FAS systems

Install and maintain

NetApp
January 09, 2026

This PDF was generated from https://docs.netapp.com/us-en/ontap-systems/fas50/install-overview.html on January 09, 2026. Always check docs.netapp.com for the latest.

Table of Contents

FAS systems . . . 1
- FAS50 systems . . . 1
  - Install and setup . . . 1
  - Maintain . . . 18
  - Key specifications for FAS50 . . . 148
- FAS70 and FAS90 systems . . . 151
  - Install and setup . . . 151
  - Maintain . . . 171
  - Key specifications . . . 286
- FAS2700 systems . . . 292
  - Install and setup . . . 292
  - Maintain . . . 308
  - Key specifications for FAS2750 . . . 393
- FAS2820 systems . . . 396
  - Install and setup . . . 396
  - Maintain . . . 409
  - Key specifications for FAS2820 . . . 493
- FAS8300 and FAS8700 systems . . . 496
  - Install and setup . . . 496
  - Maintain . . . 506
  - Key specifications . . . 639
- FAS9500 systems . . . 645
  - Install and setup . . . 645
  - Maintain . . . 661
  - Key specifications for FAS9500 . . . 754

FAS systems

FAS50 systems

Install and setup

Installation and setup workflow - FAS50

To install and set up your FAS50 storage system, you must review the installation requirements, prepare your site, install and cable the hardware components, power on the storage system, and set up the ONTAP cluster.



Review the installation requirements

Before installing your storage system, it must meet the installation requirements.



Prepare for installation

To prepare for installation, get the site ready, check environmental and electrical requirements, and ensure there's enough rack space. Then, unpack the equipment, compare contents to the packing slip, and register the hardware to access support benefits.



Install the hardware

To install the hardware, install the rail kits for your storage system and shelves, and then install and secure your storage system and shelves in the cabinet or telco rack.



Cable the hardware

To cable the hardware, connect the controllers to your network and then to your shelves.



Power on the storage system

To power on your storage system, power on each shelf and assign a unique shelf ID as needed, and then power on the controllers.



Set up your cluster

After you've powered on your storage system, you set up your cluster.

Installation requirements - FAS50

Review the requirements for your FAS50 storage system.

Equipment needed for install

To install your storage system, you need the following equipment and tools.

- Access to a Web browser to configure your storage system
- Electrostatic discharge (ESD) strap
- Flashlight
- Laptop or console with a USB/serial connection
- Phillips #2 screwdriver

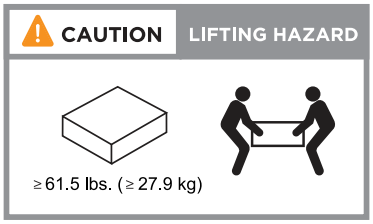
Lifting precautions

Storage systems and shelves are heavy. Exercise caution when lifting and moving these items.

Storage system weight

Take the necessary precautions when moving or lifting your storage system.

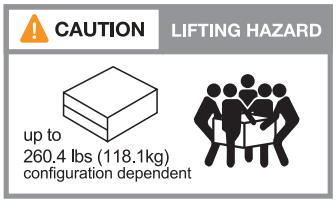
The storage system can weigh up to 53.8 lbs (24.4 kg). To lift the storage system, use two people or a hydraulic lift.



Shelf weight

Take the necessary precautions when moving or lifting your shelf.

A DS460C shelf can weigh up to 260.4 lbs (181.1 kg). To lift the shelf, you might need up to five people or a hydraulic lift. Keep all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.



Related information

- Safety information and regulatory notices

What's next?

After you've reviewed the installation requirements and considerations for your storage system, you prepare for installation.

Prepare to install - FAS50

Prepare to install your FAS50 storage system by getting the site ready, unpacking the boxes and comparing the contents of the boxes to the packing slip, and registering the

storage system to access support benefits.

Step 1: Prepare the site

To install your storage system, ensure that the site and the cabinet or rack that you plan to use meet specifications for your configuration.

Steps

1. Use NetApp Hardware Universe to confirm that your site meets the environmental and electrical requirements for your storage system.
2. Make sure you have adequate cabinet or rack space for your storage system, shelves, and any switches:
3. Install any required network switches.

   See the Switch documentation for installation instructions and NetApp Hardware Universe for compatibility information.

Step 2: Unpack the boxes

After you've ensured that the site and the cabinet or rack that you plan to use for your storage system meet the required specifications, unpack all boxes and compare the contents to the items on the packing slip.

Steps

1. Carefully open all the boxes and lay out the contents in an organized manner.
2. Compare the contents you've unpacked with the list on the packing slip.

   You can get your packing list by scanning the QR code on the side of the shipping carton.

   The following items are some of the contents you might see in the boxes.

   Ensure that everything in the boxes matches the list on the packing slip. If there are any discrepancies, note them down for further action.

Hardware

- Bezel
- Storage system
- Rail kits with instructions (optional)
- Storage shelf (if you ordered additional storage)

Cables

- Management Ethernet cables (RJ-45 cables)
- Network cables
- Power cords
- Storage cables (if you ordered additional storage)
- USB-C serial console cable

Step 3: Register your storage system

After you've ensured that your site meets the requirements for your storage system specifications, and you've verified that you have all the parts you ordered, you should register your storage system.

Steps

1. Locate the System Serial Numbers (SSN) for every controller being installed.

   You can find the serial numbers in the following locations:

   - On the packing slip
   - In your confirmation email

2. Go to the NetApp Support Site.
3. Determine whether you need to register your storage system:

If you are a…	Follow these steps…
Existing NetApp customer	a. Sign in with your username and password. b. Select **Systems** > **My Systems**. c. Confirm that the new serial numbers are listed. d. If it is not, follow the instructions for new NetApp customers.
New NetApp customer	a. Click **Register Now**, and create an account. b. Select **Systems** > **Register Systems**. c. Enter the storage system's serial numbers and requested details. After your registration is approved, you can download any required software. The approval process might take up to 24 hours.

What's next?

After you've prepared to install your storage system, you install the hardware for your storage system.

Install the hardware - FAS50

After you prepare to install your FAS50 storage system, install the hardware for the storage system. First, install the rail kits. Then install and secure your storage system in a cabinet or telco rack.

Skip this step if your storage system came in a cabinet.

Before you begin

- Make sure you have the instructions packaged with the rail kit.
- Be aware of the safety concerns associated with the weight of the storage system and shelf.
- Understand that the airflow through the storage system enters from the front where the bezel or end caps are installed and exhausts out the rear where the ports are located.

Steps

1. Install the rail kits for your storage system and shelves as needed, using the instructions included with the kits.

2. Install and secure your storage system in the cabinet or telco rack:
   a. Position the storage system onto the rails in the middle of the cabinet or telco rack, and then support the storage system from the bottom and slide it into place.
   b. Make sure that the guiding pins of the cabinet or telco rack are securley in the chassis guide slots.
   c. Secure the storage system to the cabinet or telco rack using the included mounting screws.
3. Attach the bezel to the front of the storage system.
4. Install and secure the shelf as needed.
   a. Position the back of the shelf onto the rails, and then support the shelf from the bottom and slide it into the cabinet or telco rack.

      If you are installing multiple shelves, place the first shelf directly above the controllers. Place the second shelf directly under the controllers. Repeat this pattern for any additional shelves.

   b. Secure the shelf to the cabinet or telco rack using the included mounting screws.

      What's next?

      After you've installed the hardware for your storage system, you cable the hardware.

Cable the hardware - FAS50

After you install your FAS50 storage system hardware, cable the controllers to the network and shelves.

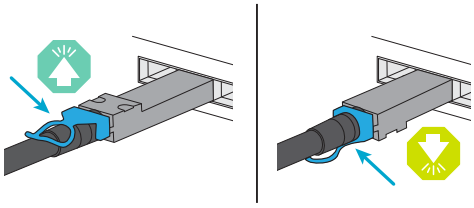
Before you begin

Contact your network administrator for information about connecting the storage system to your network switches.

About this task

- The cabling graphics have arrow icons showing the proper orientation (up or down) of the cable connector pull-tab when inserting a connector into a port.

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



- If cabling to an optical switch, insert the optical transceiver into the controller port before cabling to the switch port.

Step 1: Cable the cluster/HA connections

Create the ONTAP cluster connections. For switchless clusters, connect the controllers to each other. For switched clusters, connect the controllers to the cluster network switches.

The cluster/HA cabling examples show common configurations.

If you do not see your configuration here, go to NetApp Hardware Universe for comprehensive configuration and slot priority information to cable your storage system.

Switchless cluster cabling

FAS50 with one 2-port 40/100 GbE I/O module

Steps

1. Cable the Cluster/HA interconnect connections:

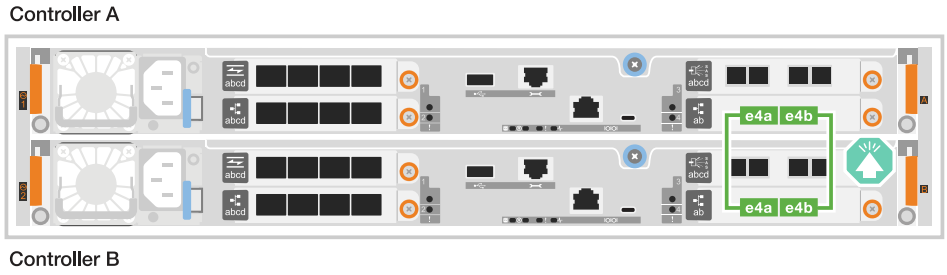
   The cluster interconnect traffic and the HA traffic share the same physical ports (on the I/O module in slot 4). The ports are 40/100 GbE.

   a. Cable controller A port e4a to controller B port e4a.

   b. Cable controller A port e4b to controller B port e4b.

   100 GbE Cluster/HA interconnect cables





Switched cluster cabling

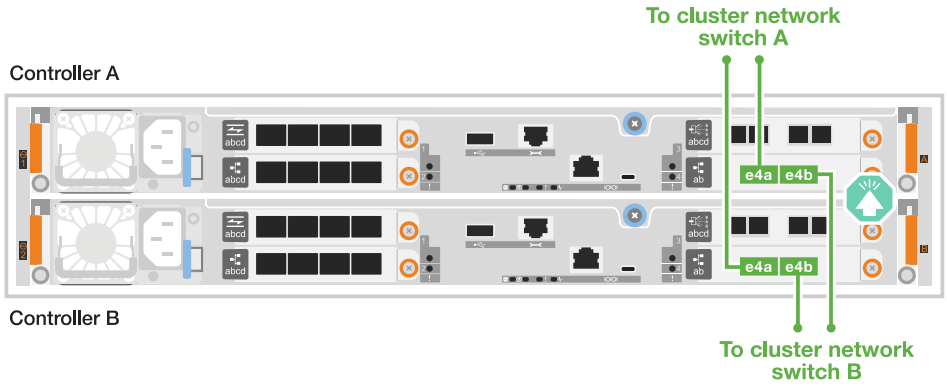
FAS50 with one 2-port 40/100 GbE I/O module

1. Cable the controllers to the cluster network switches:

   The cluster interconnect traffic and the HA traffic share the same physical ports(on the I/O module in slot 4). The ports are 40/100 GbE.

   a. Cable controller A port e4a to cluster network switch A.
   b. Cable controller A port e4b to cluster network switch B.
   c. Cable controller B port e4a to cluster network switch A.
   d. Cable controller B port e4b to cluster network switch B.

40/100 GbE Cluster/HA interconnect cables



Step 2: Cable the host network connections

Cable the controllers to your Ethernet or FC host network.

The host network cabling examples show common configurations.

If you do not see your configuration here, go to NetApp Hardware Universe for comprehensive configuration and slot priority information to cable your storage system.

Ethernet host cabling

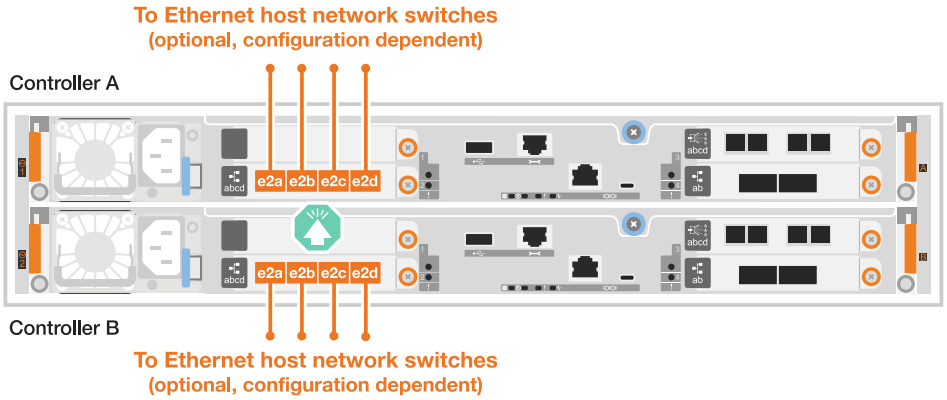
FAS50 with one 4-port 10/25 GbE I/O module

Steps

1. On each controller, cable ports e2a, e2b, e2c and e2d to the Ethernet host network switches.

10/25 GbE cables





FC host cabling

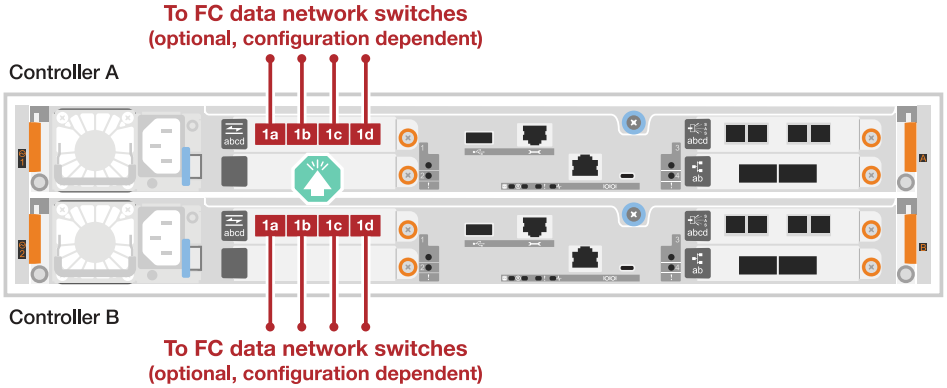
FAS50 with one 4-port 64 Gb/s FC I/O module

Steps

1. On each controller, cable ports 1a, 1b, 1c and 1d to the FC host network switches.

64 Gb/s FC cables





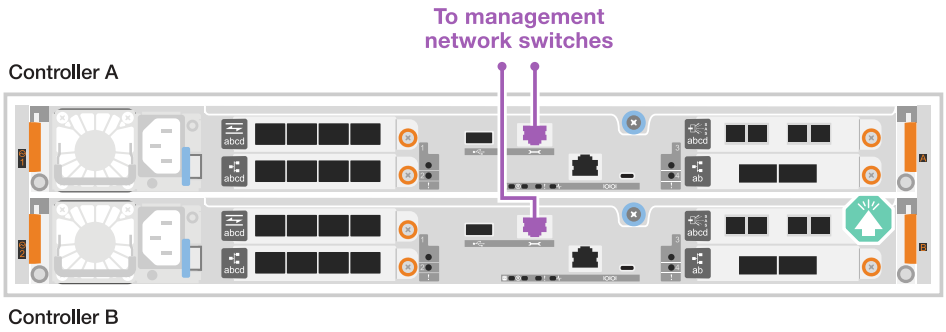
Step 3: Cable the management network connections

Cable the controllers to your management network.

1. Cable the management (wrench) ports on each controller to the management network switches.

1000BASE-T RJ-45 cables





Do not plug in the power cords yet.

Step 4: Cable the shelf connections

The following procedures show how to cable the controllers to one or two DS460C shelves.

About this task

- The cabling examples show DS460C shelves; however, additional SAS shelves are supported, see NetApp Hardware Universe.

  Cabling for other supported SAS shelves is similar. See Install and cable shelves for a new system installation. You can also refer to SAS cabling rues, worksheets, and examples overview.

- For the maximum number of shelves supported for your storage system and all of your cabling options, such as optical and switch-attached, see NetApp Hardware Universe.
- The graphics show controller A cabling in blue and controller B cabling in yellow.
- You use the storage cables that came with your storage system, which could be the following cable type:

  mini-SAS HD cable

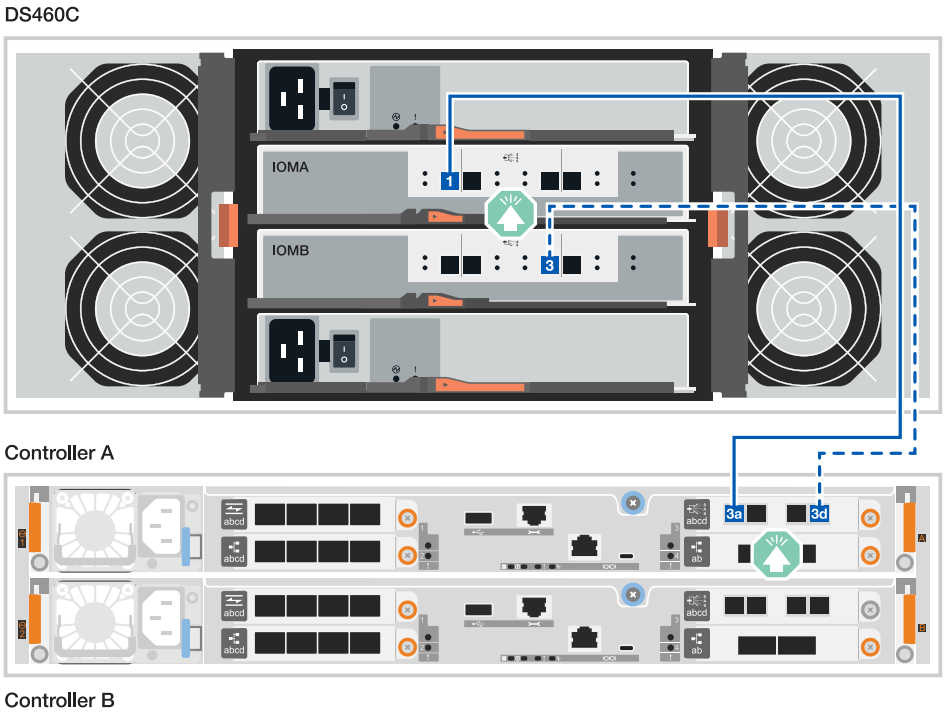


Option 1: One DS460C shelf

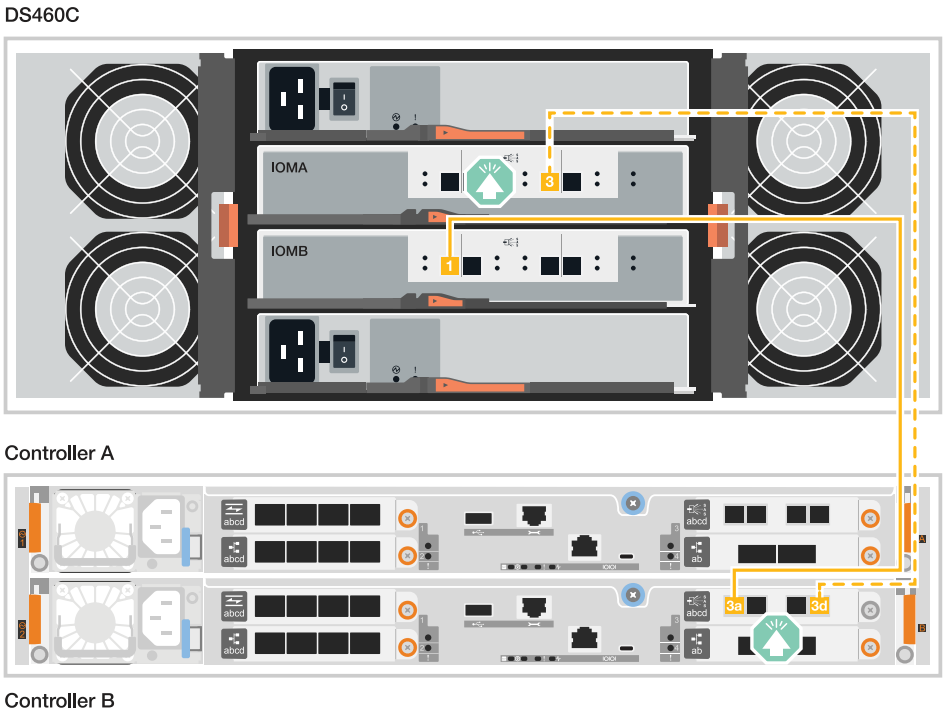
Cable each controller to each IOM12 module on the DS460C shelf.

Steps

1. Cable controller A to the shelf:
   a. Cable controller A port 3a to IOMA port 1.
   b. Cable controller A port 3d to IOMB port 3.



