEM battery from the controller module.

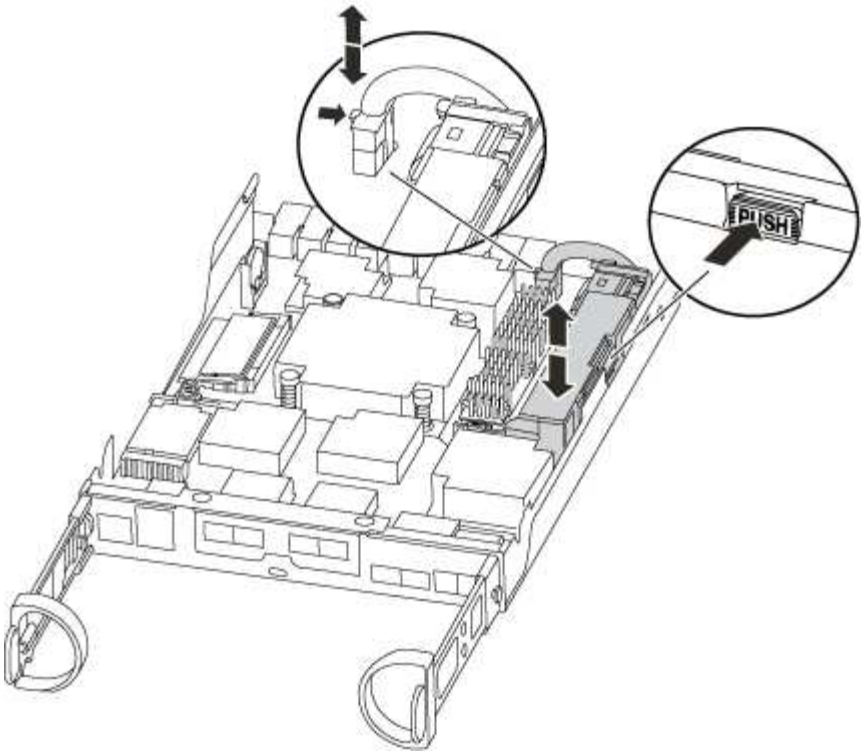
Steps

1. If you are not already grounded, properly ground yourself.
2. Check the NVMEM LED on the back of controller module.

   You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the nonvolatile memory (NVMEM). The LED is located on the back of the controller module. Look for the following icon:

3. If the NVMEM LED is not flashing, there is no content in the NVMEM; you can skip the following steps and proceed to the next task in this procedure.
4. If the NVMEM LED is flashing, there is data in the NVMEM and you must disconnect the battery to clear the memory:
   a. Locate the battery, press the clip on the face of the battery plug to release the lock clip from the plug socket, and then unplug the battery cable from the socket.



b. Confirm that the NVMEM LED is no longer lit.

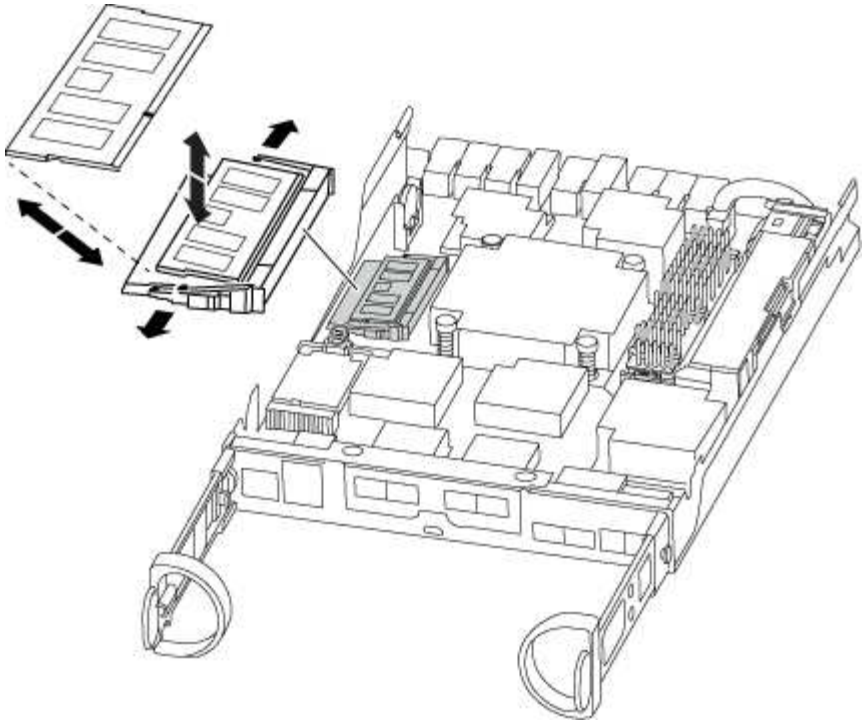
c. Reconnect the battery connector.

5. Return to Step 3: Replace the DIMMs of this procedure to recheck the NVMEM LED.
6. Locate the DIMMs on your controller module.
7. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
8. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

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

   The number and placement of system DIMMs depends on the model of your system.

   The following illustration shows the location of system DIMMs:



9. 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.

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

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

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

11. 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.
12. Locate the NVMEM battery plug socket, and then squeeze the clip on the face of the battery cable plug to insert it into the socket.

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

13. Close the controller module cover.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

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

4. Recable the system, as needed.

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

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, then turn on the power to start the boot process.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

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

```
cluster_B::>  metrocluster node show

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

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

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

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

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

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

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

6. Reestablish any SnapMirror or SnapVault configurations.

Step 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 SSD Drive or HDD Drive - AFF A150

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

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

Before you begin

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

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

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

- Determine whether SED authentication is enabled.

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

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

About this task

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

Option 1: Replace SSD

Steps

1. If you want to manually assign drive ownership for the replacement drive, you need to disable automatic drive assignment, if it is enabled.
   a. Verify whether automatic drive assignment is enabled: `storage disk option show`

      You can enter the command on either controller module.

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

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

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Physically identify the failed drive.

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

   The activity (green) LED on a failed drive can be illuminated (solid), which indicates that the drive has power, but should not be blinking, which indicates I/O activity. A failed drive has no I/O activity.
4. Remove the failed drive:
   a. Press the release button on the drive face to open the cam handle.
   b. Slide the drive out of the shelf using the cam handle and supporting the drive with your other hand.
5. Wait a minimum of 70 seconds before inserting the replacement drive.