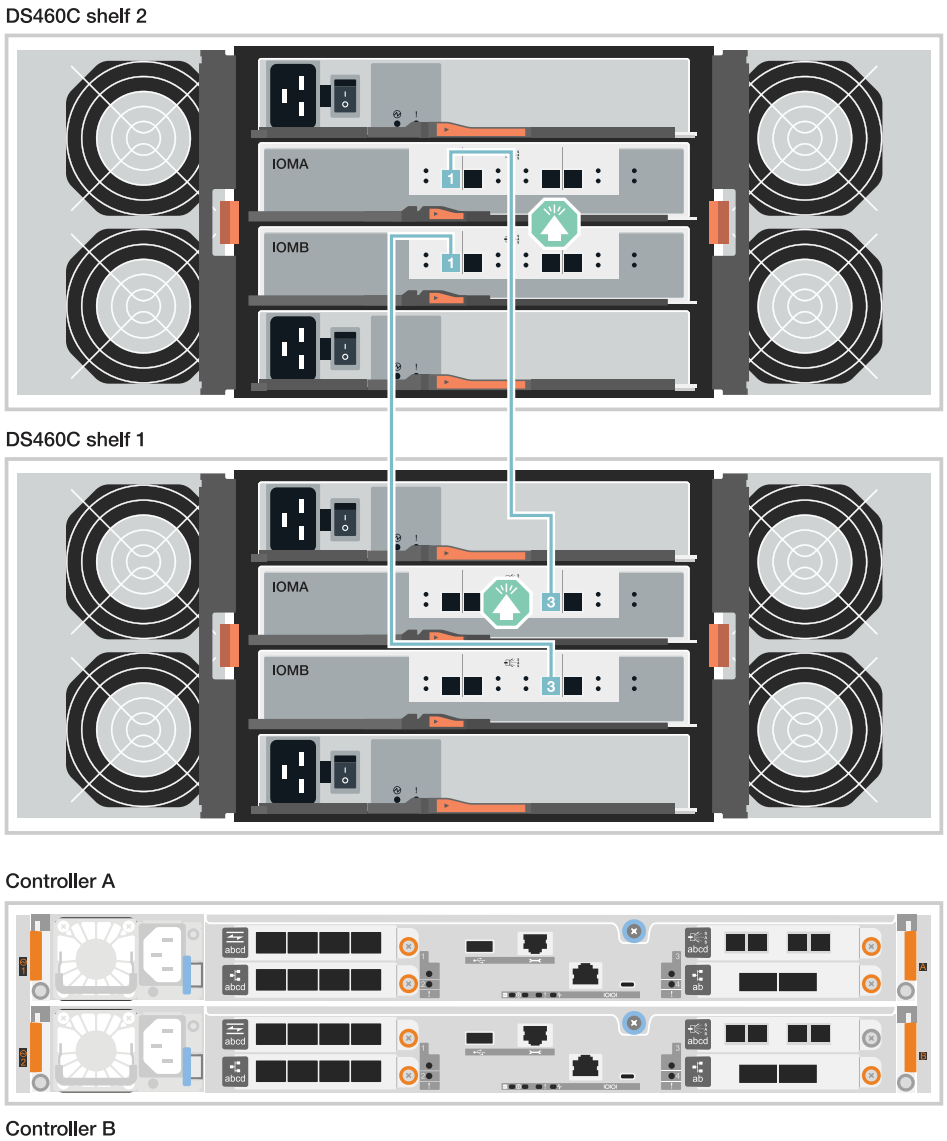
2. Cable controller B to the shelf:
   a. Cable controller B port 3a to IOMB port 1.
   b. Cable controller B port 3d to IOMA port 3.



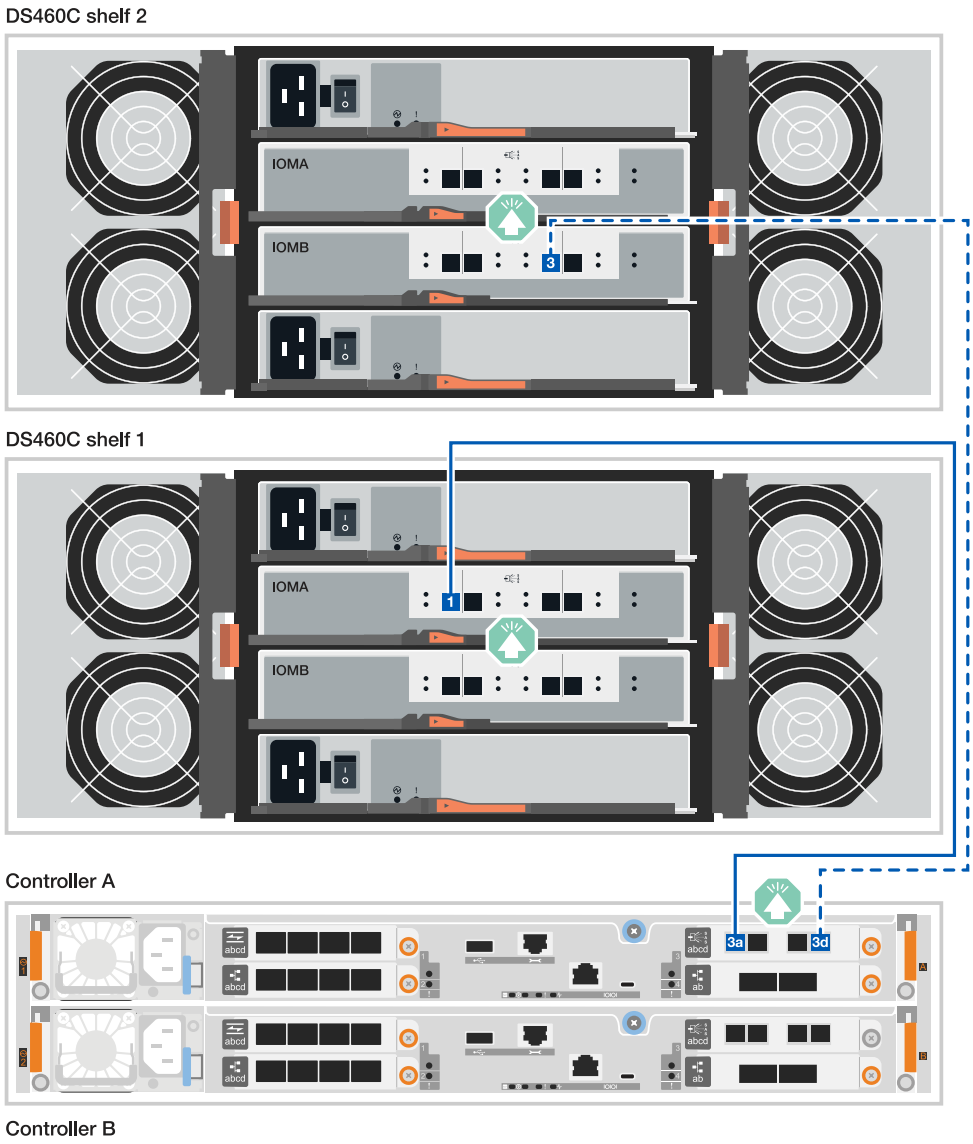
Option 2: Two DS460C shelves

Connect each controller to the IOM12 modules on both DS460C shelves.

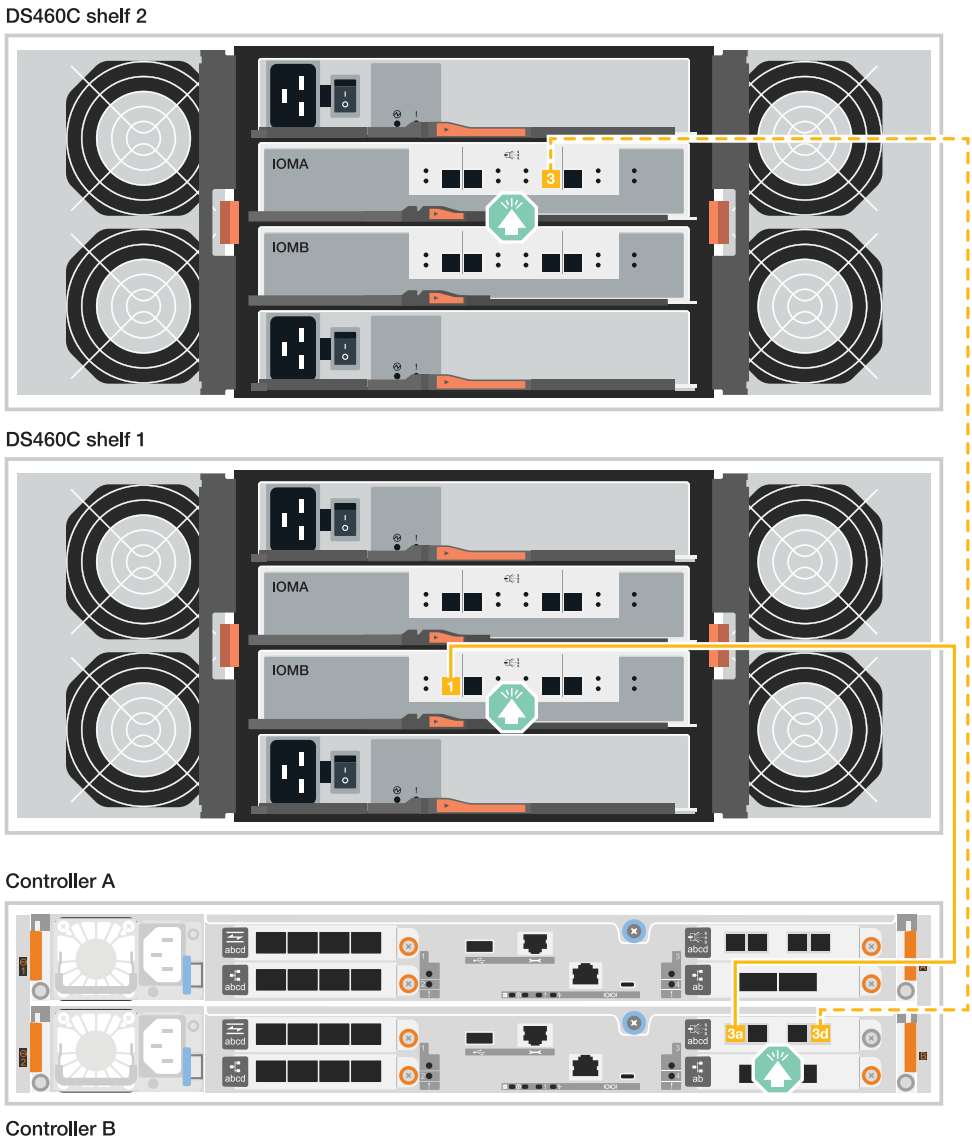
1. Cable the shelf-to-shelf connections:
   a. Cable Shelf 1 IOMA port 3 to Shelf 2 IOMA port 1.
   b. Cable Shelf 1 IOMB port 3 to Shelf 2 IOMB port 1.



2. Cable controller A to the shelves:
   a. Cable controller A port 3a to shelf 1 IOMA port 1.
   b. Cable controller A port 3d to shelf 2 IOMB port 3.



3. Cable controller B to the shelves:
   a. Cable controller B port 3a to shelf 1 IOMB port 1.
   b. Cable controller B port 3d to shelf 2 IOMA port 3.



What's next?

After you've cabled the hardware for your storage system, you power on the storage system.

Power on the storage system - FAS50

After you cable the controllers to the network and shelves in your FAS50 storage system, you power on your shelves and controllers.

Step 1: Power on the shelf and assign shelf ID

Each shelf is distinguished by a unique shelf ID. This ID ensures that the shelf is distinct within your storage system setup.

About this task

- A valid shelf ID is 01 through 99.

  If you have internal shelves (storage), which are integrated within the controllers, they are assigned a fixed

shelf ID of 00.

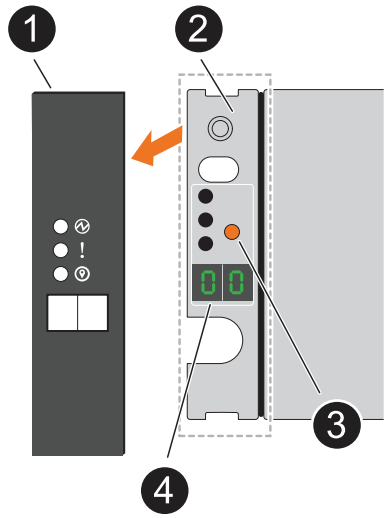
- You must power cycle a shelf (turn off the power switch on each of the power supplies of the SAS shelf, wait the appropriate amount of time, and then switch the power back on) for the shelf ID to take effect.

Steps

1. Power on the shelf by connecting the power cords first to the shelf, securing them in place with the power cord retainer, connecting the power cords to power sources on different circuits, and then turning on the power switch on each of the power supplies (at the rear of the shelf).

   The shelf powers on and boots automatically when powered on.

2. Remove the left end cap to access the orange shelf ID button on the faceplate.



1	Shelf end cap
2	Shelf faceplate
3	Shelf ID button
4	Shelf ID number

3. Change the first number of the shelf ID:
   a. Press and hold the shelf ID button until the first number on the digital display blinks, and then release the button.

      It can take up to 15 seconds for the number to blink. This activates the shelf ID programming mode.

      If the ID takes longer than 15 seconds to blink, press and hold the shelf ID button again, making sure to press it in all the way.

   b. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

Each press and release duration can be as short as one second.

The first number continues to blink.

4. Change the second number of the shelf ID:
   a. Press and hold the button until the second number on the digital display blinks.

      It can take up to three seconds for the number to blink.

      The first number on the digital display stops blinking.

   b. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

      The second number continues to blink.

5. Lock in the desired number and exit the programming mode by pressing and holding the shelf ID button until the second number stops blinking.

   It can take up to three seconds for the number to stop blinking.

   Both numbers on the digital display start blinking and the amber LED illuminates after about five seconds, alerting you that the pending shelf ID has not yet taken effect.

6. Power-cycle the shelf for at least 10 seconds to make the shelf ID take effect.
   a. Turn off the power switch on each of the power supplies.
   b. Wait 10 seconds.
   c. Turn on the power switch on each of the power supplies to complete the power cycle.

      When a power supply is powered on, the bicolored LED should illuminate green.

7. Replace the left end cap.

Step 2: Power on the controllers

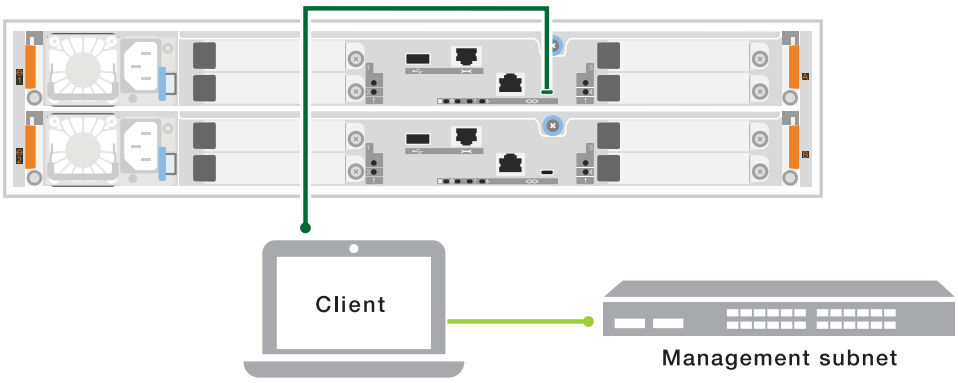
After you've powered on your shelves and assigned them unique IDs, power on the storage controllers.

Steps

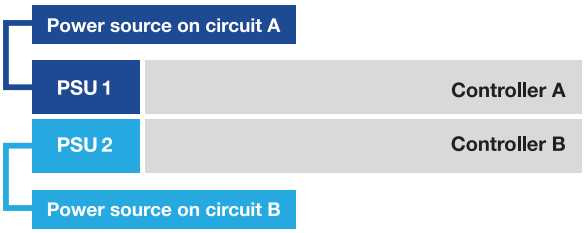
1. Connect your laptop to the serial console port. This will allow you to monitor the boot sequence when the controllers are powered on.
   a. Set the serial console port on the laptop to 115,200 baud with N-8-1.

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

   b. Using the console cable provided with your storage system, connect one end of the console cable to your laptop and the other end to the serial console port on controller A.
   c. Connect the laptop to the switch on the management subnet.



2. Assign a TCP/IP address to the laptop, using one that is on the management subnet.
3. Plug the two power cords into the controller power supplies, and then connect them to power sources on different circuits.



   - The system begins to boot. Initial booting might take up to eight minutes.
   - The LEDs flash on and the fans start, which indicates that the controllers are powering on.
   - The fans might be very noisy when they first start up. The fan noise during start-up is normal.
   - The shelf ID display on the front of the system chassis does not illuminate.
4. Secure the power cords using the securing device on each power supply.

What's next?

After you've powered on your storage system, you set up your cluster.

Maintain

Overview of hardware maintenance - FAS50

Maintain the hardware of your FAS50 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the FAS50 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the FAS50 storage systems, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot the image from the partner node and automatically run the appropriate boot menu option to install the boot image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot the image from a USB drive and restore the configuration from the partner node.
Caching module	A caching module (Flash Cache module) utilizes high-speed SSDs to store frequently accessed data for quicker retrieval.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Drive	A drive is a device that provides the physical storage needed for data.
Fan	A fan cools the controller and drives.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
NV battery	The non-volatile memory (NV) battery is responsible for providing power to the NVMEM components while data in-flight is being destaged to flash memory after a power loss.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automatic recovery workflow - FAS50

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your FAS50 storage system.

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

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

1 **Review the boot media requirements**

Review the requirements for boot media replacement.

2 **Shut down the controller**

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

3 **Replace the boot media**

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

4 **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5 **Return the failed part to NetApp**

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

Requirements for automated boot media recovery - FAS50

Before replacing the boot media in your FAS50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0M (wrench) port on the impaired controller is working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

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

Review the following requirements.

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

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

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:

- /cfcard/kmip/servers.cfg file.
- /cfcard/kmip/certs/client.crt file.
- /cfcard/kmip/certs/client.key file.
- /cfcard/kmip/certs/CA.pem file.

- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - FAS50

Shut down the impaired controller in your FAS50 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

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

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

About this task

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

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

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

Steps

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

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

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

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

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - FAS50

The boot media in your FAS50 storage system stores essential firmware and configuration data. The replacement process involves removing the controller module, removing the impaired boot media, installing the replacement boot media, and then reinstalling the controller module.

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

About this task

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

A platform chassis has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Step 1: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

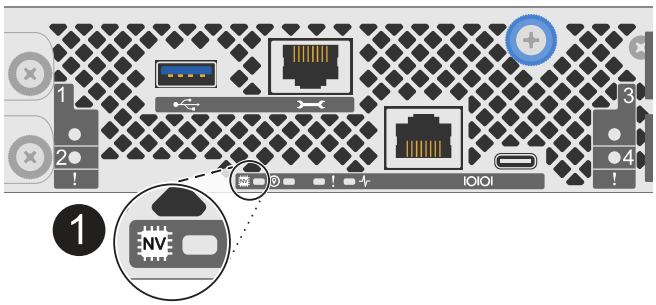
Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

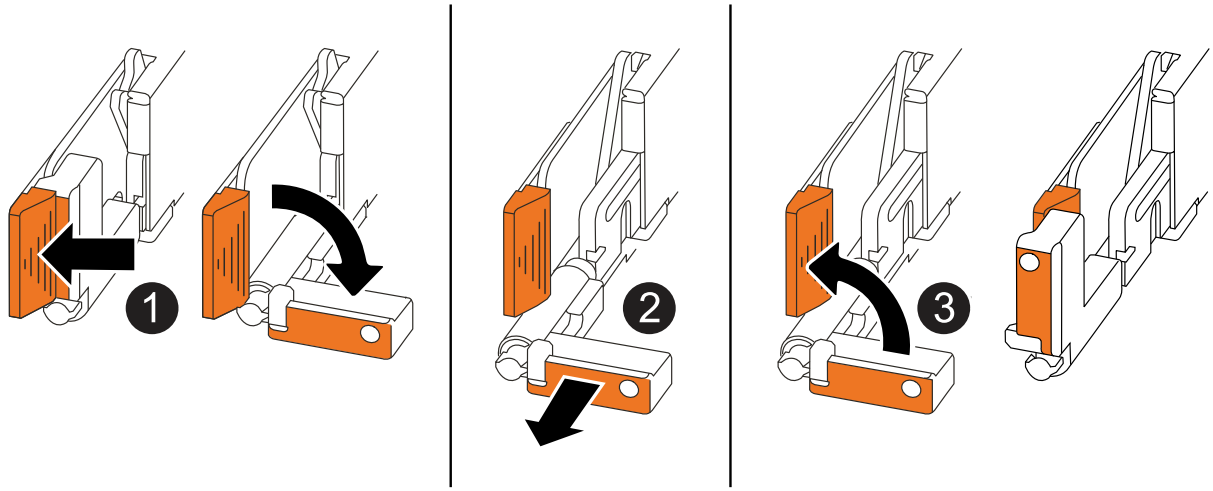
If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



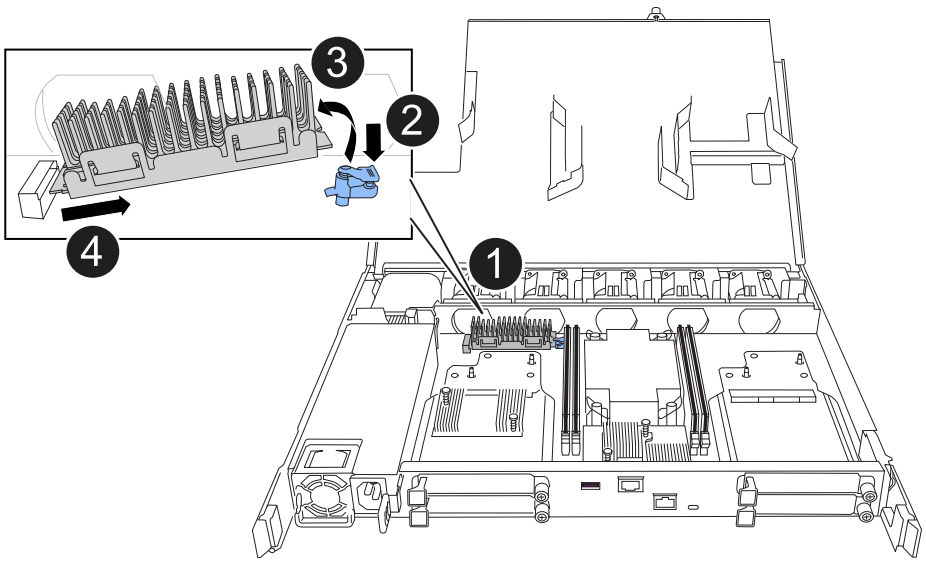
1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Place the controller on an anti-static mat.
7. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Replace the boot media

To replace the boot media, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the boot media:



1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

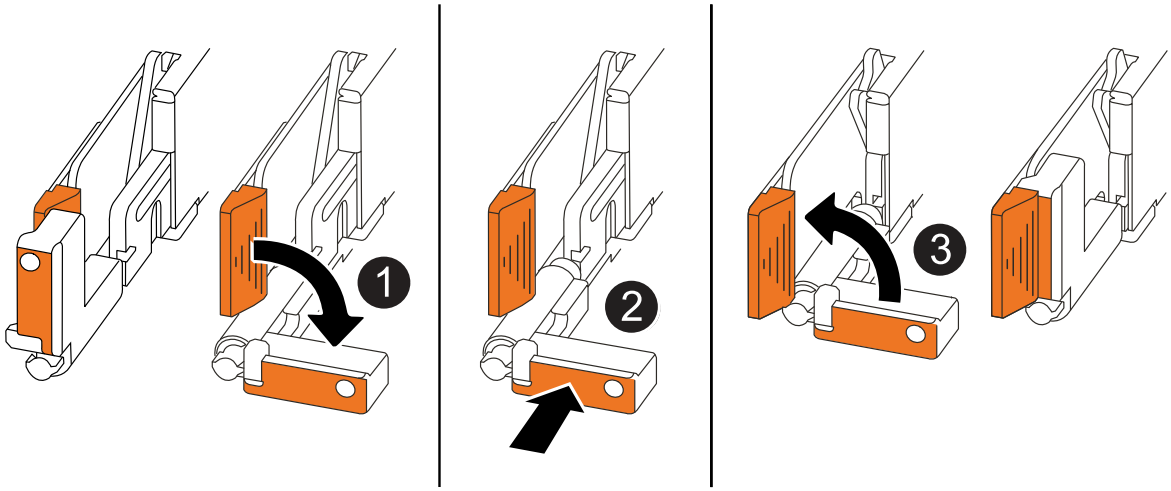
3. Install the replacement boot media:
   a. Remove the boot media from its package.
   b. Slide the socket end of the boot media into its socket.
   c. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 3: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   > Do not completely insert the controller in the chassis until instructed to do so later in this procedure.