   This allows the system to recognize that a drive was removed.
6. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the replacement drive.
   b. Push until the drive stops.
   c. Close the cam handle so that the drive is fully seated into the midplane and the handle clicks into place.

      Be sure to close the cam handle slowly so that it aligns correctly with the face of the drive.
7. Verify that the drive’s activity (green) LED is illuminated.

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

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

      `storage disk show -container-type unassigned`

      You can enter the command on either controller module.

   b. Assign each drive:

      `storage disk assign -disk disk_name -owner node_name`

      You can enter the command on either controller module.

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

   c. Reenable automatic drive assignment if needed:

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

      You must reenable automatic drive assignment on both controller modules.
10. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

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

Option 2: Replace HDD

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

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

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

      You can enter the command on either controller module.

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

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

      You must disable automatic drive assignment on both controller modules.
2. Properly ground yourself.
3. Gently remove the bezel from the front of the platform.
4. Identify the failed disk drive from the system console warning message and the illuminated fault LED on the disk drive

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

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

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

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

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

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

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

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

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

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

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

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

       You can enter the command on either controller module.

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

       You can enter the command on either controller module.

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

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

You must reenable automatic drive assignment on both controller modules.

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

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

Replace the NVMEM battery - AFF A150

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

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

Step 1: Shut down the impaired controller

Take over and halt the impaired controller so that the healthy controller continues to serve data from the impaired controller's storage. To do this, you suppress automatic case creation in AutoSupport, disable automatic giveback, and bring the impaired controller to the LOADER prompt. The LOADER prompt is the safe halted state from which you can replace the FRU.

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.

4. If the system has only one controller module in the chassis, turn off the power supplies, and then unplug the impaired controller's power cords from the power source.

Step 2: Remove controller module

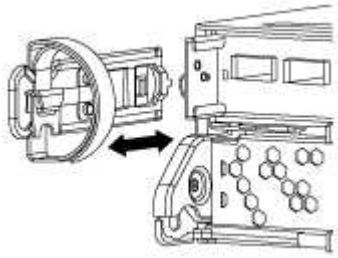
To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

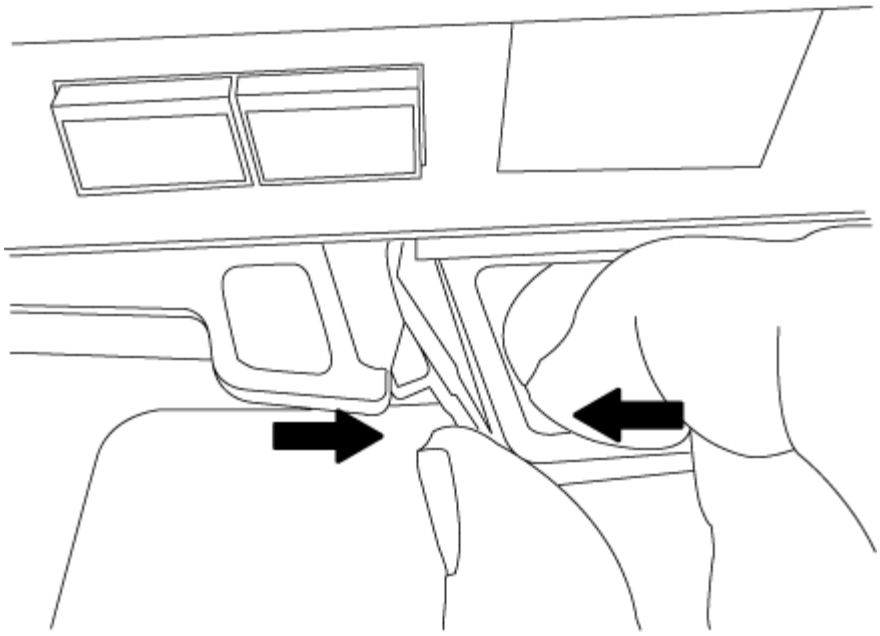
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

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

3. Remove and set aside the cable management devices from the left and right sides of the controller module.



4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.



5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.

Step 3: Replace the NVMEM battery

To replace the NVMEM battery in your system, you must remove the failed NVMEM battery from the system and replace it with a new NVMEM battery.

Steps

1. If you are not already grounded, properly ground yourself.
2. Check the NVMEM LED:
   - If your system is in an HA configuration, go to the next step.
   - If your system is in a stand-alone configuration, cleanly shut down the controller module, and then check the NVRAM LED identified by the NV icon.

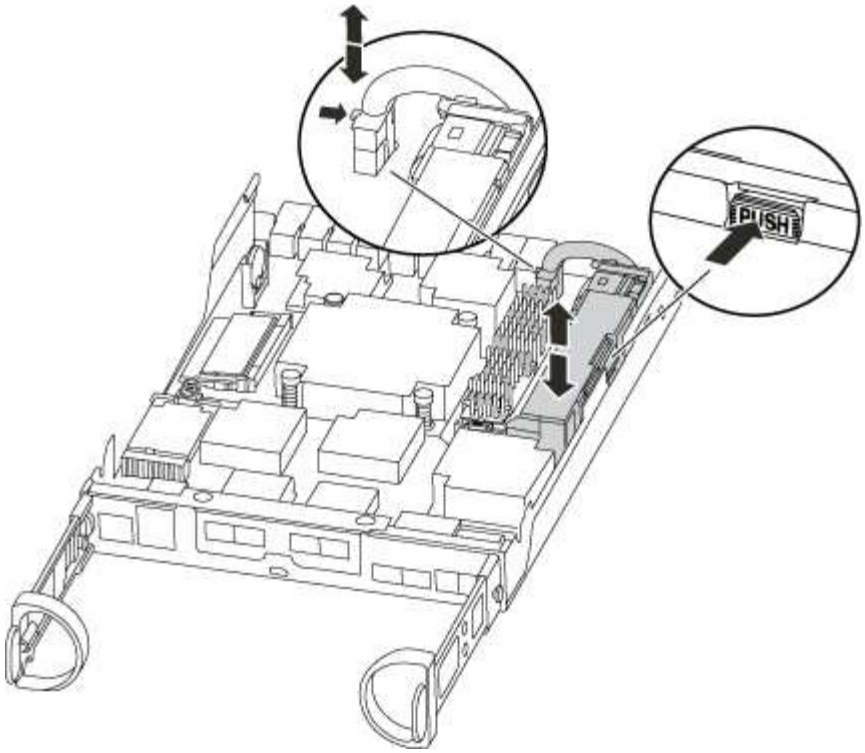


The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - If power is lost without a clean shutdown, the NVMEM LED flashes until the destage is complete, and then the LED turns off.
   - If the LED is on and power is on, unwritten data is stored on NVMEM.

     This typically occurs during an uncontrolled shutdown after ONTAP has successfully booted.

3. Locate the NVMEM battery in the controller module.



4. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
5. Remove the battery from the controller module and set it aside.
6. Remove the replacement battery from its package.
7. Loop the battery cable around the cable channel on the side of the battery holder.
8. Position the battery pack by aligning the battery holder key ribs to the “V” notches on the sheet metal side wall.
9. Slide the battery pack down along the sheet metal side wall until the support tabs on the side wall hook into the slots on the battery pack, and the battery pack latch engages and clicks into the opening on the side wall.
10. Plug the battery plug back into the controller module.

Step 4: Reinstall the controller module

After you replace components in the controller module, reinstall it into the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. If you have not already done so, replace the cover on the controller module.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

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

4. Recable the system, as needed.

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

optic cables.

5. Complete the reinstallation of the controller module:

If your system is in…	Then perform these steps…
An HA pair	The controller module begins to boot as soon as it is fully seated in the chassis. a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. The controller begins to boot as soon as it is seated in the chassis. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap.
A stand-alone configuration	a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors. b. If you have not already done so, reinstall the cable management device. c. Bind the cables to the cable management device with the hook and loop strap. d. Reconnect the power cables to the power supplies and to the power sources, turn on the power to start the boot process.

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

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