3. Reconnect the cables to the controller; however, do not plug in the power cord to the power supply (PSU) at this time.

   > Make sure that the console cable is connected to the controller because you want to catch and log the boot sequence later in the boot media replacement procedure when you fully seat the controller in the chassis and it begins to boot.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

      > The controller boots to the LOADER prompt when fully seated in the chassis. It gets its power from the partner controller.

b. Rotate the controller handles up and lock in place with the tabs.

5. Reconnect the power cord to the PSU on the impaired controller.

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

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - FAS50

After installing the new boot media device in your FAS50 storage system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

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

Before you begin

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

Steps

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

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

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

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

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

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. a. Wait for the login prompt to display. b. Log into the node and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* c. Go to re-enabling automatic giveback if it was disabled.
`key manager is configured.`	Encryption is installed. Go to restoring the key manager.

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

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

Onboard Key Manager (OKM)

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

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

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

a. Enter `y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the passphrase for onboard key management when prompted.

c. Enter the passphrase again when prompted to confirm.

d. Enter the backup data for onboard key manager when prompted.

Show example of passphrase and backup data prompts

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

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

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

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

Successfully recovered keymanager secrets.
```

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

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

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

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

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

Go to re-enabling automatic giveback if it was disabled.

External Key Manager (EKM)

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

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

a. Enter the EKM configuration settings when prompted:

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

Show example of client certificate contents

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

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

Show example of client key file contents

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

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

Show example of KMIP server file contents

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

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

Show example of server configuration file contents

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

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

Show example of ONTAP Cluster UUID prompt

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


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

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

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

Show example of temporary network setting prompts

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

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

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

b. Verify the key restoration status:

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

Show example of key recovery error and warning messages

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

Uptime: 11m32s
System halting...

LOADER-B>
```

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

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

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

Go to re-enabling automatic giveback if it was disabled.

5. If automatic giveback was disabled, reenable it:

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

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

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

What's next

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

Return the failed boot media part to NetApp - FAS50

If a component in your FAS50 storage system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - FAS50

The manual recovery of the boot image involves using a USB drive to reinstall ONTAP onto the FAS50 system's replacement boot media. You must download the appropriate ONTAP recovery image from the NetApp Support Site and copy it to a USB drive. This prepared USB drive is then used to perform the recovery and restore the system to operational status.

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

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



Review the boot media requirements

Review the requirements for replacing the boot media.



Check onboard encryption keys

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



Shut down the impaired controller

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



Replace the boot media

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



Boot the recovery image

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



Restore encryption

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



Return the failed part to NetApp

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

Requirements for manual boot media recovery - FAS50

Before replacing the boot media in your FAS50 storage system, ensure you meet the

necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

USB flash drive

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

File preparation

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

Component replacement

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

Controller identification

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

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

What's next?

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

Check encryption support for manual boot media recovery - FAS50

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.

What's next?

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

Shut down the controller for manual boot media recovery - FAS50

Shut down the impaired controller in your FAS50 storage system to prevent data loss and maintain system stability during the manual boot media recovery process.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

What's next?

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

Replace the boot media and prepare for manual boot recovery - FAS50

The boot media in your FAS50 system stores essential firmware and configuration data. The replacement process involves removing the controller module, removing the impaired boot media, installing the replacement boot media, and then manually transferring the ONTAP image to the replacement boot media using a USB flash drive.

About this task

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

A platform chassis has three location LEDs: one on the operator display panel and one on each controller. Location LEDs remain illuminated for 30 minutes.

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

Step 1: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

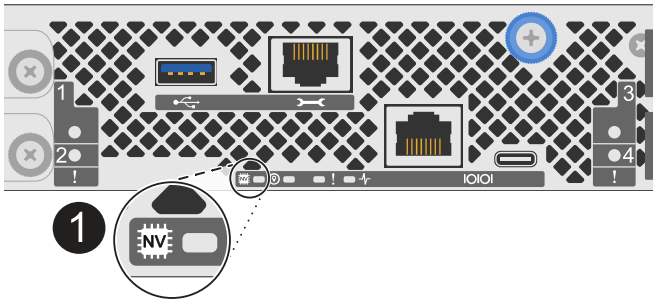
Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

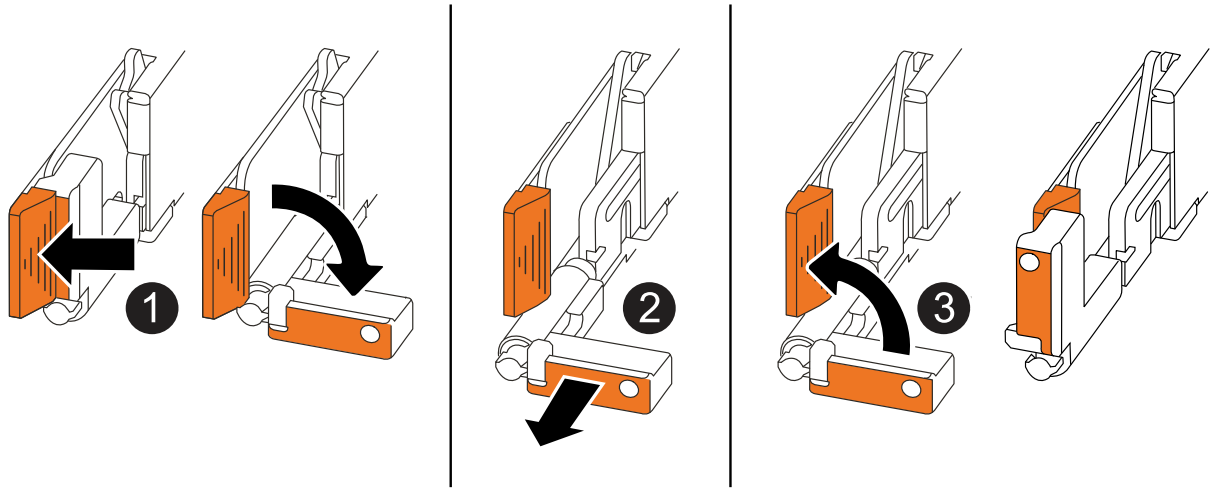
If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



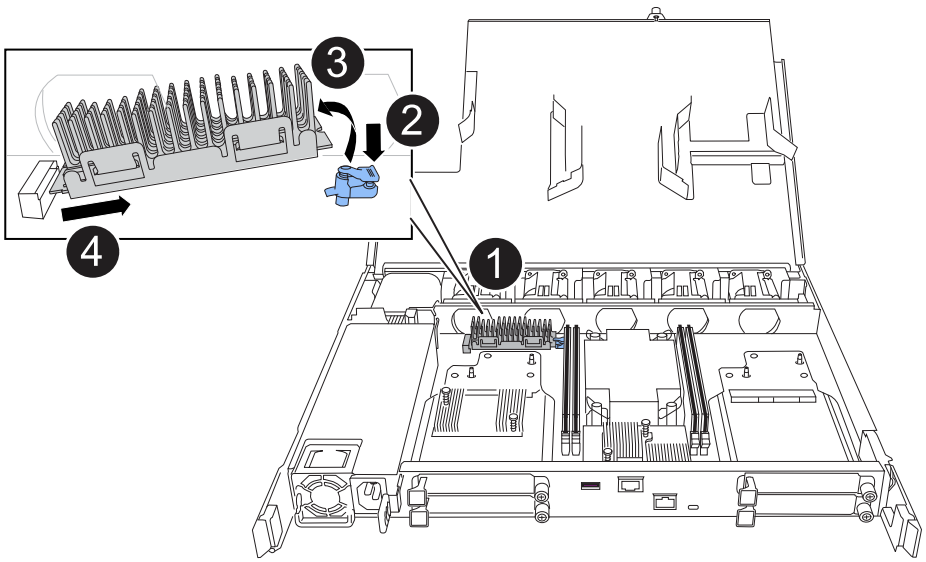
1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Place the controller on an anti-static mat.
7. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Replace the boot media

To replace the boot media, locate it inside the controller and follow the specific sequence of steps.

1. If you are not already grounded, properly ground yourself.
2. Remove the boot media:



1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

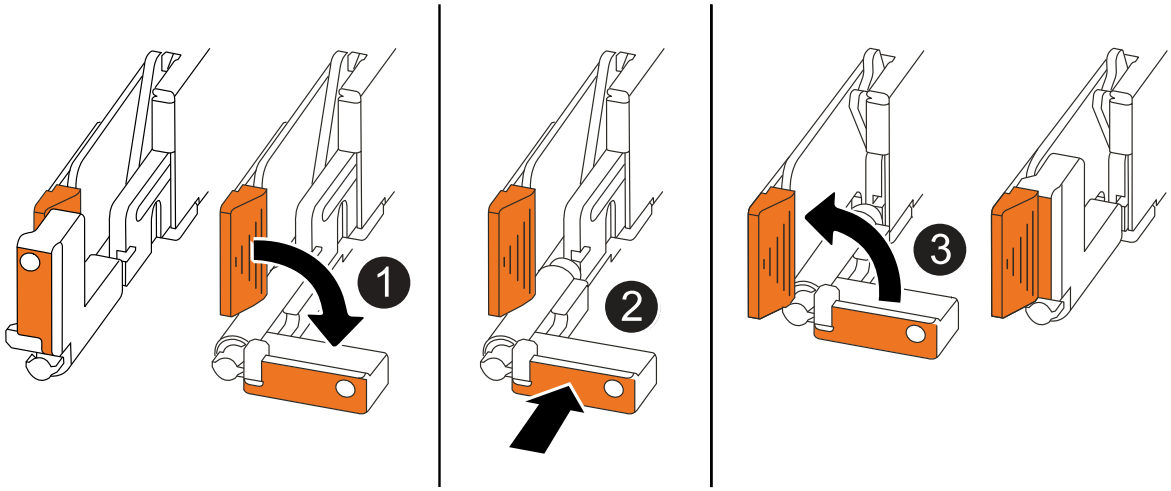
3. Install the replacement boot media:
   a. Remove the boot media from its package.
   b. Slide the socket end of the boot media into its socket.
   c. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 3: Reinstall the controller

Reinstall the controller into the chassis, but do not reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

   > Do not completely insert the controller in the chassis until instructed to do so later in this procedure.

3. Reconnect the cables to the controller; however, do not plug in the power cord to the power supply (PSU) at this time.

   > Make sure that the console cable is connected to the controller because you want to catch and log the boot sequence later in the boot media replacement procedure when you fully seat the controller in the chassis and it begins to boot.

Step 4: Transfer the boot image to the boot media

The replacement boot media that you installed is without an ONTAP image so you need to transfer an ONTAP image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- You must have 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 supported, download the image with NetApp Volume Encryption, as indicated in the download button.
- If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.

- You must have a network connection between the node management ports of the controllers (typically the e0M interfaces).

Steps

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

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

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

   c. Remove the USB flash drive from your laptop.
2. Insert the USB flash drive into the USB-A port on the impaired controller.

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

3. Fully seat the impaired controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

      The controller boots when fully seated in the chassis. It gets its power from the partner controller.

   b. Rotate the controller handles up and lock in place with the tabs.
4. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

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

5. Reconnect the power cord to the power supply (PSU) on the impaired controller.

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

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.

If you are reconnecting a…	Then…
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next?

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

Manual boot media recovery from a USB drive - FAS50

After installing the new boot media device in your FAS50 storage system, you can boot the recovery image manually from a USB drive to restore the configuration from the partner node.

Before you begin

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

Steps

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

   ```
   boot_recovery
   ```

   The recovery image is downloaded from the USB flash drive.

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

ONTAP 9.16.0 or earlier

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

a. **On the impaired controller:** Press `Y` when you see `Do you want to restore the backup configuration now?`

b. **On the impaired controller:** If prompted, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. **On the partner controller:** Set the impaired controller to advanced privilege level:

```
set -privilege advanced
```

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

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

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

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

```
set -privilege admin
```

f. **On the impaired controller:** Press `Y` when you see `Was the restore backup procedure successful?`

g. **On the impaired controller:** Press `Y` when you see `...would you like to use this restored copy now?`

h. **On the impaired controller:** Press `Y` when prompted to reboot, then press `Ctrl-C` when you see the Boot Menu.

i. **On the impaired controller:** Do one of the following:

- If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.
- If the system uses encryption, go to Restore encryption.

ONTAP 9.16.1 or later

Complete the following steps on the impaired controller:

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

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

b. Press `Y` when prompted to confirm that the restore backup was successful.

c. Press `Y` when prompted to use the restored configuration.

d. Press `Y` when prompted to reboot the node.

e. Press `Y` when prompted to reboot again, then press `Ctrl-C` when you see the Boot Menu.

f. Do one of the following:

- If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.

- If the system uses encryption, go to Restore encryption.

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

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

6. If you disabled automatic giveback, reenable it:

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

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

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

What's next?

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

Restore encryption keys after manual boot recovery - FAS50

Restore encryption on the replacement boot media in your FAS50 storage system to ensure continued data protection. The replacement process involves verifying key availability, reapplying encryption settings, and confirming secure access to your data.

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
```

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

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

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

   The recovery process completes and displays the following message:

   ```
   Successfully recovered keymanager secrets.
   ```

Show example prompt

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

Successfully recovered keymanager secrets.

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

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

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

Show example prompt

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


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

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

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

   On the partner controller:

8. Giveback the impaired controller:

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

   On the impaired controller:

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

   `security key-manager onboard sync`

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

Show example prompt

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

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

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

11. Verify that all keys are synced:

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

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

    On the partner controller:

12. Giveback the impaired controller:

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

13. Restore automatic giveback if you disabled it:

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

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

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

External Key Manager (EKM)

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

Before you begin

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

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

Steps

On the impaired controller:

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

   Show example boot menu

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

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

   Show example prompt

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

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

Show example

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

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

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

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

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

The recovery process completes and displays the following message:

```
Successfully recovered keymanager secrets.
```

Show example

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

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

Show example prompt

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

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

6. Restore automatic giveback if you disabled it:

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

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

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

What's next?

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

Return the failed part to NetApp - FAS50

If a component in your FAS50 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Hot-swap a caching module - FAS50

You can hot-swap an NVMe SSD caching module (Flash Cache module) of the same capacity from the same or different supported vendor for your FAS50 storage system.

Before you begin

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

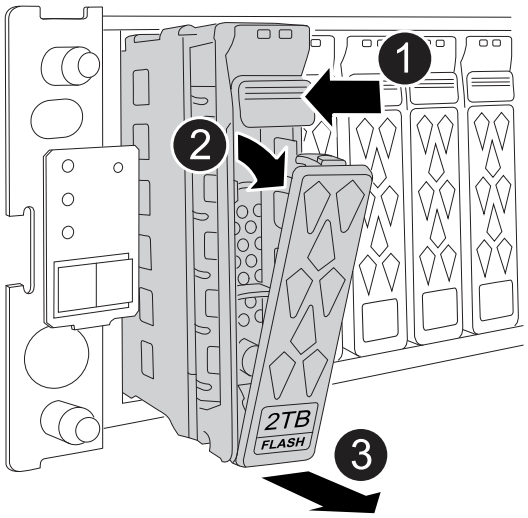
- Your storage system must have the appropriate operating system for the caching module you are installing.
- The replacement caching module must have the same capacity as the failed caching module, but can be from a different supported vendor.
- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing.

Steps

1. Properly ground yourself.
2. Remove the bezel from the front of the storage system.
3. Locate the failed caching module by the lit amber Attention LED on the front of the caching module.

   A caching module can be in drive bay 0 or 23.

4. Remove the caching module:



1	Press the release button on the module face to open the cam handle.
2	Rotate the cam handle downward to disengage the module from the midplane.
3	Slide the module out of the drive bay using the cam handle and supporting the module with your other hand. When removing a module, always use two hands to support its weight.

5. Wait a minimum of 70 seconds before inserting the replacement caching module.
6. Install the replacement caching module:
   a. With the cam handle in the open position, use both hands to insert the module.
   b. Gently push until the module stops.
   c. Close the cam handle so that the module 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 module.

7. Verify that the module's activity (green) LED is illuminated.
8. Reinstall the bezel on the front of the storage system.
9. 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

Chassis replacement workflow - FAS50

Follow these workflow steps to replace your FAS50 storage system chassis.

1. **Review the chassis replace requirements**

To replace the chassis, you must meet certain requirements.

2. **Shut down the controllers**

Shut down the controllers so you can perform maintenance on the chassis.

3. **Replace the chassis**

Replacing the chassis includes moving the drive blanks, any caching modules, controllers (with the power supplies), and bezel from the impaired chassis to the new chassis, and swapping out the impaired chassis with the new chassis of the same model as the impaired chassis.

4. **Complete chassis replacement**

Verify the HA state of the chassis and return the failed part to NetApp.

Chassis replacement requirements - FAS50

Before replacing the chassis of your FAS50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement chassis, and the necessary tools.

Review the following requirements and considerations.

Requirements

- The replacement chassis must be the same model as the impaired chassis. This procedure is for a like-for-like replacement, not for an upgrade.
- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

Considerations

- The chassis replacement procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.
- You can use the chassis replacement procedure with all versions of ONTAP supported by your storage system.
- The chassis replacement procedure is written with the assumption that you are moving the bezel, drives, any drive blanks, and controllers to the new chassis.

What's next?

After you've reviewed the requirements to replace the chassis, you need to shut down the controllers.

Shut down the controllers - FAS50

Shut down the controllers in your FAS50 storage system to prevent data loss and ensure system stability when replacing the chassis.

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

Before you begin

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

Steps

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

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

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

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

6. Exit the cluster shell:

```
exit
```

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

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

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

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

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

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

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

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

What's next?

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

Replace the chassis - FAS50

Replace the chassis of your FAS50 storage system when a hardware failure requires it. The replacement process involves removing the controllers, removing any caching modules and drive blanks, installing the replacement chassis, and reinstalling the chassis components.

About this task

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

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

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

Step 1: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

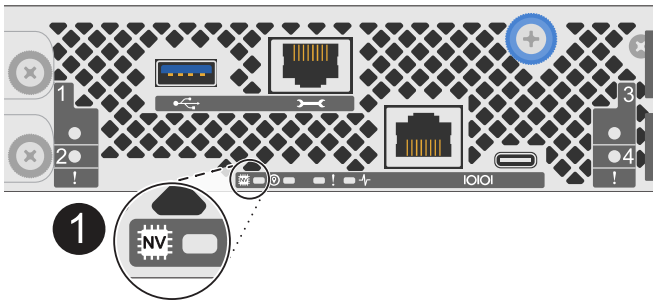
Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

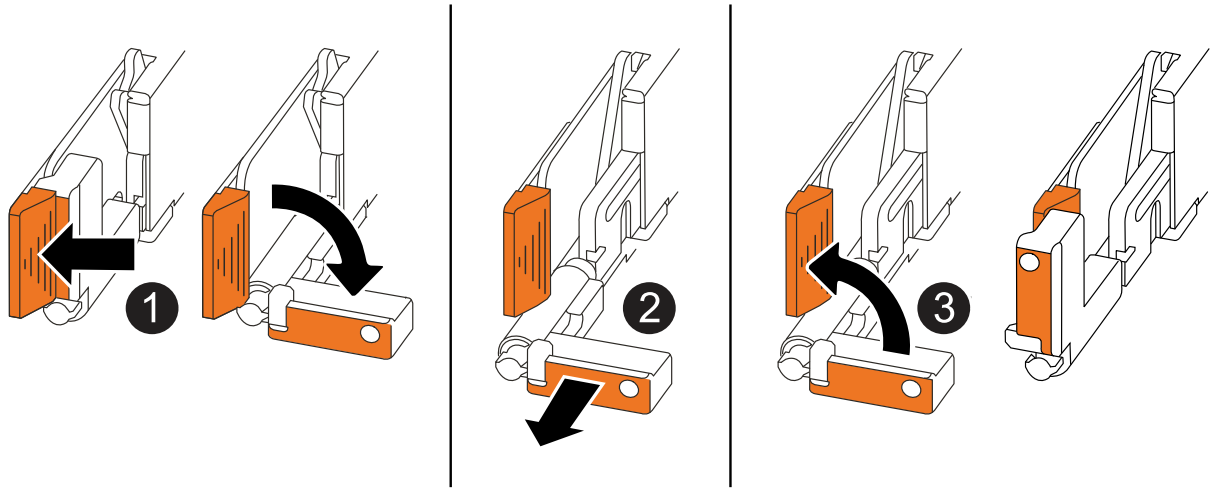
If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

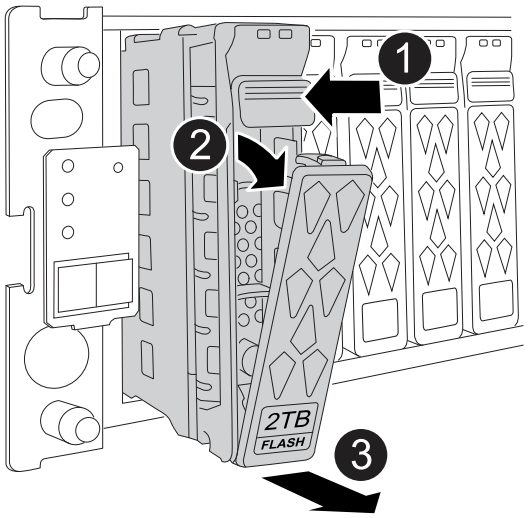
6. Repeat these steps for the other controller in the chassis.

Step 2: Remove the caching modules from the impaired chassis

You need to remove any caching modules and drive blanks from the impaired chassis so that later in the procedure you can install them in the replacement chassis.

1. Gently remove the bezel from the front of the storage system.
2. Remove the caching modules and drive blanks:

Keep track of what drive bay each caching module was removed from because they must be installed in the same drive bays in the replacement chassis.



1	Press the release button on the caching module face to open the cam handle.
2	Rotate the cam handle downward to disengage the caching module from the midplane.
3	Slide the caching module out of the drive bay using the cam handle and supporting the caching module with your other hand. When removing a caching module, always use two hands to support its weight.

3. Set the caching modules aside on a static-free cart or table.

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

You remove the impaired chassis from the equipment rack or system cabinet, install the replacement chassis, install the controllers, install any caching modules and drive blanks and then install the bezel.

1. Remove the screws from the impaired chassis mount points.

   Set the screws aside to use later in this procedure.

   If the storage system shipped in a NetApp system cabinet, you must remove additional screws at the rear of the chassis before the chassis can be removed.

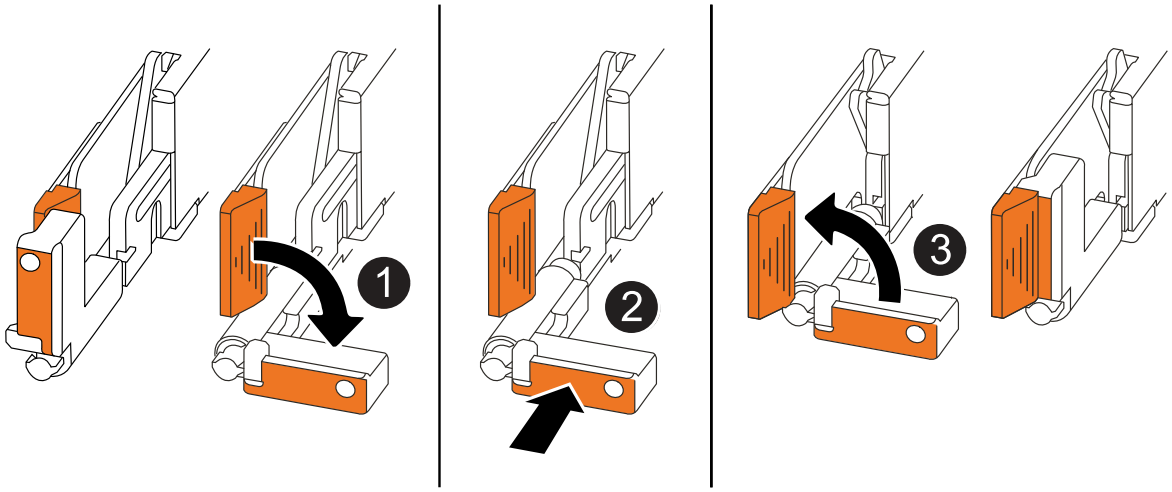
2. Using two people or a power lift, remove the impaired chassis from the equipment rack or system cabinet by sliding it off the rails, and then set it aside.
3. Using two people, install the replacement chassis into the equipment rack or system cabinet by sliding it onto the rails.
4. Secure the front of the replacement chassis to the equipment rack or system cabinet using the screws you removed from the impaired chassis.

Step 4: Install the controllers

Install the controllers into the replacement chassis and reboot them.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when installing a controller, and can be used as a reference for the rest of the controller installation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis and push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

1. Insert one of the controllers into the chassis:
   a. Align the back of the controller with the opening in the chassis.
   b. Firmly push on the handles until the controller meets the midplane and is fully seated in the chassis.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   c. Rotate the controller handles up and lock in place with the tabs.
2. Recable the controller, as needed, except for the power cords.
3. Repeat these steps to install the second controller into the chassis.
4. Install the caching modules and drive blanks you removed from the impaired chassis into the replacement chassis:

   The caching modules and drive blanks must be installed in the same drive bays in the replacement chassis.

   a. With the cam handle in the open position, use both hands to insert the caching module.
   b. Gently push until the caching module stops.
   c. Close the cam handle so that the caching module 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 caching module.

d. Repeat the process for the remaining caching module, if needed.

5. Install the bezel.
6. Reconnect the power cords to the power supplies (PSU) in the controllers.

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

   The controllers begin to boot as soon as the power is restored.

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

7. If controllers boot to the LOADER prompt, reboot the controllers:

   ```
   boot_ontap
   ```

8. Turn AutoSupport back on:

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

What's next?

After you've replaced the impaired FAS50 chassis and reinstalled the components into it, you need to complete the chassis replacement.

Complete chassis replacement - FAS50

Verify the HA state of the chassis and then return the failed part to NetApp to complete the final step in the FAS50 chassis replacement procedure.

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 storage system configuration.

1. In Maintenance mode, from either controller, display the HA state of the local controller 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 storage system configuration:
   a. Set the HA state for the chassis:

```
ha-config modify chassis HA-state
```

The value for HA-state should be *ha*.
The value for HA-state can be one of the following:
* ***ha***
* *mcc* (not supported in ASA)

b. Confirm that the setting has changed:

```
ha-config show
```

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

Step 2: Return the failed part to NetApp

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

Controller

Controller replacement workflow - FAS50

Follow these workflow steps to replace your controller in your FAS50 storage system.



Review the controller replacement requirements

To replace the controller, you must meet certain requirements.



Shut down the impaired controller

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



Replace the controller

Replacing the controller includes removing the impaired controller, moving FRU components to the replacement controller, installing the replacement controller in the chassis, setting the time and date, and then recabling.



Restore and verify the system configuration

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



Give back the controller

Transfer the ownership of storage resources back to the replacement controller.

6 Complete controller replacement

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

Requirements to replace the controller - FAS50

Before replacing the controller in your FAS50 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements and considerations for the controller replacement procedure.

Requirements

- All shelves must be working properly.
- 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.
- You must replace a controller with a controller of the same model type. You cannot upgrade your system by just replacing the controller.
- You cannot change any drives or shelves as part of this procedure.
- You must always capture the controller's console output to a text log file.

  The console output provides you with a record of the procedure you can use to troubleshoot issues you might encounter during the replacement process.

Considerations

It is important that you apply the commands in this procedure to the correct controller:

- 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.

What's next?

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

Shut down the impaired controller - FAS50

Shut down the impaired controller in your FAS50 storage system to prevent data loss and ensure system stability when replacing the controller.

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

What's next?

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

Replace the controller - FAS50

Replace the controller in your FAS50 storage system when a hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

About this task

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

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

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

Step 1: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

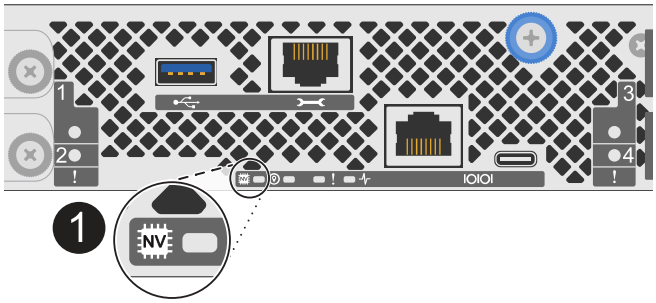
Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

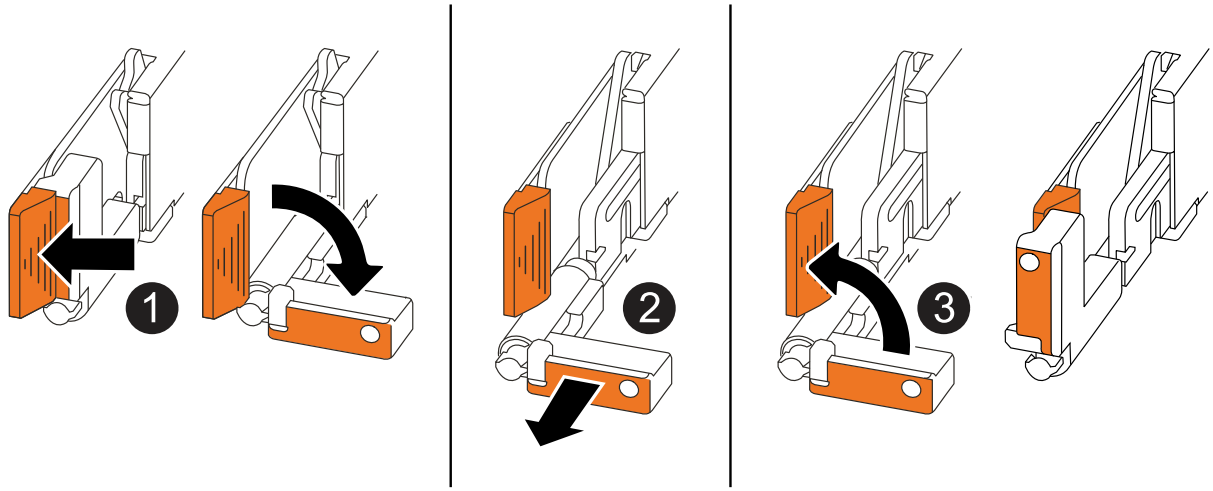
If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 2: Move the power supply

Move the power supply (PSU) to the replacement controller.

1. Move the PSU from the impaired controller:

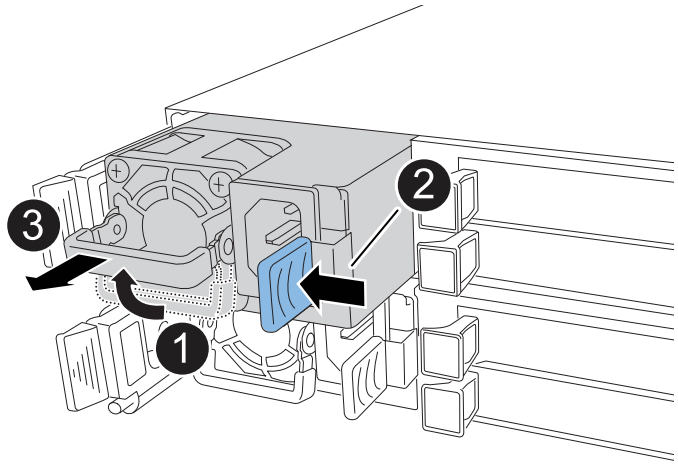
   Make sure the left side controller handle is in the upright position to allow you access to the PSU.

Option 1: Move an AC PSU

To move an AC PSU, complete the following steps.

Steps

1. Remove the AC PSU from the impaired controller:



1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight. The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

2. Insert the PSU into the replacement controller:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently push the PSU into the controller until the locking tab clicks into place.

      A PSU will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

   c. Rotate the handle down, so it is out of the way of normal operations.

Option 2: Move a DC PSU

To move a DC PSU, complete the following steps.

Steps

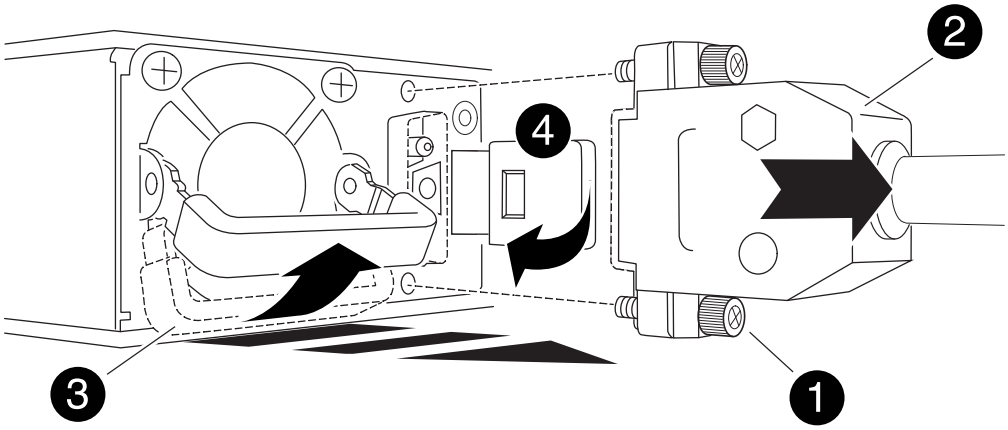
1. Remove the DC PSU from the impaired controller:

a. Rotate the handle up, to its horizontal position, and then grasp it.

b. With your thumb, press the terracotta tab to release the locking mechanism.

c. Pull the PSU out of the controller while using your other hand to support its weight.

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



1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

2. Insert the PSU into the replacement controller:

a. Using both hands, support and align the edges of the PSU with the opening in the controller.

b. Gently slide the PSU into the controller until the locking tab clicks into place.

A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.

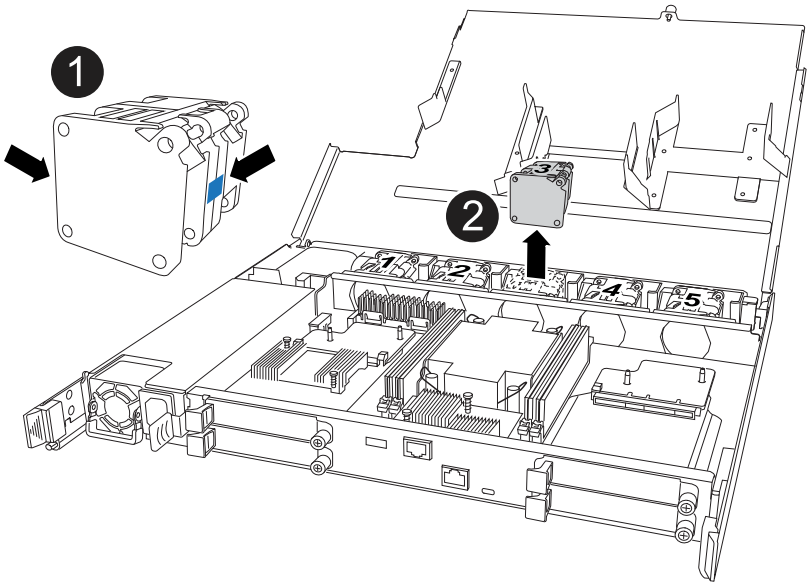
> To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

c. Rotate the handle down, so it is out of the way of normal operations.

Step 3: Move the fans

Move the fans to the replacement controller.

1. Remove one of the fans from the impaired controller:



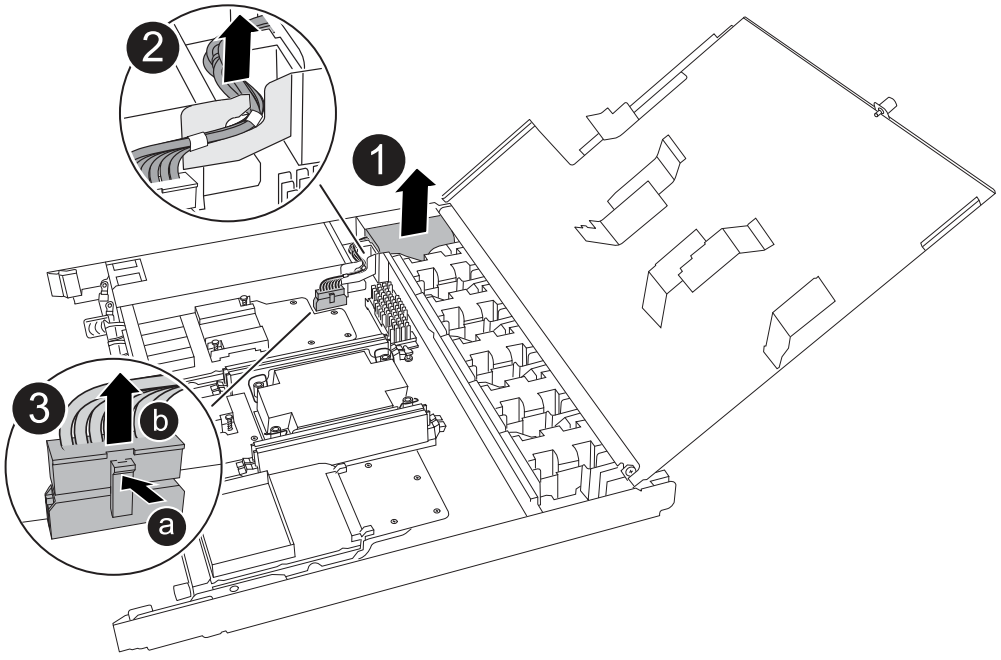
1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

2. Insert the fan into the replacement controller by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.
3. Repeat these steps for the remaining fans.

Step 4: Move the NV battery

Move the NV battery to the replacement controller.

1. Remove the NV battery from the impaired controller:



1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	1. Push in and hold the tab on the connector. 2. Pull the connector up and out of the socket. As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.

2. Install the NV battery into the replacement controller:
   a. Plug the wiring connector into its socket.
   b. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
   c. Place the NV battery into the compartment.

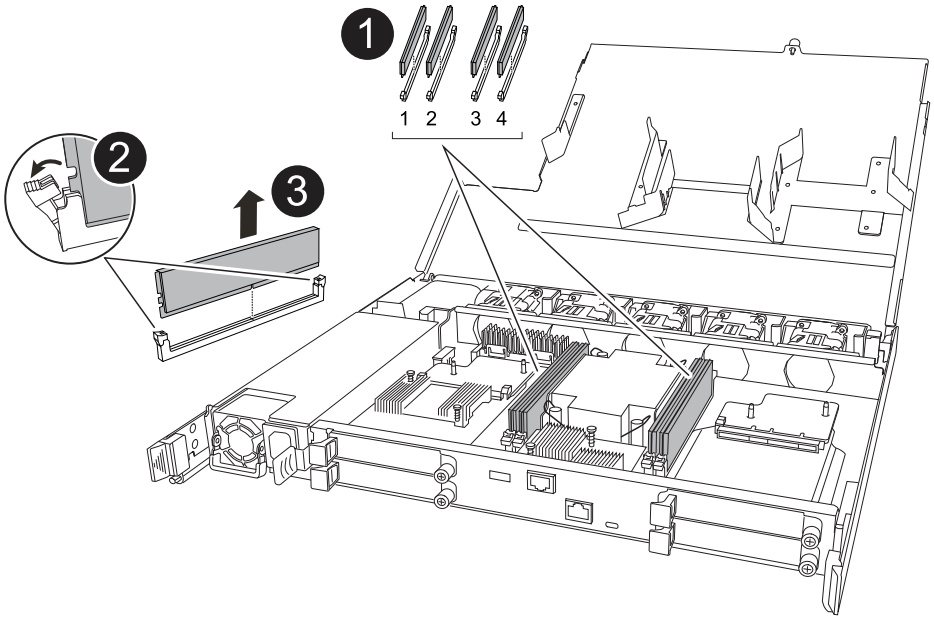
      The NV battery should sit flush in its compartment.

Step 5: Move system DIMMs

Move the DIMMs to the replacement controller.

If you have DIMM blanks, you do not need to move them, the replacement controller should come with them installed.

1. Remove one of the DIMMs from the impaired controller:



1	DIMM slot numbering and positions. Depending on your storage system model, you will have two or four DIMMs.
2	• Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller in the proper orientation. • Eject the DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot. Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

2. Install the DIMM in the replacement controller:
   a. Make sure that the DIMM ejector tabs on the connector are in the open position.
   b. Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

      The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

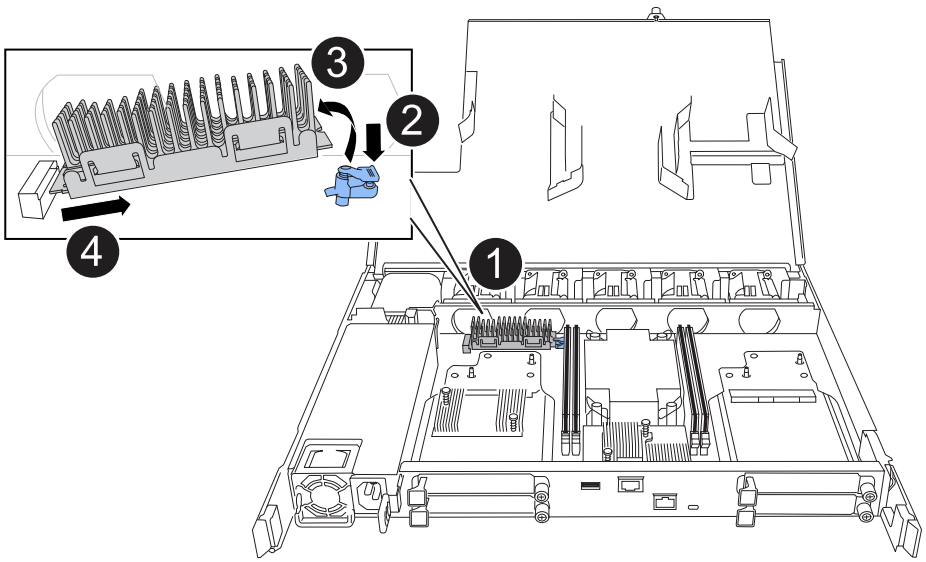
      When inserted correctly, the DIMM goes in easily but fits tightly in the slot. If not, reinsert the DIMM.

   c. Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
   d. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.
3. Repeat these steps for the remaining DIMMs.

Step 6: Move the boot media

Move the boot media to the replacement controller.

1. Remove the boot media from the impaired controller:



1	Boot media location
2	Press down on the blue tab to release the right end of the boot media.
3	Lift the right end of the boot media up at a slight angle to get a good grip along the sides of the boot media.
4	Gently pull the left end of the boot media out of its socket.

2. Install the boot media into the replacement controller:
   a. Slide the socket end of the boot media into its socket.
   b. At the opposite end of the boot media, press down and hold the blue tab (in the open position), gently push down on that end of the boot media until it stops, and then release the tab to lock the boot media into place.

Step 7: Move the I/O modules

Move the I/O modules and any I/O blanking modules to the replacement controller.

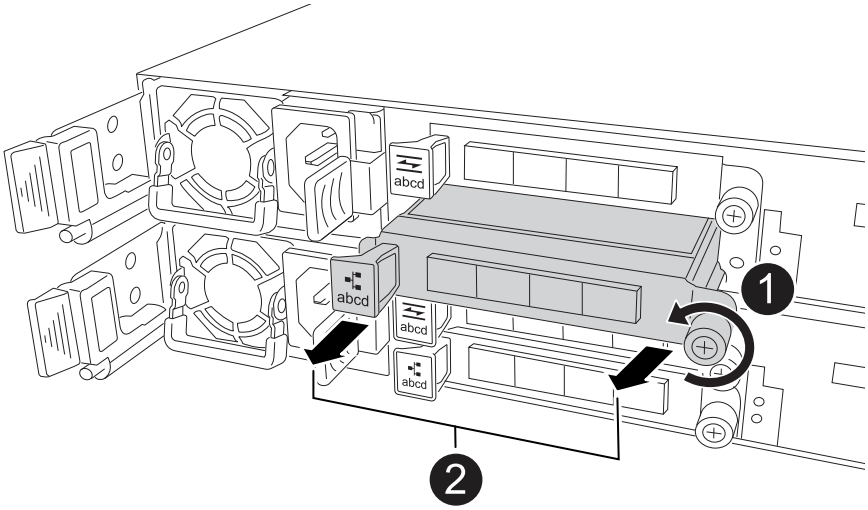
1. Unplug cabling from one of the I/O modules.

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

2. Remove the I/O module from the impaired controller:

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

   If you are removing the I/O module in slot 4, make sure the right side controller handle is in the upright position to allow you access to the I/O module.



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

3. Install the I/O module into the replacement controller:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

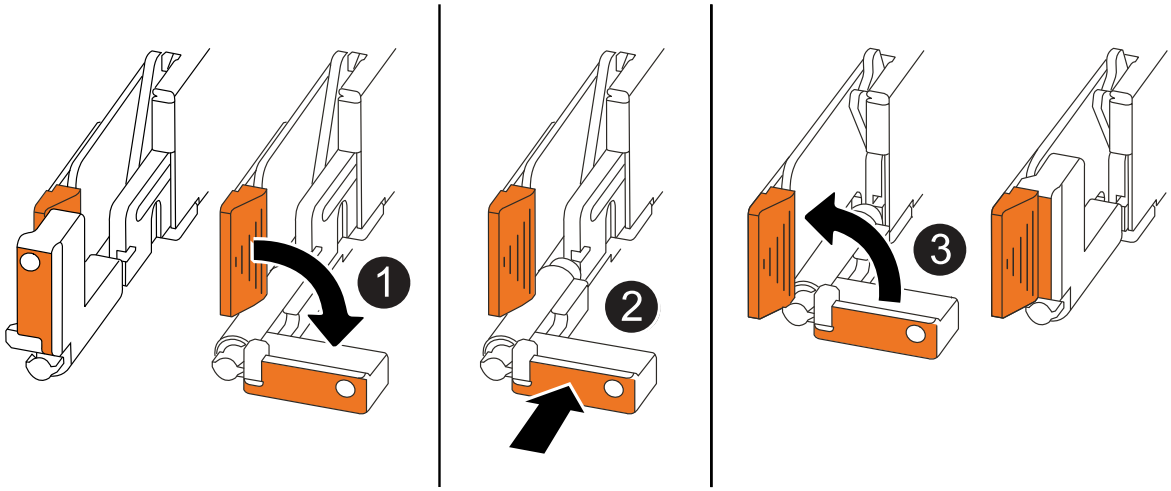
   c. Turn the thumbscrew clockwise to tighten.
4. Repeat these steps to move the remaining I/O modules and any I/O blanking modules to the replacement controller.

Step 8: Install the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

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

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Take the controller to the LOADER prompt by pressing CTRL-C to abort AUTOBOOT.
6. Set the time and date on the controller:

   Make sure you are at the controller's LOADER prompt.

   a. Display the date and time on the controller:

      ```
      show date
      ```

      Time and date default is in GMT. You have the option to display in local time and in 24hr mode.

   b. Set the current time in GMT:

      ```
      set time hh:mm:ss
      ```

      You can get the current GMT from the healthy node:

      ```
      date -u
      ```

   c. Set the current date in GMT:

      ```
      set date mm/dd/yyyy
      ```

      You can get the current GMT from the healthy node:
      ```
      date -u
      ```
7. Recable the controller as needed.
8. Reconnect the power cord to the power supply (PSU).

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

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

What's next?

After you've replaced the impaired FAS50 controller, you need to restore the system configuration.

Restore and verify the system configuration - FAS50

Verify that the controller's HA configuration is active and functioning correctly in your FAS50 storage system, and confirm that the system's adapters list all the paths to the disks.

Step 1: Verify HA config settings

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

1. Boot to maintenance mode:

   ```
   boot_ontap maint
   ```

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

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

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

   If you see *PERSONALITY MISMATCH* contact customer support.

3. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.
4. Verify that all components display the same `HA` state:

   ```
   ha-config show
   ```

   The HA state should be the same for all components.

5. If the displayed system state of the controller does not match your storage system configuration, set the `HA` state for the controller:

   ```
   ha-config modify controller ha
   ```

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

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

6. Confirm that the setting has changed:

   ```
   ha-config show
   ```

Step 2: Verify disk list

1. Verify that the adapter lists the paths to all disks:

   ```
   storage show disk -p
   ```

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

2. Exit Maintenance mode:

```
halt
```

What's next?

After you've restored and verified the system configuration for your FAS50 system, you need to give back the controller.

Give back the controller - FAS50

Return control of storage resources to the replacement controller so your FAS50 system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption, Onboard Key Manager (OKM) encryption, or External Key Manager (EKM) encryption.

No encryption

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

Steps

1. From the LOADER prompt, enter `boot_ontap`.
2. Press <enter> when console messages stop.
   - If you see the *login* prompt, go to the next step at the end of this section.
   - If you see *Waiting for giveback*, press the <enter> key, log into the partner node, and then go to the next step at the end of this section.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

1. From the LOADER prompt, enter `boot_ontap maint`.
2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap menu` and select option 10.
3. Enter the OKM passphrase.

   You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option `1` for normal boot.
6. Press <enter> when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as `admin`.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`

   If you encounter errors, contact NetApp Support.

9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
    a. Move the console cable back to the replacement controller.
    b. Synchronize missing keys: `security key-manager onboard sync`

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

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

   The output should show no results when when properly synchronized.

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

11. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
12. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
13. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

External key manager (EKM)

Reset encryption and return the controller to normal operation.

Steps

1. If the root volume is encrypted with External Key Manager and the console cable is connected to the replacement node, enter `boot_ontap menu` and select option `11`.
2. If these questions appear, answer `y` or `n` as appropriate:

   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}

   Do you know the KMIP server address? {y/n}

   Do you know the KMIP port? {y/n}

   Contact NetApp Support if you have issues.

3. Supply the information for:
   - The client certificate (client.crt) file contents
   - The client key (client.key) file contents
   - The KMIP server CA(s) (CA.pem) file contents
   - The IP address for the KMIP server
   - The port for the KMIP server
4. Once the system processes, you see the Boot Menu. Select '1' for normal boot.
5. Check the takeover status: `storage failover show`
6. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
7. If automatic giveback was disabled, reenable it: `storage failover modify -node local`

```
-auto-giveback true
```

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

What's next?

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

Complete controller replacement - FAS50

To complete the controller replacement for your FAS50 system, first restore the NetApp Storage Encryption configuration (if necessary) and install the required licenses on the new controller. Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, register the new controller's serial number and then return the failed part to NetApp.

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.

Before you begin

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 ONTAP platforms. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

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.

- 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 the node is in a MetroCluster configuration and all nodes at a site have been replaced, license keys must be installed on the *replacement* node or nodes prior to switchback.

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, register the serial number, and check cluster health

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

Steps

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

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

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

Step 3: Return the failed part to NetApp

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

Replace a DIMM - FAS50

Replace a DIMM in your FAS50 storage system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

Before you begin

- All other components in the storage system must be working correctly; if not, contact NetApp Support before continuing.
- You must replace the failed FRU component with a replacement FRU component you received from your provider.

About this task

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

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

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

Step 2: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

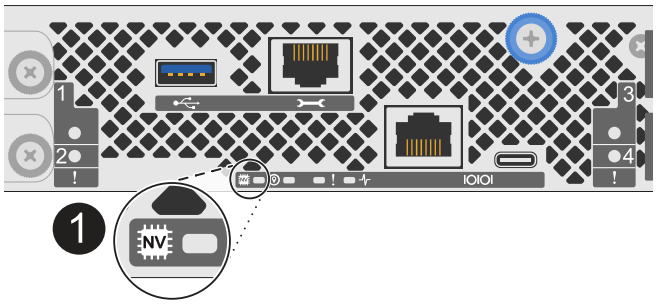
Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

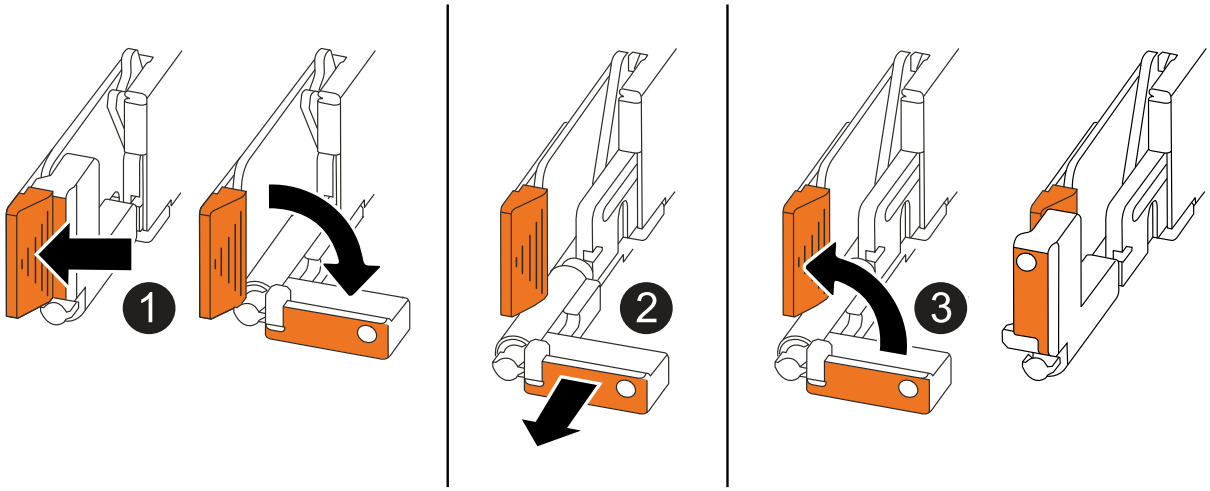
If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace a DIMM

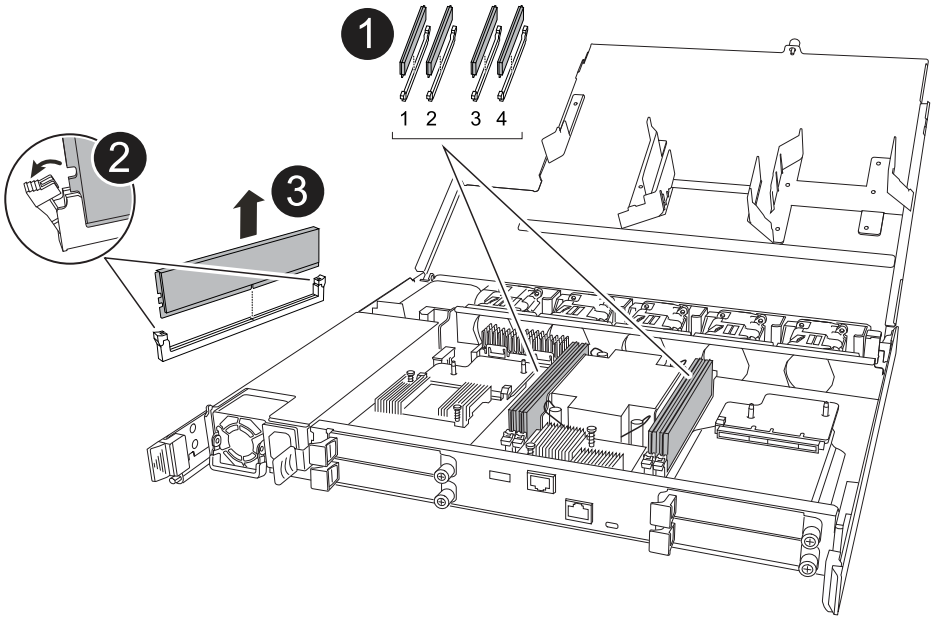
To replace a DIMM, locate the faulty DIMM inside the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller and identify the faulty DIMM.

> Consult either the Netapp Hardware Universe or the FRU map on the cover of the controller for exact DIMM locations.

3. Remove the faulty DIMM:



1	DIMM slot numbering and positions. Depending on your storage system model you will have two or four DIMMs.
2	• Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM using the same orientation. • Eject the faulty DIMM by slowly pushing apart the two DIMM ejector tabs on both ends of the DIMM slot. Carefully hold the DIMM by the corners or edges to avoid pressure on the DIMM circuit board components.
3	Lift the DIMM up and out of the slot. The ejector tabs remain in the open position.

4. Install the replacement DIMM:
   a. Remove the replacement DIMM from its antistatic shipping bag.
   b. Make sure that the DIMM ejector tabs on the connector are in the open position.
   c. Hold the DIMM by the corners, and then insert the DIMM squarely into the slot.

      The notch on the bottom of the DIMM, among the pins, should line up with the tab in the slot.

      When inserted correctly, the DIMM goes in easily but fits tightly in the slot. Reinsert the DIMM if you feel it is not inserted correctly.

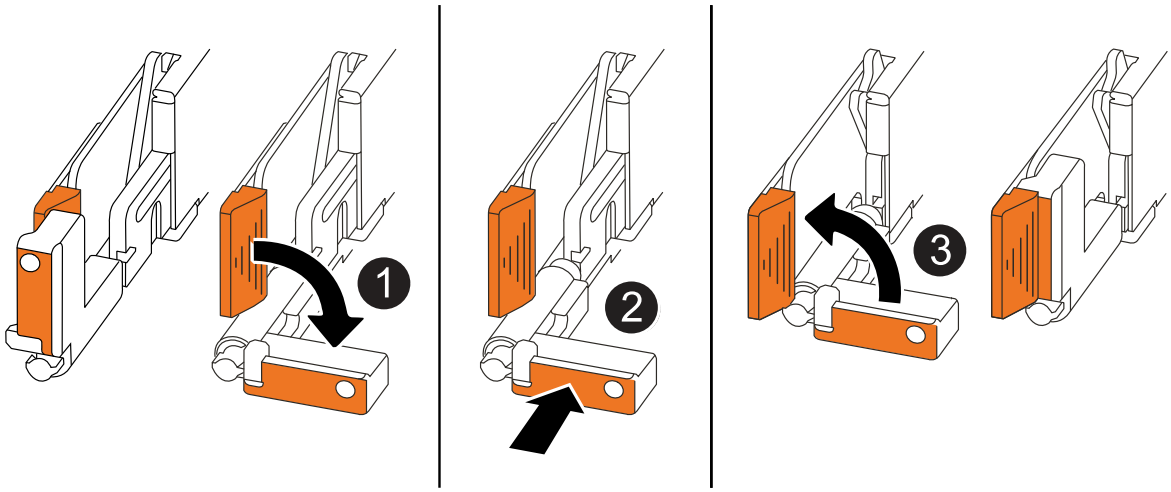
   d. Visually check the DIMM to make sure it is evenly aligned and fully inserted into the slot.
   e. Push down carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at both ends of the DIMM.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

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

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

> Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

> Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

> The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

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

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

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

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

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

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

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

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

Step 5: Return the failed part to NetApp

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

Replace a drive - FAS50

Replace a drive in your FAS50 storage system when a drive fails or requires an upgrade. The replacement process involves identifying the faulty drive, safely removing it, and installing a new drive to ensure continued data access and system performance.

You can replace a failed SSD drive nondisruptively while I/O is in progress.

Before you begin

- The drive that you are installing must be supported by your storage system.

  NetApp Hardware Universe

- If self-encrypting drive (SED) authentication is enabled, you must use the SED replacement instructions in the ONTAP documentation.

  Instructions in the ONTAP documentation describe additional steps you must perform before and after replacing an SED.

  NetApp encryption overview with the CLI

- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.
- Verify that the drive you are removing is failed.

  You can verify that the drive is failed by running the `storage disk show -broken` command. 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 drive type and capacity, it can take up to several hours for the drive to appear in the list of failed drives.

About this task

- When replacing a failed drive, you must wait 70 seconds between the removal of the drive and the insertion of the replacement drive to allow the storage system to recognize that a drive was removed.
- The best practice is to have the current version of the Disk Qualification Package (DQP) installed before hot-swapping a drive.

  Having the current version of the DQP installed allows your system to recognize and use newly qualified drives. This avoids system event messages about having noncurrent drive information and prevention of drive partitioning because drives are not recognized. The DQP also notifies you of noncurrent drive firmware.

  NetApp Downloads: Disk Qualification Package

- The best practice is to have current versions of NVMe shelf module (NSM) firmware and drive firmware on your system before replacing FRU components.

NetApp Downloads: Disk Shelf Firmware

NetApp Downloads: Disk Drive Firmware

> Do not revert firmware to a version that does not support your shelf and its components.

- Drive firmware is automatically updated (nondisruptively) on new drives that have non-current firmware versions.

  > Drive firmware checks occur every two minutes.

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

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

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

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.

   > 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.

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

   b. If automatic drive assignment is enabled, disable it:

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

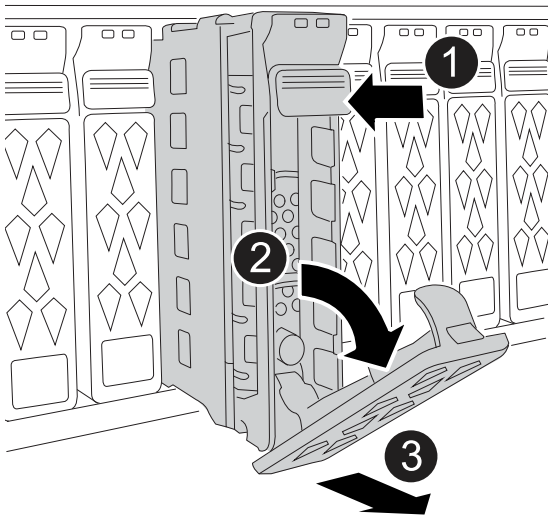
      You must disable automatic drive assignment on both controllers.

2. Properly ground yourself.
3. Remove the bezel from the front of the storage system.
4. 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.

5. Remove the failed drive:



1	Press the release button on the drive face to open the cam handle.
2	Rotate the cam handle downward to disengage the drive from the midplane.
3	Slide the drive out of the drive bay using the cam handle and supporting the drive with your other hand. When removing a drive, always use two hands to support its weight. Because drives are fragile, minimize handling to avoid damaging them.

6. Wait a minimum of 70 seconds before inserting the replacement drive.
7. Insert the replacement drive:
   a. With the cam handle in the open position, use both hands to insert the drive.
   b. Gently 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.
8. 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.

9. If you are replacing another drive, repeat the preceding steps.
10. Reinstall the bezel on the front of the storage system.
11. If you disabled automatic drive assignment earlier in this procedure, 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.

    b. Assign each drive:

       `storage disk assign -disk disk_name -owner owner_name`

       You can enter the command on either controller.

       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 controllers.
12. 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.//2025-11-17 ontap-systems-internal/issues/1391

Replace a fan module - FAS50

Replace a fan module in your FAS50 system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

About this task

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

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

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

Step 2: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

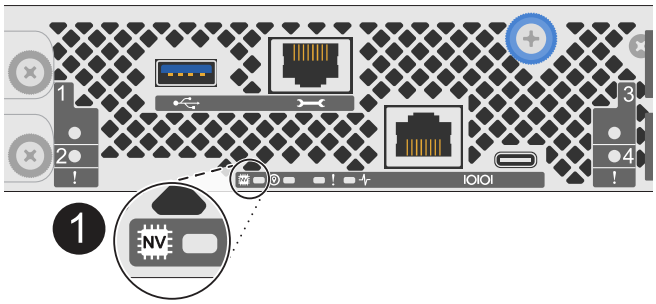
Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

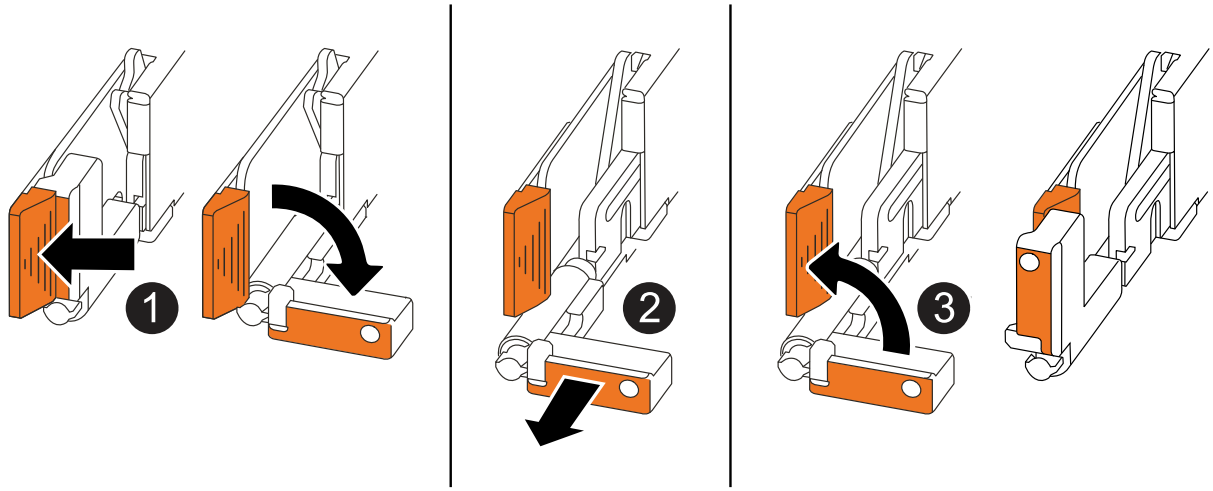
If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

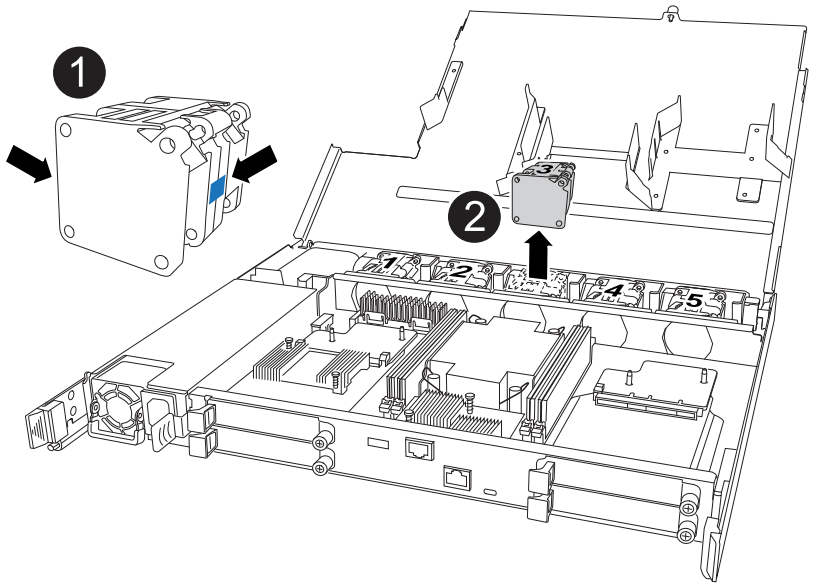
6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace fan

To replace a fan, remove the failed fan and replace it with a new fan.

Steps

1. Identify the fan that you must replace by checking the console error messages.
2. Remove the failed fan:



1	Hold both sides of the fan at the blue touch points.
2	Pull the fan straight up and out its socket.

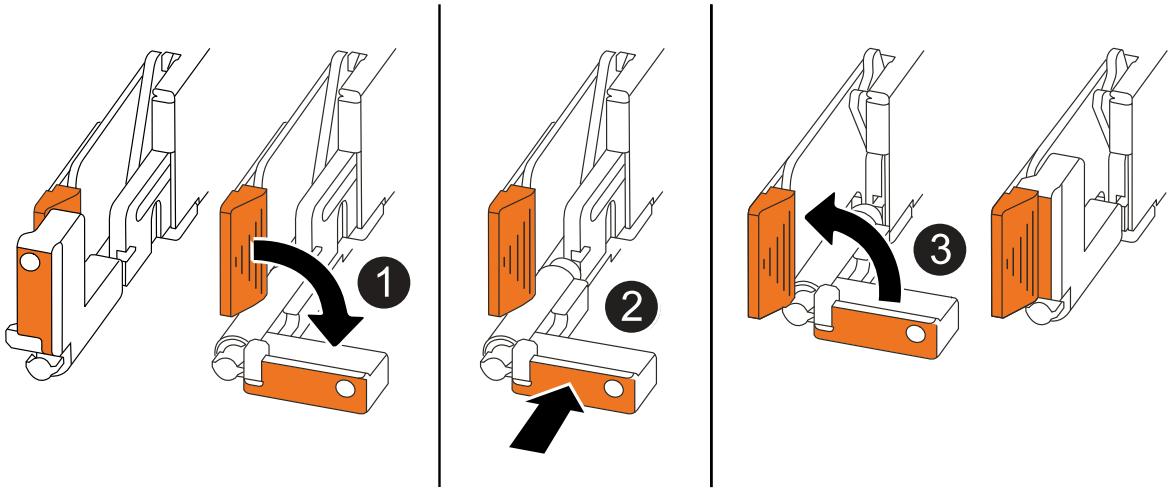
3. Insert the replacement fan by aligning it within the guides, and then push down until the fan connector is fully seated in the socket.

Step 4: Reinstall the controller module

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

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

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

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

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

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

   `storage failover giveback -ofnode impaired_node_name`

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

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

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

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

Step 5: Return the failed part to NetApp

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

I/O module

Overview of I/O module maintenance - FAS50

The FAS50 storage systems offer flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding, hot-swapping, or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your storage system with the same type of I/O module, or with a different type of I/O module. You can hot-swap a cluster and HA I/O module when your storage system meets specific requirements. You can also add an I/O module to a storage system with available slots.

- Add an I/O module

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

- Hot-swap a cluster and HA I/O module

Hot-swapping a failed cluster and HA I/O module can restore the storage system to its optimal operating state. Hot-swapping is done without having to manually take over the impaired controller.

To use this procedure, your storage system must be running ONTAP 9.17.1 or later and meet specific system requirements.

- Replace an I/O module

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

Add an I/O module - FAS50

Add an I/O module to your FAS50 storage system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your FAS50 storage system when there are slots available or when all slots are fully populated (by removing an existing I/O module and installing a new one in its place).

About this task

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

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

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

Step 1: Shut down the impaired controller module

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: MetroCluster configuration

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

Step 2: Add the new I/O module

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

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- All other components in the storage system must be functioning properly; if not, contact NetApp Support before continuing with this procedure.

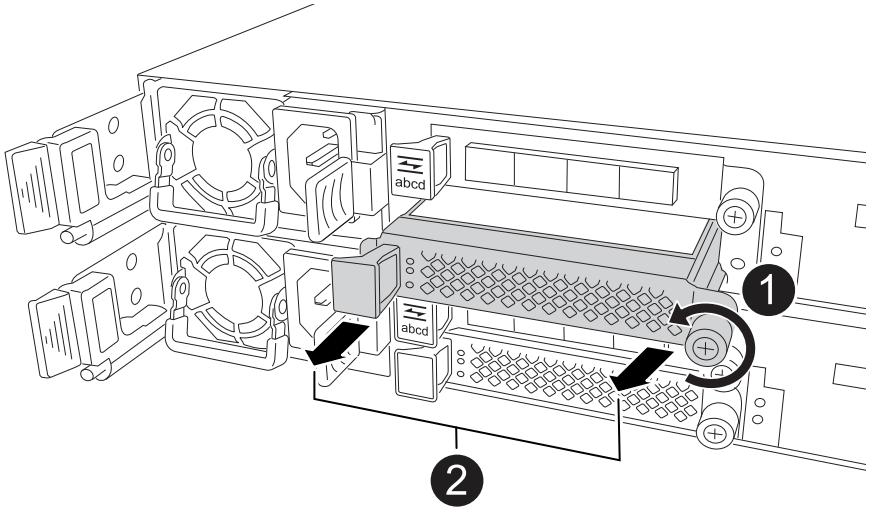
Add I/O module to an available slot

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

Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, remove the I/O blanking module from the target slot.

   Unused I/O slots should have blanking module installed to prevent possible thermal issues and for EMC compliance.



1	On the I/O blanking module, turn the thumbscrew counterclockwise to loosen.
2	Pull the I/O blanking module out of the controller using the tab on the left and the thumbscrew.

3. Install the new I/O module:
   a. Align the I/O module with the edges of the controller slot opening.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O module.

   c. Turn the thumbscrew clockwise to tighten.
4. Cable the I/O module to the designated devices.

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

5. Reboot the impaired controller from the LOADER prompt: `bye`

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

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

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

7. Repeat these steps to add an I/O module to the other controller.
8. Restore automatic giveback from the console of the healthy controller:

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

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation: +

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

Add I/O module to a fully-populated system

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

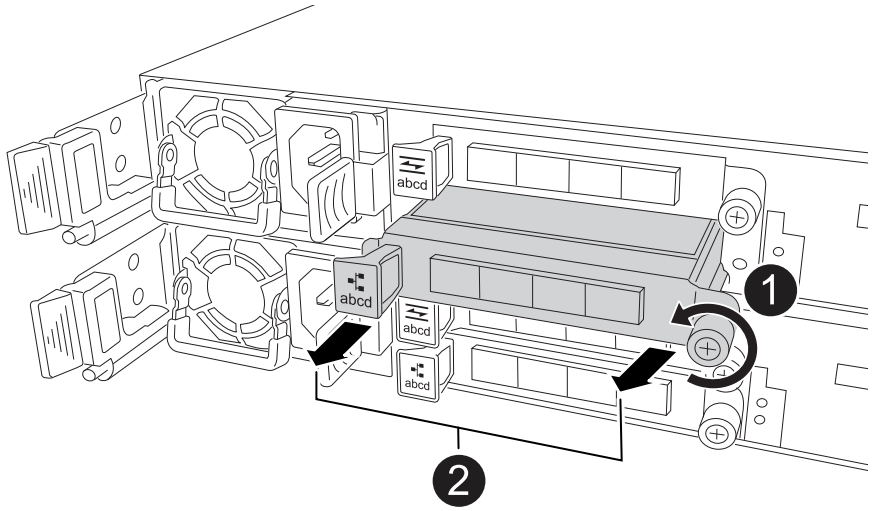
About this task

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

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

Steps

1. If you are not already grounded, properly ground yourself.
2. On the impaired controller, unplug any cabling on the target I/O module.
3. Remove the target I/O module from the controller:



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the new I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

      You can use the tab on the left and the thumbscrew to push in the I/O Module.

   c. Turn the thumbscrew clockwise to tighten.
5. Cable the I/O module to the designated devices.

   If you installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.
6. Repeat the I/O module remove and install steps to add any additional I/O modules in the controller.
7. Reboot the impaired controller from the LOADER prompt:

   ```
   bye
   ```

   Rebooting the impaired controller also reinitializes the I/O modules and other components.
8. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```
9. Restore automatic giveback from the console of the healthy controller:

   ```
   storage failover modify -node local -auto-giveback true
   ```
10. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```
11. If you installed a NIC module, specify the usage mode for each port as *network*:

    ```
    storage port modify -node node_name -port port_name -mode network
    ```
12. Repeat these steps for the other controller.

Hot-swap the I/O module used for cluster and HA traffic - FAS50

The cluster and HA I/O module supports interconnects for clustering and high-availability. You can hot-swap the module in your FAS50 storage system when the module fails and if your storage system meets specific requirements.

To hot-swap a module, you ensure your storage system meets the procedure requirements, prepare the storage system and I/O module in slot 4, hot-swap the failed module for an equivalent one, bring the replacement module online, restore the storage system to normal operation, and return the failed module to NetApp.

About this task

- Hot-swapping the cluster and HA I/O module means that you do not have to perform a manual takeover; the impaired controller (the controller with the failed cluster and HA I/O module) has automatically taken over the healthy controller.

  When the impaired controller has taken over the healthy controller, the only way to recover without an outage is to hot-swap the module.

- It is critical to apply the commands to the correct controller when you are hot-swapping the cluster and HA I/O module:
  - The *impaired controller* is the controller on which you are hot-swapping the cluster and HA I/O module and it is the controller that has taken over the healthy controller.
  - The *healthy controller* is the HA partner of the impaired controller and it is the controller that was taken over by the impaired controller.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

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

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

Step 1: Ensure the storage system meets the procedure requirements

To use this procedure, make sure your storage system meets all requirements.

> If your storage system does not meet all requirements, you must use the replace an I/O module procedure.

- Your storage system must be running ONTAP 9.17.1 or later.
- The I/O module that failed must be a cluster and HA I/O module in slot 4 and you must be replacing it with an equivalent cluster and HA I/O module. You cannot change the I/O module type.
- Your storage system configuration must have only one cluster and HA I/O module located in slot 4, not two cluster and HA I/O modules.
- Your storage system must be a two-node (switchless or switched) cluster configuration.
- The controller with the failed cluster and HA I/O module (the impaired controller) must have already taken over the healthy partner controller. The takeover should have occurred automatically if the I/O module is failed.

  For two-node clusters, the storage system cannot discern which controller has the failed I/O module, so either controller might initiate the takeover. The cluster and HA I/O module hot-swap procedure is only supported when the controller with the failed I/O module (the impaired controller) has taken over the healthy controller.

You can verify that the impaired controller successfully took over the healthy controller by entering the `storage failover show` command.

If you are not sure which controller has the failed I/O module, contact NetApp Support.

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

Step 2: Prepare the storage system and I/O module slot 4

Prepare the storage system and I/O module slot 4 so that it is safe to remove the failed cluster and HA I/O module:

Steps

1. Properly ground yourself.
2. Unplug cabling from the failed cluster and HA I/O module.

   Make sure to label the cables so that later in this procedure you can reconnect them to the same ports.

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

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

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

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

4. Disable automatic giveback:
   a. Enter the following command from the console of the impaired controller:

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

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
5. Prepare the failed cluster and HA module in slot 4 for removal by removing it from service and powering it off:
   a. Enter the following command:

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

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

      For example, the following command prepares the module in slot 4 on node 2 (the impaired controller) for removal, and displays a message that it is safe to remove:

```
node2::> system controller slot module remove -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered off for removal.

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

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

6. Verify the failed cluster and HA module in slot 4 is powered off:

   `system controller slot module show`

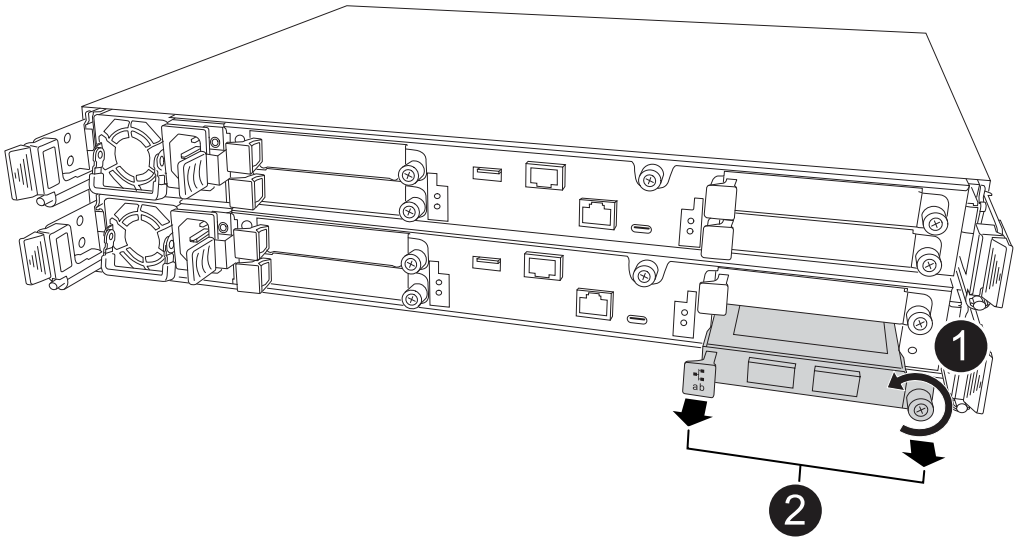
   The output should show *`powered-off`* in the status column for the failed module in slot 4.

Step 3: Replace the failed cluster and HA I/O module

Replace the failed cluster and HA I/O module in slot 4 with an equivalent I/O module:

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the failed cluster and HA I/O module from the impaired controller:



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew on the right.

3. Install the replacement cluster and HA I/O module into slot 4:

a. Align the I/O module with the edges of the slot.
b. Gently push the I/O module all the way into the slot, making sure to properly seat the I/O module into the connector.

   You can use the tab on the left and the thumbscrew on the right to push in the I/O module.

c. Turn the thumbscrew clockwise to tighten.

4. Cable the cluster and HA I/O module.

Step 4: Bring the replacement cluster and HA I/O module online

Bring the replacement cluster and HA I/O module in slot 4 online, verify the module ports initialized successfully, verify slot 4 is powered on, and then verify the module is online and recognized.

Steps

1. Bring the replacement cluster and HA I/O module online:
   a. Enter the following command:

      `system controller slot module insert -node` *`impaired_node_name`* `-slot` *`slot_name`*

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

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

      For example, the following command brings slot 4 on node 2 (the impaired controller) online, and displays a message that the process was successful:

```
node2::> system controller slot module insert -node node2 -slot 4

Warning: IO_2X_100GBE_NVDA_NIC module in slot 4 of node node2 will be
powered on and initialized.

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

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

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

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

   > It might take several minutes to allow for any required firmware updates and port initialization.

   The output should show a hotplug.init.success EMS event logged for each port on the cluster and HA I/O module with *`hotplug.init.success:`* in the *`Event`* column.

For example, the following output shows initialization succeeded for cluster and HA I/O module ports e4b and e4a:

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

Time                Node             Severity      Event

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

7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e4b" in slot 4 succeeded

7/11/2025 16:04:06  node2      NOTICE        hotplug.init.success:
Initialization of ports "e4a" in slot 4 succeeded

2 entries were displayed.
```

3. Verify I/O module slot 4 is powered on and ready for operation:

   `system controller slot module show`

   The output should show slot 4 status as *`powered-on`* and therefore ready for operation of the replacement cluster and HA I/O module.

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

   Enter the command from the console of the impaired controller:

   `system controller config show -node local -slot4`

   If the replacement cluster and HA I/O module was successfully brought online and is recognize, the output shows I/O module information, including port information, for slot 4.

   For example, you should see output similar to the following:

```
node2::> system controller config show -node local -slot 4

Node: node2
Sub- Device/
Slot slot Information
---- ---- ------------------------------
   4    - Dual 40G/100G Ethernet Controller CX6-DX
                  e4a MAC Address: d0:39:ea:59:69:74 (auto-100g_cr4-fd-
up)
                          QSFP Vendor:        CISCO-BIZLINK
                          QSFP Part Number:   L45593-D218-D10
                          QSFP Serial Number: LCC2807GJFM-B
                  e4b MAC Address: d0:39:ea:59:69:75 (auto-100g_cr4-fd-
up)
                          QSFP Vendor:        CISCO-BIZLINK
                          QSFP Part Number:   L45593-D218-D10
                          QSFP Serial Number: LCC2809G26F-A
                  Device Type:        CX6-DX PSID(NAP0000000027)
                  Firmware Version:   22.44.1700
                  Part Number:        111-05341
                  Hardware Revision:  20
                  Serial Number:      032403001370
```

Step 5: Restore the storage system to normal operation

Restore your storage system to normal operation by giving back storage to the healthy controller, restoring automatic giveback, and reenabling AutoSupport automatic case creation.