```
cluster_B::>  metrocluster node show

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

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

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

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

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

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

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

6. Reestablish any SnapMirror or SnapVault configurations.

Step 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.

Hot-swap a power supply - AFF A150

Swapping out a power supply involves turning off, disconnecting, and removing the old power supply and installing, connecting, and turning on the replacement power supply.

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

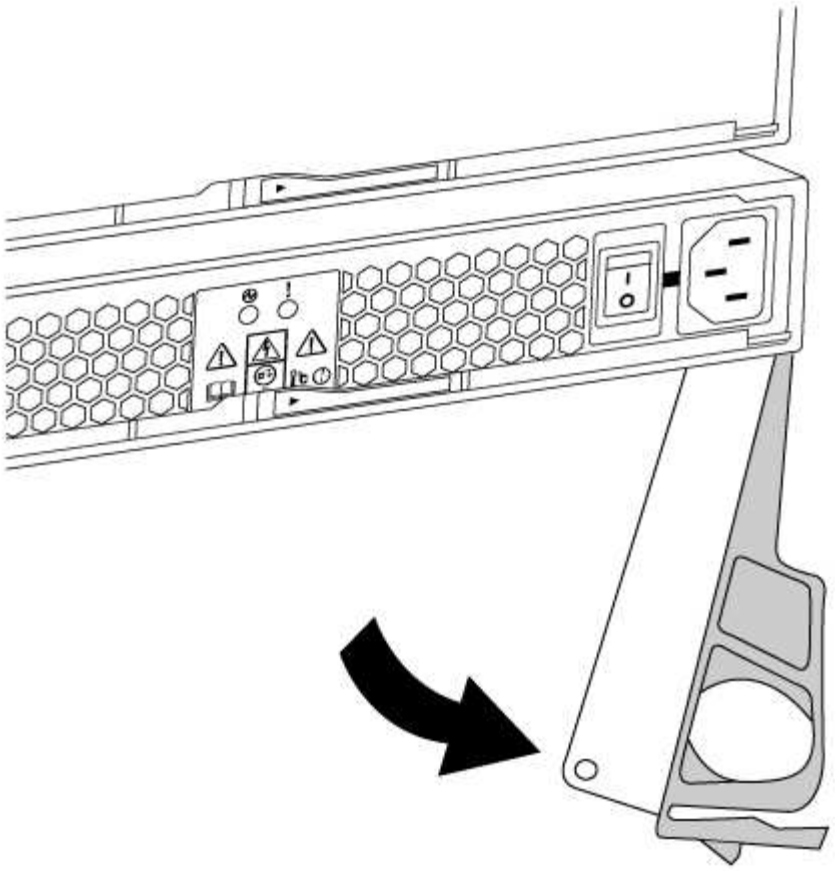
- The power supplies are redundant and hot-swappable. You do not have to shut down the controller to replace a PSU.
- This procedure is written for replacing one power supply at a time.

> Cooling is integrated with the power supply, so you must replace the power supply within two minutes of removal to prevent overheating due to reduced airflow. Because the chassis provides a shared cooling configuration for the two HA nodes, a delay longer than two minutes will shut down all controller modules in the chassis. If both controller modules do shut down, make sure that both power supplies are inserted, turn both off for 30 seconds, and then turn both on.

- Power supplies are auto-ranging.

Steps

1. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
2. If you are not already grounded, properly ground yourself.
3. Turn off the power supply and disconnect the power cables:
   a. Turn off the power switch on the power supply.
   b. Open the power cable retainer, and then unplug the power cable from the power supply.
   c. Unplug the power cable from the power source.
4. Squeeze the latch on the power supply cam handle, and then open the cam handle to fully release the power supply from the mid plane.



5. Use the cam handle to slide the power supply out of the system.

   When removing a power supply, always use two hands to support its weight.

6. Make sure that the on/off switch of the new power supply is in the Off position.
7. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis using the cam handle.

   The power supplies are keyed and can only be installed one way.

   Do not use excessive force when sliding the power supply into the system. You can damage the connector.

8. Close the cam handle so that the latch clicks into the locked position and the power supply is fully seated.
9. Reconnect the power supply cabling:

   a. Reconnect the power cable to the power supply and the power source.
   b. Secure the power cable to the power supply using the power cable retainer.

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

10. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

    The power supply LEDs are lit when the power supply comes online.

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

Replace the real-time clock battery - AFF A150

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

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

You must use an approved RTC battery.

Step 1: Shut down the impaired controller

Take over and halt the impaired controller so that the healthy controller continues to serve data from the impaired controller's storage. To do this, you suppress automatic case creation in AutoSupport, disable automatic giveback, and bring the impaired controller to the LOADER prompt. The LOADER prompt is the safe halted state from which you can replace the FRU.

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

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

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

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

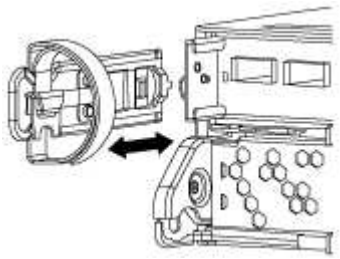
Step 2: Remove controller module

To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

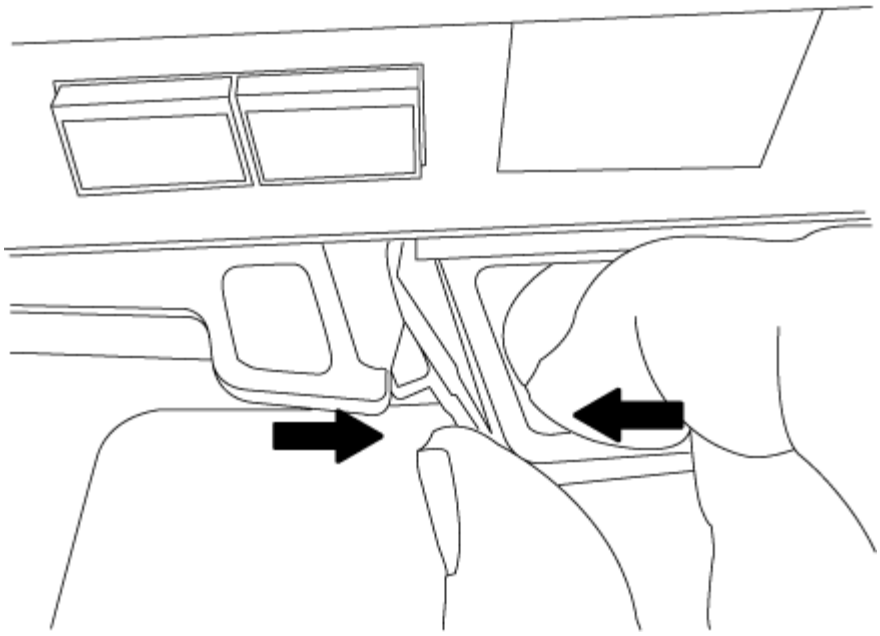
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

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

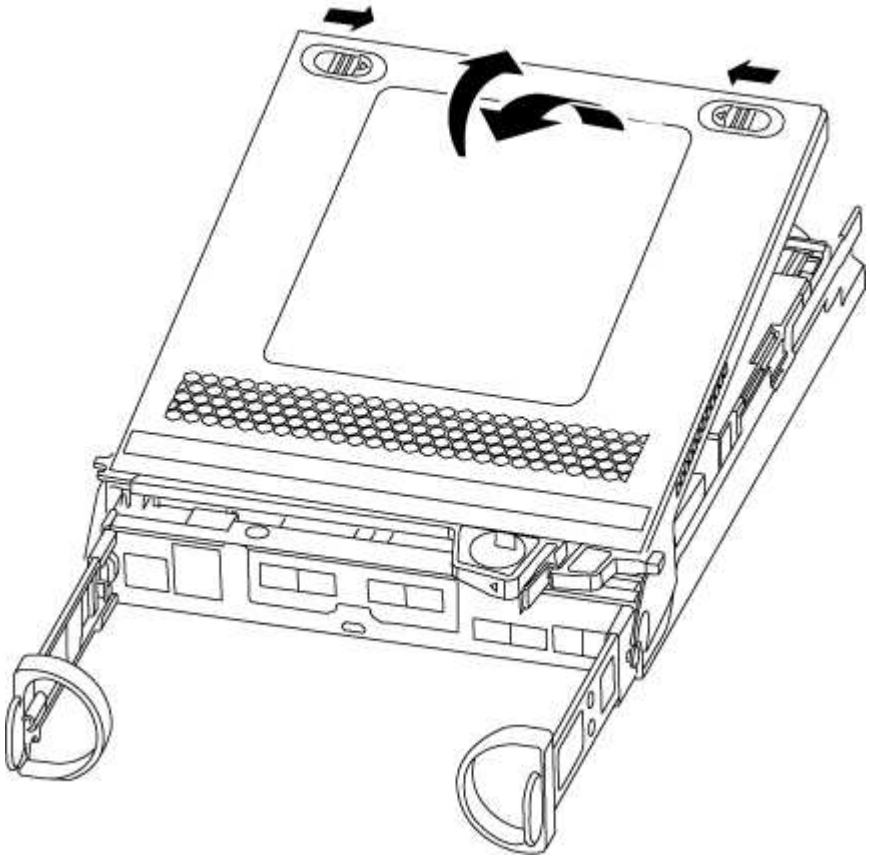
3. Remove and set aside the cable management devices from the left and right sides of the controller module.



4. Squeeze the latch on the cam handle until it releases, open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.



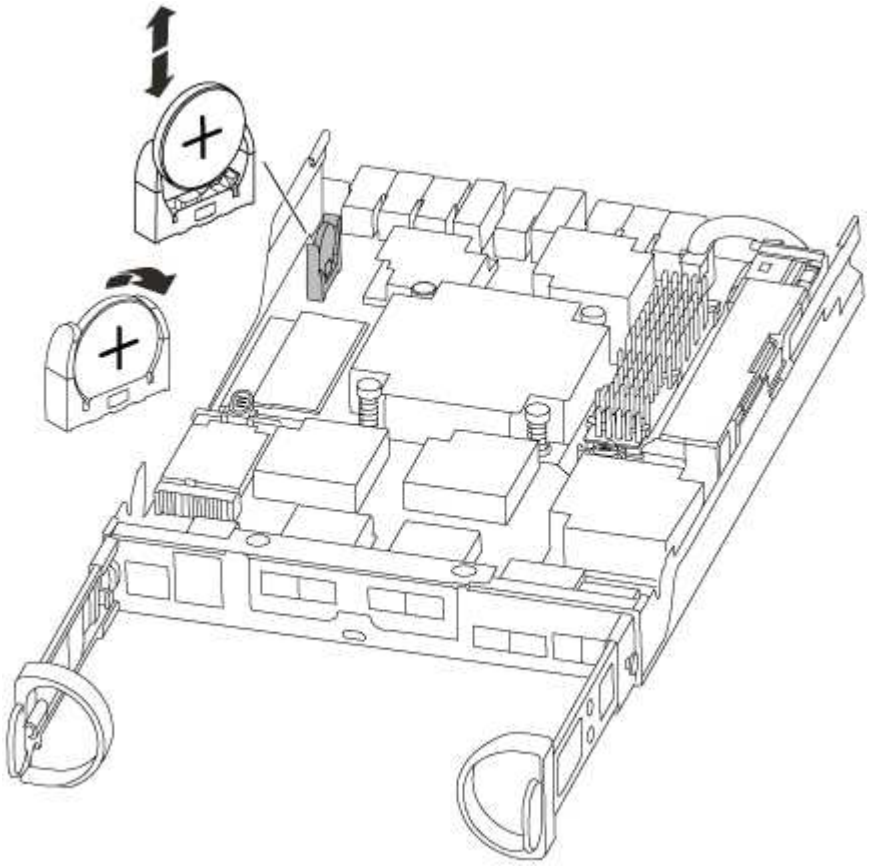
5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.



Step 3: Replace the RTC battery

To replace the RTC battery, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.



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. Locate the empty battery holder in the controller module.
6. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
7. 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 and set time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

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

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

3. Recable the system, as needed.

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

4. If the power supplies were unplugged, plug them back in and reinstall the power cable retainers.
5. Complete the reinstallation of the controller module:
   a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position.

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

   b. If you have not already done so, reinstall the cable management device.
   c. Bind the cables to the cable management device with the hook and loop strap.
   d. Reconnect the power cables to the power supplies and to the power sources, and then turn on the power to start the boot process.
   e. Halt the controller at the LOADER prompt.
6. Reset the time and date on the controller:
   a. Check the date and time on the healthy controller with the `show date` command.
   b. At the LOADER prompt on the target controller, check the time and date.
   c. If necessary, modify the date with the `set date mm/dd/yyyy` command.
   d. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.
   e. Confirm the date and time on the target controller.
7. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
8. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
9. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

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

```
cluster_B::>  metrocluster node show

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

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

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

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

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

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

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

6. Reestablish any SnapMirror or SnapVault configurations.

Step 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.

Key specifications for AFF A150

The following are select specifications for the AFF A150. Visit NetApp Hardware Universe (HWU) for a complete list of AFF A150 specifications. This page is reflective of a single high availability pair.

Key specifications for AFF A150

Platform Configuration: AFF A150 UTA2 Single Chassis HA Pair

Max Raw Capacity: 0.5472 PB

Memory: 64.0000 GB

Form Factor: 2U chassis with 2 HA controllers and 24 drive slots

ONTAP Version: b_startONTAP: 9.16.1P2b_end

PCIe Expansion Slots:

Minimum ONTAP Version: ONTAP 9.10.1P12

Scaleout Maximums

Type	HA Pairs	Raw Capacity	Max Memory
NAS	12	6.6 PB / 5.8 PiB	768 GB
SAN	6	3.3 PB / 2.9 PiB	384 GB
HA Pair		0.5 PB / 0.5 PiB	64.0000

IO

Onboard IO

Protocol	Ports
Ethernet 10 Gbps	12
FC 16 Gbps	8
SAS 12 Gbps	4

Total IO

Protocol	Ports
Ethernet 10 Gbps	12

FC 16 Gbps	8
SAS 12 Gbps	4

Management Ports

Protocol	Ports
Ethernet 1 Gbps	2
RS-232 115 Kbps	4
USB 12 Mbps	2

Storage Networking Supported

CIFS; FC; FCoE; iSCSI; NFS v3; NFS v4.0; NFS v4.1; NFS v4.2; NVMe/TCP; S3; S3 with NAS; SMB 2.0; SMB 2.1; SMB 2.x; SMB 3.0; SMB 3.1; SMB 3.1.1;

System Environment Specifications

- Typical Power: 1209 BTU/hr
- Worst-case Power: 1676 BTU/hr
- Weight: 60.8 lb 27.6 kg
- Height: 2U
- Width: 19” IEC rack-compliant (17.6” 44.7 cm)
- Depth: 19.0” (24.1” with cable management bracket)
- Operating Temp/Altitude/Humidity: 5°C to 45°C (41°F to 113°F) at up to 3048m (10000 ft) elevation;8% to 90% relative humidity, noncondensing
- Non-operating Temp/Humidity: -40°C to 70°C (-40°F to 158°F) up to 12192m (40000 ft) 10% to 95% relative humidity, noncondensing, in original container
- Acoustic Noise: Declared sound power (LwAd): 6.9; Sound pressure (LpAm) (bystander positions): 51.0 dB

Compliance

- Certifications EMC/EMI: AMCA, FCC, ICES, KC, Morocco, VCCI
- Certifications safety: BIS, CB, CSA, G_K_U-SoR, IRAM, NOM, NRCS, SONCAP, TBS
- Certifications Safety/EMC/EMI: EAC, UKRSEPRO
- Certifications Safety/EMC/EMI/RoHS: BSMI, CE DoC, UKCA DoC
- Standards EMC/EMI: BS-EN-55024, BS-EN55035, CISPR 32, EN55022, EN55024, EN55032, EN55035, EN61000-3-2, EN61000-3-3, FCC Part 15 Class A, ICES-003, KS C 9832, KS C 9835
- Standards Safety: ANSI/UL60950-1, ANSI/UL62368-1, BS-EN62368-1, CAN/CSA C22.2 No. 60950-1, CAN/CSA C22.2 No. 62368-1, CNS 14336, EN60825-1, EN62368-1, IEC 62368-1, IEC60950-1, IS 13252(part 1)

High Availability

Ethernet based baseboard management controller (BMC) and ONTAP management interface; Redundant hot-swappable controllers; Redundant hot-swappable power supplies; SAS in-band management over SAS

connections;

Copyright information

Copyright © 2026 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at http://www.netapp.com/TM are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.