Steps

1. Return the healthy controller (the controller that was taken over) to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode healthy_node_name
   ```

2. Restore automatic giveback from the console of the impaired controller (the controller that took over the healthy controller):

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

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

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

Step 6: Return the failed part to NetApp

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

Replace an I/O module in your FAS50 storage system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

Before you begin

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

About this task

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

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

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

Step 2: Replace a failed I/O module

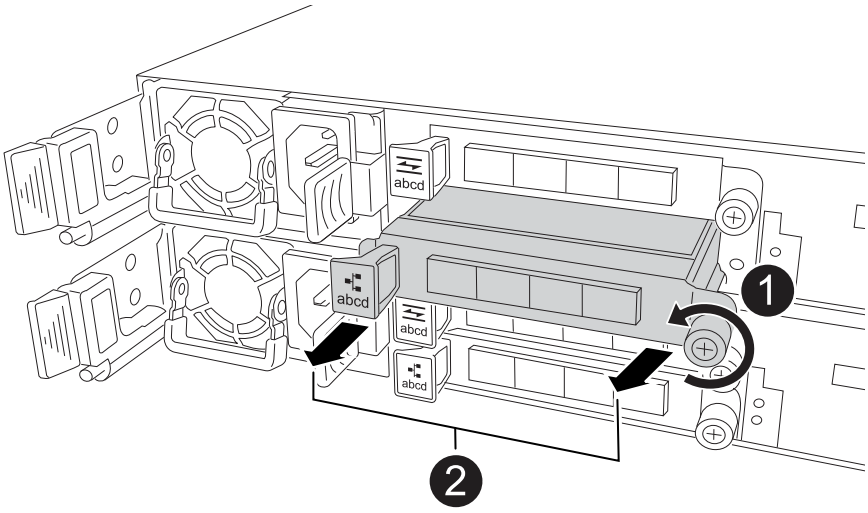
To replace a failed I/O module, locate it in the controller and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug cabling from the failed I/O module.

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

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



1	Turn the I/O module thumbscrew counterclockwise to loosen.
2	Pull the I/O module out of the controller using the port label tab on the left and the thumbscrew.

4. Install the replacement I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Gently push the I/O module all the way into the slot, making sure to properly seat the module into the connector.

You can use the tab on the left and the thumbscrew to push in the I/O Module.

c. Turn the thumbscrew clockwise to tighten.

5. Cable the I/O module.

Step 3: Reboot the controller

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

Steps

1. Reboot the controller from the LOADER prompt:

```
bye
```

Rebooting the impaired controller also reinitializes the I/O modules and other components.

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

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

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

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

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

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

Step 4: Return the failed part to NetApp

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

Replace the NV battery - FAS50

Replace the NV battery in your FAS50 storage system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

Before you begin

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

About this task

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

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

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

Step 2: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

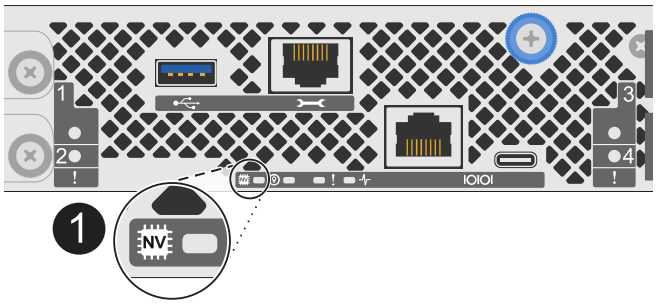
Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   > If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   > Power supplies (PSUs) do not have a power switch.

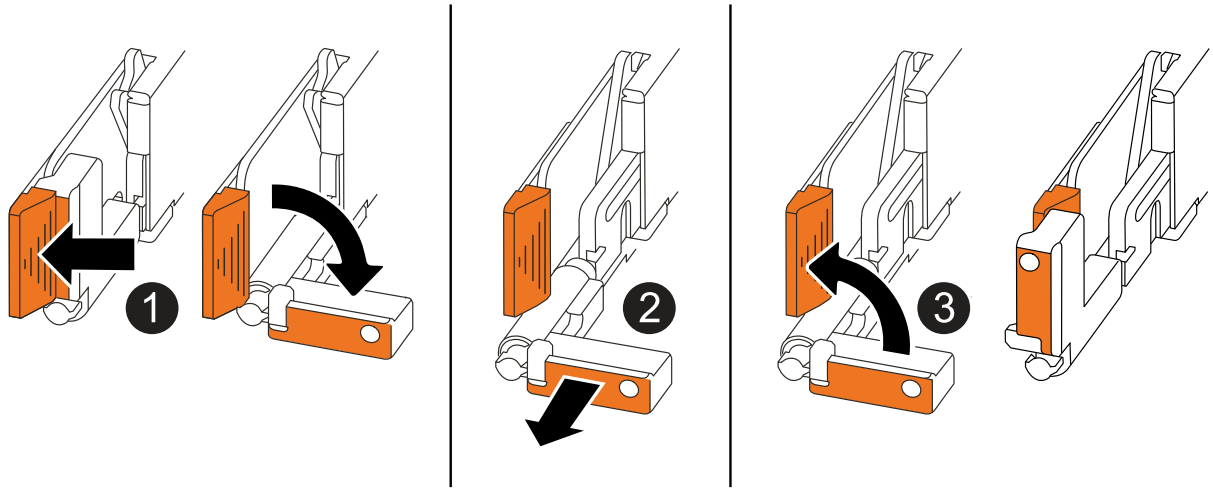
If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

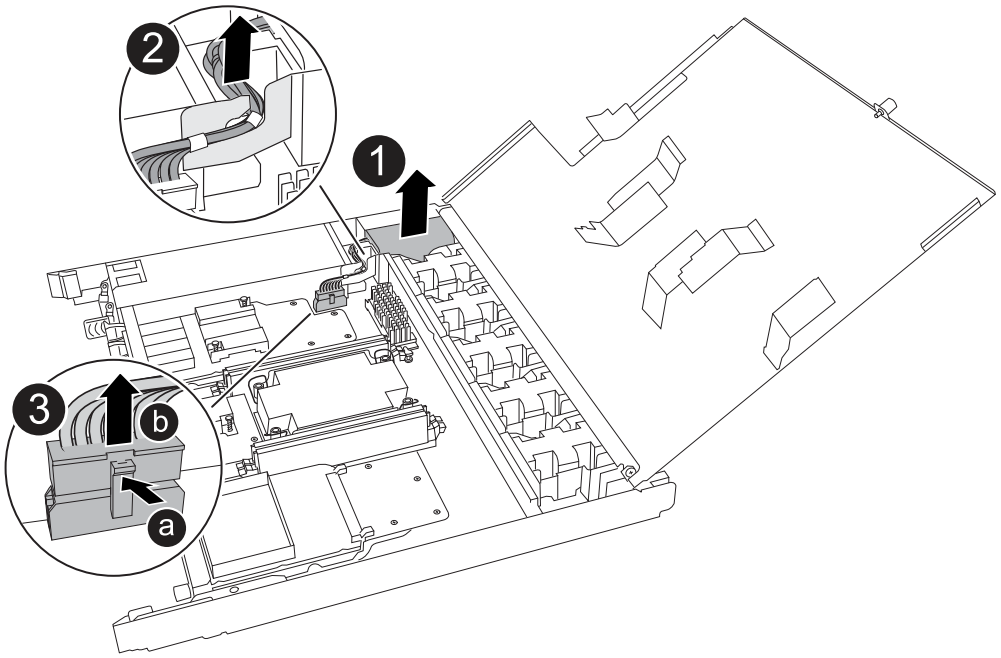
6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

Step 3: Replace the NV battery

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

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the NV battery.
3. Remove the NV battery:



1	Lift the NV battery up and out of its compartment.
2	Remove the wiring harness from its retainer.
3	1. Push in and hold the tab on the connector. 2. Pull the connector up and out of the socket. As you pull up, gently rock the connector from end to end (lengthwise) to unseat it.

4. Install the replacement NV battery:
   a. Remove the replacement battery from its package.
   b. Plug the wiring connector into its socket.
   c. Route the wiring along the side of the power supply, into its retainer, and then through the channel in front of the NV battery compartment.
   d. Place the NV battery into its compartment.

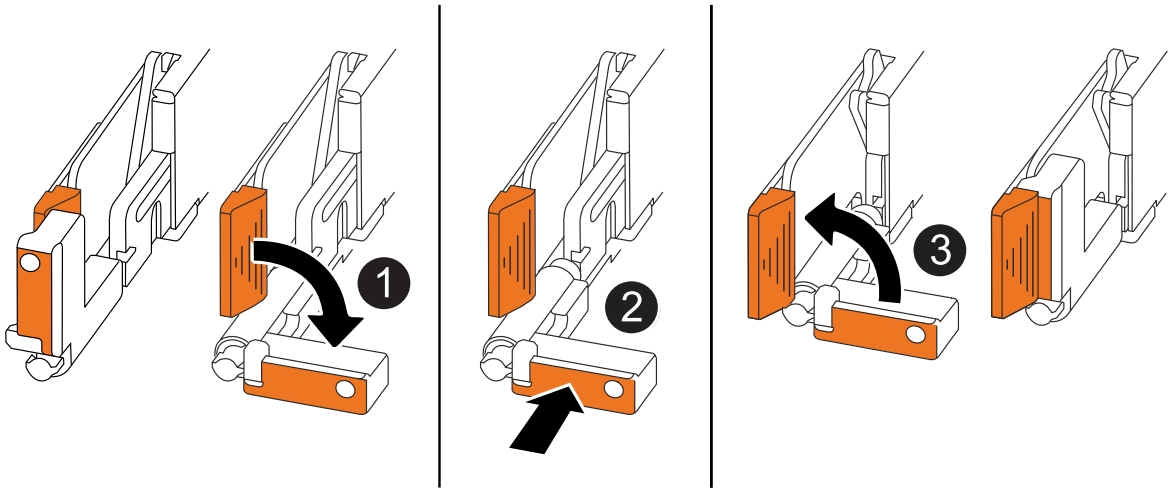
      The NV battery should sit flush in its compartment.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

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

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

b. Rotate the controller handles up and lock in place with the tabs.

The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

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

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

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

   `storage failover giveback -ofnode impaired_node_name`

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

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

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

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

Step 5: Return the failed part to NetApp

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

Replace a power supply - FAS50

Replace an AC or DC power supply unit (PSU) in your FAS50 storage system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cord, replacing the faulty PSU, and then reconnecting it to the power source.

About this task

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

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

- IMPORTANT: Do not mix PSUs with different efficiency ratings or different input types. Always replace like for like.
- Use the appropriate procedure for your type of PSU: AC or DC.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led *on*` command.

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

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

Option 1: Hot-swap an AC PSU

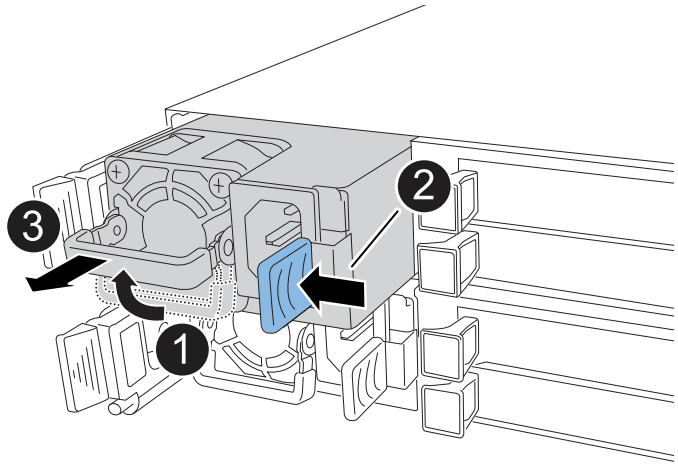
To replace an AC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the power cord from the PSU by opening the power cord retainer, and then unplug the power cord from the PSU.

   PSUs do not have a power switch.

4. Remove the PSU:



1	Rotate the PSU handle up, to its horizontal position, and then grasp it.
2	With your thumb, press the blue tab to release the PSU from the controller.
3	Pull the PSU out of the controller while using your other hand to support its weight. The PSU is short. Always use two hands to support it when removing it from the controller so that it does not suddenly swing free from the controller and injure you.

5. Install the replacement PSU:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently push the PSU into the controller until the locking tab clicks into place.

      A PSU will only properly engage with the internal connector and lock in place one way.

To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

c. Rotate the handle down, so it is out of the way of normal operations.

6. Reconnect the power cord to the PSU and secure the power cord with the power cord retainer.

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

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

Option 2: Hot-swap a DC PSU

To replace a DC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the faulty PSU based on console error messages or through the red Attention LED on the PSU.
3. Disconnect the PSU:

   PSUs do not have a power switch.

   a. Unscrew the two thumb screws on the D-SUB DC power cord connector.

      The illustration and table in step 4 shows the two thumb screws (item #1) and the D-SUB DC power cord connector (item #2).

   b. Unplug the cord from the PSU and set it aside.

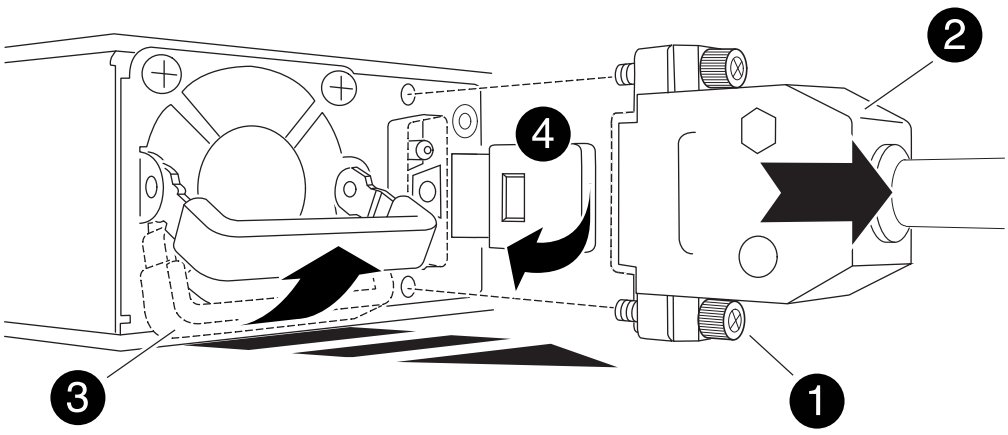
4. Remove the PSU:

   a. Rotate the handle up, to its horizontal position, and then grasp it.

   b. With your thumb, press the terracotta tab to release the locking mechanism.

   c. Pull the PSU out of the controller while using your other hand to support its weight.

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



1	Thumb screws
2	D-SUB DC power PSU cord connector
3	Power supply handle
4	Terracotta PSU locking tab

5. Insert the replacement PSU:
   a. Using both hands, support and align the edges of the PSU with the opening in the controller.
   b. Gently slide the PSU into the controller until the locking tab clicks into place.

      A PSU must properly engage with the internal connector and locking mechanism. Repeat this step if you feel the PSU is not properly seated.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the controller.

   c. Rotate the handle down, so it is out of the way of normal operations.
6. Reconnect the D-SUB DC power cord:

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

   a. Plug the D-SUB DC power cord connector into the PSU.
   b. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - FAS50

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your FAS50 storage system to ensure that services and applications relying on accurate time synchronization remain operational.

Before you begin

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

About this task

- You can use this procedure with all versions of ONTAP supported by your storage system.
- If needed, you can turn on the storage system location (blue) LEDs to aid in physically locating the affected storage system. Log into the BMC using SSH and enter the `system location-led on` command.

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

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

Step 1: Shut down the impaired controller

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

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

Step 2: Remove the controller

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

Before you begin

Make sure all other components in the storage system are functioning properly; if not, you must contact NetApp Support before continuing with this procedure.

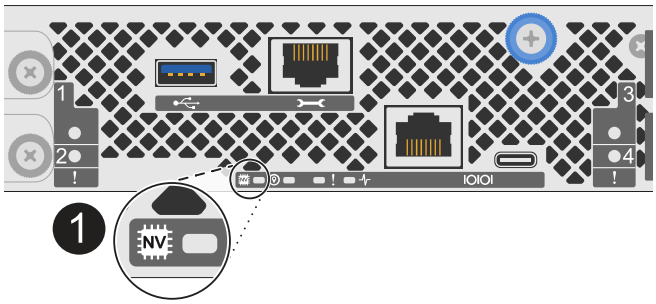
Steps

1. On the impaired controller, make sure the NV LED is off.

   When the NV LED is off, destaging is complete and it is safe to remove the impaired controller.

   If the NV LED is flashing (green), destage is in progress. You must wait for the NV LED to turn off. However, if the flashing continues for longer than five minutes, contact NetApp Support before continuing with this procedure.

   The NV LED is located next to the NV icon on the controller.



1	NV icon and LED on the controller

2. If you are not already grounded, properly ground yourself.
3. Disconnect the power on the impaired controller:

   Power supplies (PSUs) do not have a power switch.

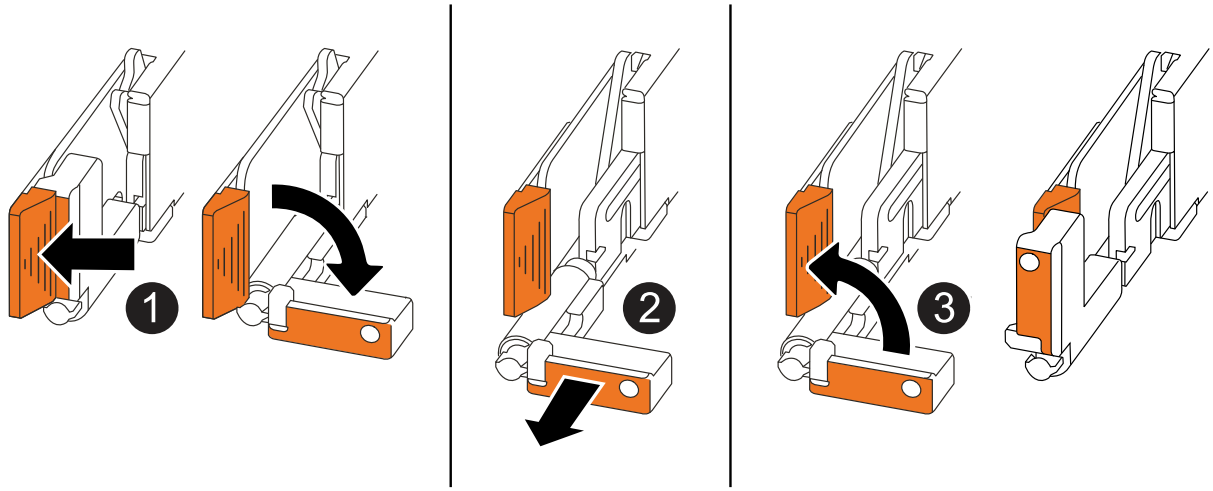
If you are disconnecting a…	Then…
AC PSU	1. Open the power cord retainer. 2. Unplug the power cord from the PSU and set it aside.
DC PSU	1. Unscrew the two thumb screws on the D-SUB DC power cord connector. 2. Unplug the power cord from the PSU and set it aside.

4. Unplug all cables from the impaired controller.

   Keep track of where the cables were connected.

5. Remove the impaired controller:

   The following illustration shows the operation of the controller handles (from the left side of the controller) when removing a controller:



1	On both ends of the controller, push the vertical locking tabs outward to release the handles.
2	• Pull the handles towards you to unseat the controller from the midplane. As you pull, the handles extend out from the controller and then you feel some resistance, keep pulling. • Slide the controller out of the chassis while supporting the bottom of the controller, and place it on a flat, stable surface.
3	If needed, rotate the handles upright (next to the tabs) to move them out of the way.

6. Open the controller cover by turning the thumbscrew counterclockwise to loosen, and then open the cover.

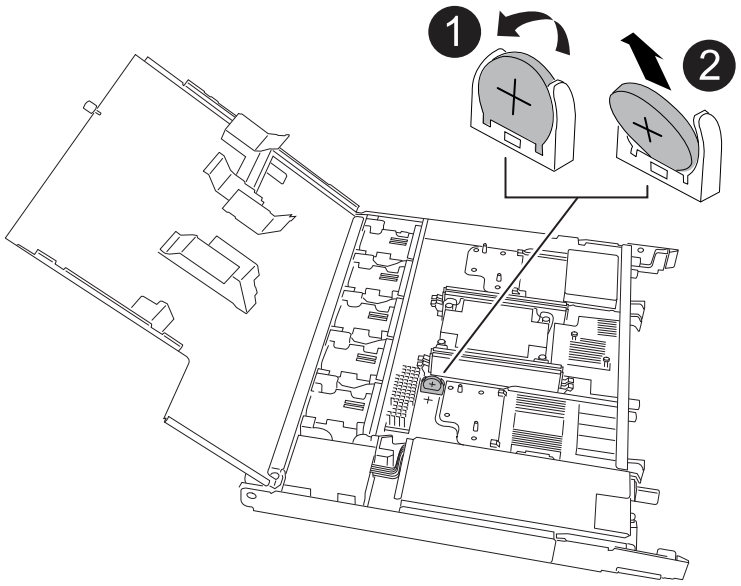
Step 3: Replace the RTC battery

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

You must use an approved RTC battery.

Steps

1. Locate the RTC battery.
2. Remove the RTC battery:



1	Gently rotate the RTC battery at an angle away from its holder.
2	Lift the RTC battery out of its holder.

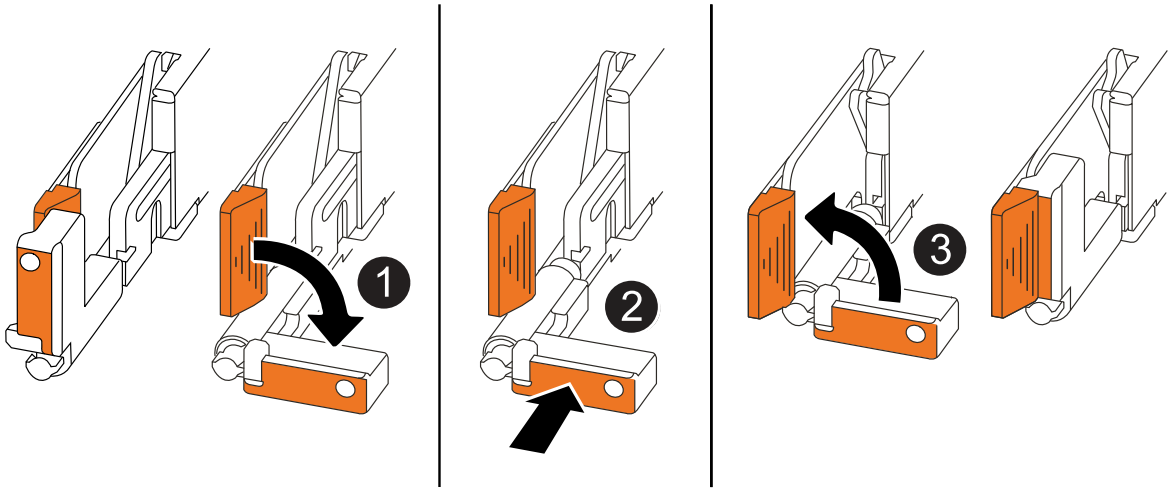
3. Install the replacement RTC battery:
   a. Remove the replacement battery from the antistatic shipping bag.
   b. Position the battery so that the plus sign on the battery faces out to correspond with the plus sign on the motherboard.
   c. Insert the battery into the holder at an angle, and then push it into an upright position so it is fully seated in the holder.
   d. Visually inspect the battery to make sure that it is completely seated in its holder and that the polarity is correct.

Step 4: Reinstall the controller

Reinstall the controller into the chassis and reboot it.

About this task

The following illustration shows the operation of the controller handles (from the left side of a controller) when reinstalling the controller, and can be used as a reference for the rest of the controller reinstallation steps.



1	If you rotated the controller handles upright (next to the tabs) to move them out of the way while you serviced the controller, rotate them down to the horizontal position.
2	Push the handles to reinsert the controller into the chassis halfway and then, when instructed, push until the controller is fully seated.
3	Rotate the handles to the upright position and lock in place with the locking tabs.

Steps

1. Close the controller cover and turn the thumbscrew clockwise until tightened.
2. Insert the controller halfway into the chassis.

   Align the rear of the controller with the opening in the chassis, and then gently push the controller using the handles.

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

3. Connect the console cable to the console port on the controller and to the laptop so that the laptop receives console messages when the controller reboots.

   Do not connect any other cables or power cords at this time.

4. Fully seat the controller in the chassis:
   a. Firmly push on the handles until the controller meets the midplane and is fully seated.

      Do not use excessive force when sliding the controller into the chassis; it could damage the connectors.

   b. Rotate the controller handles up and lock in place with the tabs.

      The replacement controller receives power from the healthy controller and begins to boot as soon as it is fully seated in the chassis.

5. Recable the controller as needed.
6. Reconnect the power cord to the power supply (PSU).

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

If you are reconnecting a…	Then…
AC PSU	1. Plug the power cord into the PSU. 2. Secure the power cord with the power cord retainer.
DC PSU	1. Plug the D-SUB DC power cord connector into the PSU. 2. Tighten the two thumb screws to secure the D-SUB DC power cord connector to the PSU.

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

   `storage failover giveback -ofnode impaired_node_name`

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

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

9. If AutoSupport is enabled, restore (unsuppress) automatic case creation:

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

Step 5: Reset the time and date on the controller

After replacing the RTC battery, inserting the controller and powering on first BIOS reset, you will see the following error messages:
`RTC date/time error. Reset date/time to default`
`RTC power failure error`
These messages are expected and you can continue with this procedure.

1. On the healthy controller, check the date and time:

   `cluster date show`

   If your storage system stops at the boot menu, select the option for `Reboot node` and respond *y* when prompted, then boot to LOADER by pressing *Ctrl-C*.

2. On the impaired controller, at the LOADER prompt, check the time and date:

   `cluster date show`

   a. If necessary, modify the date:

      `set date mm/dd/yyyy`

   b. If necessary, set the time, in GMT:

```
set time hh:mm:ss
```

c. Confirm the date and time.

3. At the LOADER prompt, enter `bye` to reinitialize the I/O modules, other components, and let the controller reboot.

Step 6: Return the failed part to NetApp

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

Key specifications for FAS50

The following are select specifications for the FAS50. Visit NetApp Hardware Universe (HWU) for a complete list of FAS50 specifications. This page is reflective of a single high availability pair.

Key specifications for FAS50

Platform Configuration: FAS50 Single Chassis HA Pair

Max Raw Capacity: 10.5600 PB

Memory: 128.0000 GB

Form Factor: 2U chassis with 2 HA controllers

ONTAP Version: b_startONTAP: 9.16.1P2b_end

PCIe Expansion Slots: 12

Minimum ONTAP Version: ONTAP 9.16.1

Scaleout Maximums

Type	HA Pairs	Raw Capacity	Max Memory
NAS	4	42.2 PB / 37.5 PiB	512 GB
SAN	4	42.2 PB / 37.5 PiB	512 GB
HA Pair		10.6 PB / 9.4 PiB	128.0000

IO

Onboard IO

No onboard IO data.

Total IO

Protocol	Ports
Ethernet 100 Gbps	16

Ethernet 25 Gbps	24
Ethernet 10 Gbps	24
FC 64 Gbps	24
NVMe/FC 64 Gbps	24
	0
SAS 12 Gbps	16

Management Ports

Protocol	Ports
Ethernet 1 Gbps	2
RS-232 115 Kbps	4
USB 600 Mbps	2

Storage Networking Supported

CIFS;
FC;
iSCSI;
NFS v3;
NFS v4.0;
NFS v4.1;
NFS v4.2;
NVMe/FC ;
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: 2319 BTU/hr
- Worst-case Power: 3872 BTU/hr
- Weight: 52.2 lb 23.7 kg
- Height: 2U
- Width: 19” IEC rack-compliant (17.6” 44.7 cm)
- Depth:
- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°F) at up to 3048m (10000 ft) elevation;8% to 80% 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): 8.0; Sound pressure (LpAm) (bystander positions): 70.5 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-55032,
  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 15598-1,
  EN60825-1,
  EN62368-1,
  IEC 62368-1,
  IEC60950-1,
  IS 13252(part 1)

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 for external shelves;

FAS70 and FAS90 systems

Install and setup

Installation and configuration workflow - FAS70 and FAS90

To install and configure your FAS70 or FAS90 system, you review the hardware requirements, prepare your site, install and cable the hardware components, power on the system, and set up your ONTAP cluster.



Review installation requirements

Review the equipment and tools needed to install your storage system and storage shelves and review the lifting and safety precautions.



Prepare to install the FAS70 or FAS90 storage system

To prepare to install your system, you need to get the site ready, check the environmental and electrical requirements, and ensure there's enough rack space. Then, unpack the equipment, compare its contents to the packing slip, and register the hardware to access support benefits.



Install the hardware for the FAS70 or FAS90 storage system

To install the hardware, install the rail kits for your storage system and shelves, and then install and secure your storage system in the cabinet or telco rack. Next, slide the shelves onto the rails. Finally, attach cable management devices to the rear of the storage system for organized cable routing.



Cable the controllers and storage shelves for the FAS70 or FAS90 storage system

To cable the hardware, first connect the storage controllers to your network and then connect the controllers to your storage shelves.



Power on the FAS70 or FAS90 storage system

Before you power on the controllers, power on each NS224 shelf and assign a unique shelf ID to ensure each shelf is uniquely identified within the setup, connect the laptop or console to the controller, and then connect the controllers to the power sources.



Set up your cluster

After you've powered on your storage system, you set up your cluster.

Installation requirements - FAS70 and FAS90

Review the equipment needed and the lifting precautions for your FAS70 or FAS90 storage system and storage shelves.

Equipment needed for install

To install your storage system, you need the following equipment and tools.

- Access to a Web browser to configure your storage system
- Electrostatic discharge (ESD) strap
- Flashlight
- Laptop or console with a USB/serial connection
- Phillips #2 screwdriver

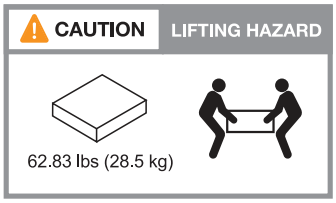
Lifting precautions

Storage systems and shelves are heavy. Exercise caution when lifting and moving these items.

Storage system weight

Take the necessary precautions when moving or lifting your storage system.

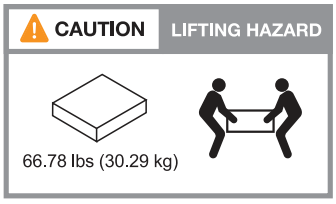
A FAS70 or FAS90 storage system can weigh up to 62.83 lbs (28.5 kg). To lift the system, use two people or a hydraulic lift.



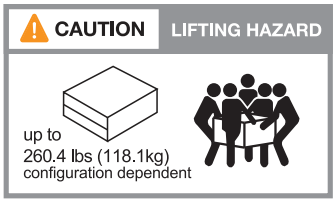
Shelf weight

Take the necessary precautions when moving or lifting your shelf.

An NS224 shelf can weigh up to 66.78 lbs (30.29 kg). To lift the shelf, use two people or a hydraulic lift. Keep all components in the shelf (both front and rear) to prevent unbalancing the shelf weight.



A DS460C shelf can weigh up to 260.4 lbs (181.1 kg). To lift the storage shelf, you might need up to five people or a hydraulic lift. Keep all components in the storage shelf (both front and rear) to prevent unbalancing the shelf weight.



Related information

- Safety information and regulatory notices

What's next?

After you've reviewed the hardware requirements, you prepare to install your FAS70 or FAS90 storage system.

Prepare to install - FAS70 and FAS90

Prepare to install your FAS70 or FAS90 storage system by getting the site ready, unpacking the boxes and comparing the contents of the boxes to the packing slip, and registering the system to access support benefits.

Step 1: Prepare the site

To install your storage system, ensure that the site and the cabinet or rack that you plan to use meet specifications for your configuration.

Steps

1. Use NetApp Hardware Universe to confirm that your site meets the environmental and electrical requirements for your storage system.
2. Make sure you have adequate cabinet or rack space for your storage system, shelves, and any switches:
   - 4U in an HA configuration
   - 2U for each NS224 storage shelf
3. Install any required network switches.

   See the Switch documentation for installation instructions and NetApp Hardware Universe for compatibility information.

Step 2: Unpack the boxes

After you've ensured that the site and the cabinet or rack that you plan to use for your storage system meet the required specifications, unpack all boxes and compare the contents to the items on the packing slip.

Steps

1. Carefully open all the boxes and lay out the contents in an organized manner.
2. Compare the contents you've unpacked with the list on the packing slip.

   You can get your packing list by scanning the QR code on the side of the shipping carton.

   The following items are some of the contents you might see in the boxes.

   Ensure that everything in the boxes matches the list on the packing slip. If there are any discrepancies, note them down for further action.

Hardware

- Bezel
- Cable management device
- Storage system
- Rail kits with instructions (optional)
- Storage shelf (if you ordered additional storage)

Cables

- Management Ethernet cables (RJ-45 cables)
- Network cables
- Power cords
- Storage cables (if you ordered additional storage)
- USB-C serial console cable

Step 3: Register your storage system

After you've ensured that your site meets the requirements for your storage system specifications, and you've verified that you have all the parts you ordered, you should register your storage system.

Steps

1. Locate the System Serial Numbers (SSN) for every controller being installed.

   You can find the serial numbers in the following locations:

   - On the packing slip
   - In your confirmation email
   - On each controller's System Management module

2. Go to the NetApp Support Site.
3. Determine whether you need to register your storage system:

If you are a…	Follow these steps…
Existing NetApp customer	a. Sign in with your username and password. b. Select **Systems** > **My Systems**. c. Confirm that the new serial numbers are listed. d. If it is not, follow the instructions for new NetApp customers.
New NetApp customer	a. Click **Register Now**, and create an account. b. Select **Systems** > **Register Systems**. c. Enter the storage system's serial numbers and requested details. After your registration is approved, you can download any required software. The approval process might take up to 24 hours.

What's next?

After you've prepared to install your FAS70 or FAS90 hardware, you install the hardware for your FAS70 or FAS90 storage system.

Install the hardware - FAS70 and FAS90

After you prepare to install your FAS70 or FAS90 storage system, install the hardware for the system. First, install the rail kits. Then install and secure your platform in a cabinet or telco rack.

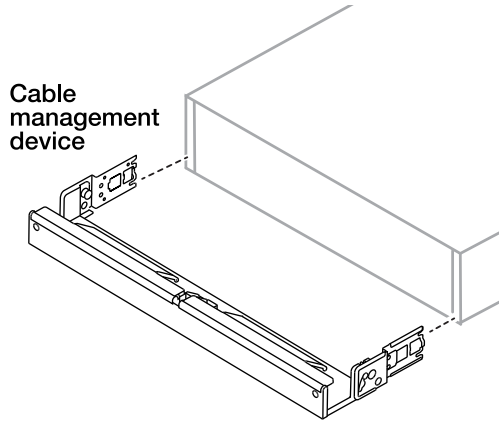
Skip this step if your cabinet is pre-populated.

Before you begin

- Make sure you have the instructions packaged with the rail kit.
- Be aware of the safety concerns associated with the weight of the storage system and shelf.
- Understand that the airflow through the storage system enters from the front where the bezel or end caps are installed and exhausts out the rear where the ports are located.

Steps

1. Install the rail kits for your storage system and shelves as needed, using the instructions included with the kits.
2. Install and secure your storage system in the cabinet or telco rack:
   a. Position the storage system onto the rails in the middle of the cabinet or telco rack, and then support the storage system from the bottom and slide it into place.
   b. Make sure that the guiding pins of the cabinet or telco rack are securley in the chassis guide slots.
   c. Secure the storage system to the cabinet or telco rack using the included mounting screws.
3. Attach the bezel to the front of the storage system.
4. Attach the cable management devices to the rear of the storage system.



5. Install and secure the shelf as needed.
   a. Position the back of the shelf onto the rails, and then support the shelf from the bottom and slide it into the cabinet or telco rack.

      If you are installing multiple shelves, place the first shelf directly above the controllers. Place the second shelf directly under the controllers. Repeat this pattern for any additional shelves.

   b. Secure the shelf to the cabinet or telco rack using the included mounting screws.

What's next?

After you've installed the hardware for your FAS70 or FAS90 storage system, you cable the hardware for your FAS70 or FAS90 storage system..

Cable the hardware - FAS70 and FAS90

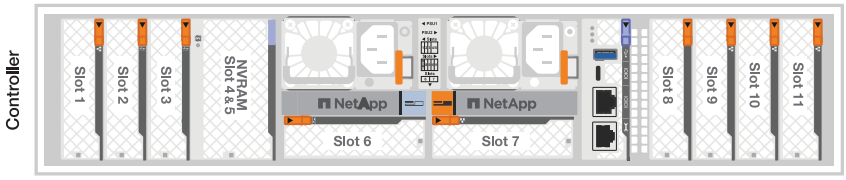
After you install the rack hardware for your FAS70 or FAS90 storage system, install the network cables for the controllers, and connect the cables between the controllers and storage shelves.

Before you begin

Contact your network administrator for information about connecting the storage system to the switches.

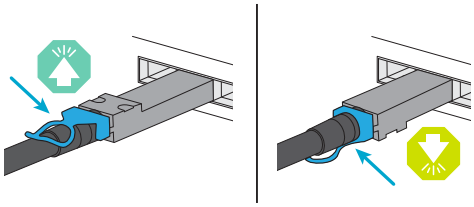
About this task

- These procedures show common configurations. The specific cabling depends on the components ordered for your storage system. For comprehensive configuration and slot priority details, see NetApp Hardware Universe.
- The I/O slots on FAS70 and FAS90 controllers are numbered 1 through 11.



- The cabling graphics have arrow icons showing the proper orientation (up or down) of the cable connector pull-tab when inserting a connector into a port.

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



- If cabling to an optical switch, insert the optical transceiver into the controller port before cabling to the switch port.

Step 1: Connect the storage controllers to your network

Cable the controllers to your ONTAP cluster. This procedure differs depending on your storage system model and I/O module configuration.

The cluster interconnect traffic and the HA traffic share the same physical ports.

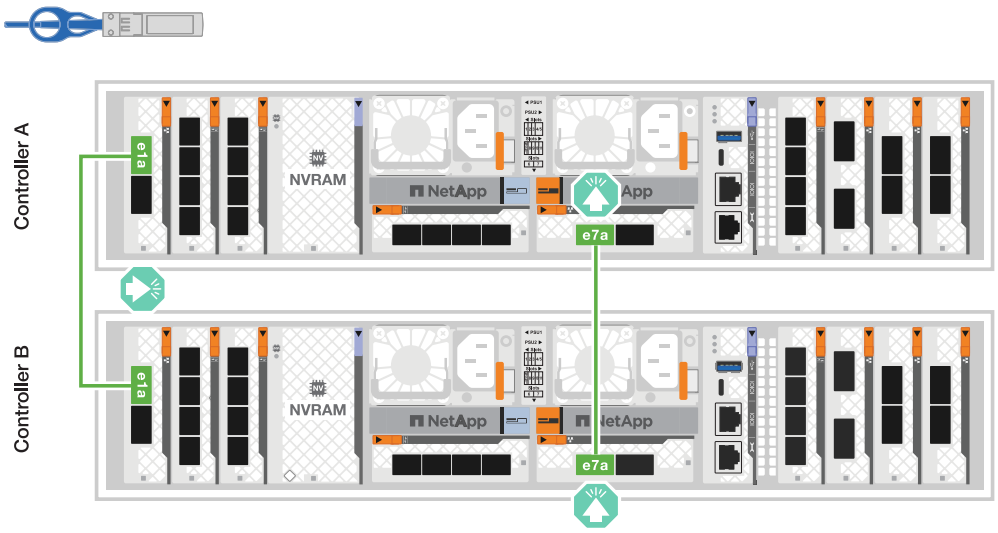
Switchless cluster cabling

Use the Cluster/HA interconnect cable to connect ports e1a to e1a and ports e7a to e7a.

Steps

1. Connect port e1a on Controller A to port e1a on Controller B.
2. Connect port e7a on Controller A to port e7a on Controller B.

Cluster/HA interconnect cables



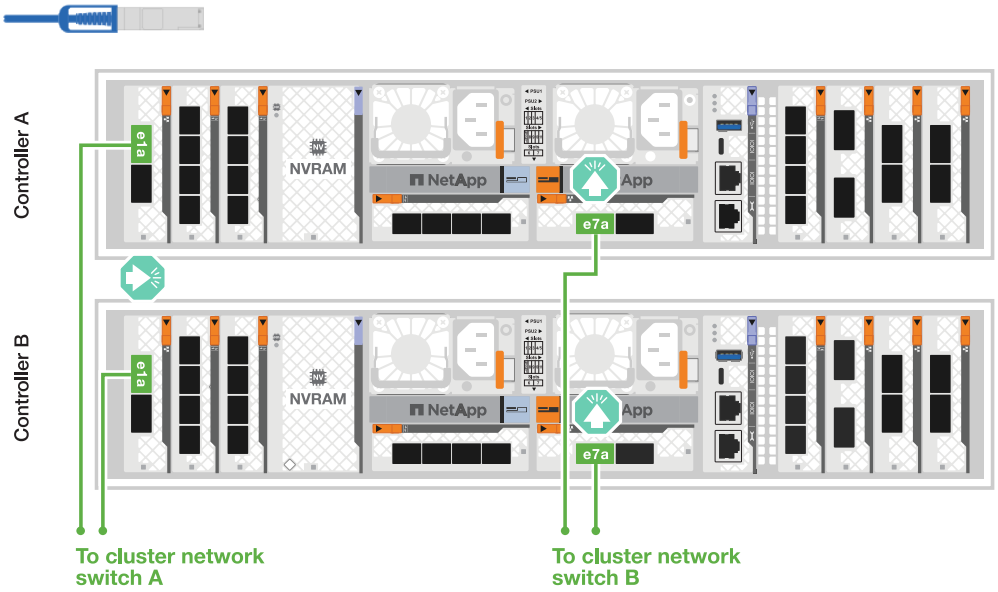
Switched cluster cabling

Use the 100 GbE cable to connect ports e1a to e1a and ports e7a to e7a.

Steps

1. Connect port e1a on Controller A and port e1a on Controller B to cluster network switch A.
2. Connect port e7a on Controller A and port e7a on Controller B to cluster network switch B.

100 GbE cable



Step 2: Cable the host network connections

Connect the Ethernet module ports to your host network.

The following are some typical host network cabling examples. See NetApp Hardware Universe for your specific system configuration.

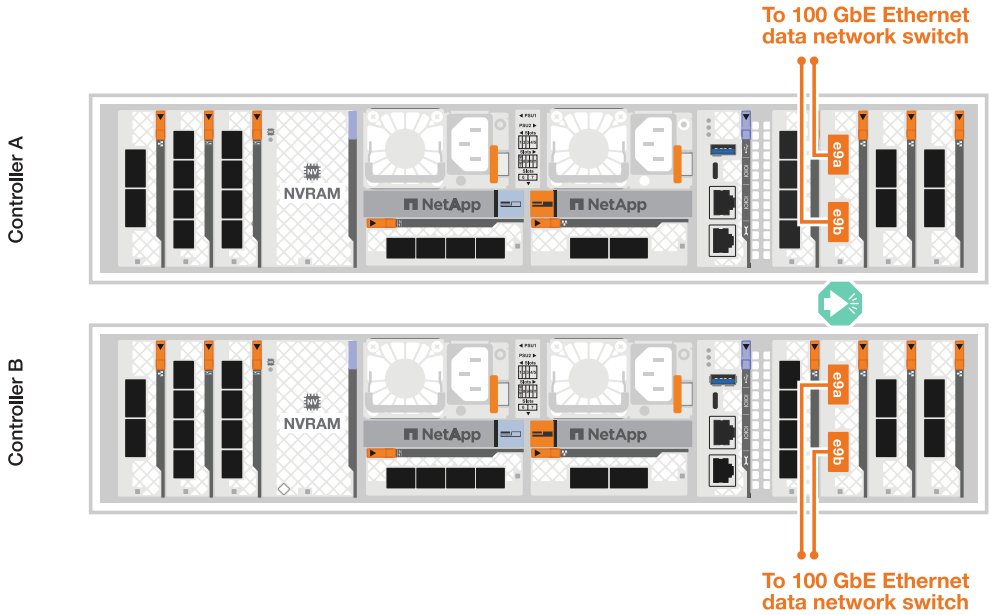
Steps

1. Connect ports e9a and e9b to your Ethernet data network switch.

   Do not use ports e1b and e7b ports for host network connections. Use a separate host card.

   100 GbE cable

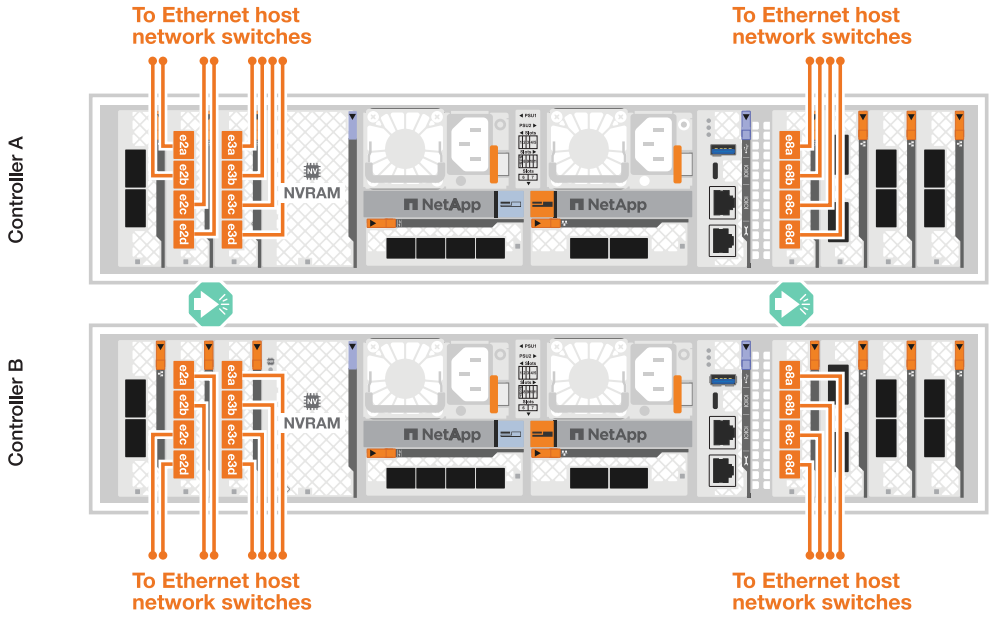
   

2. Connect your 10/25 GbE host network switches.

   4-ports, 10/25 GbE Host

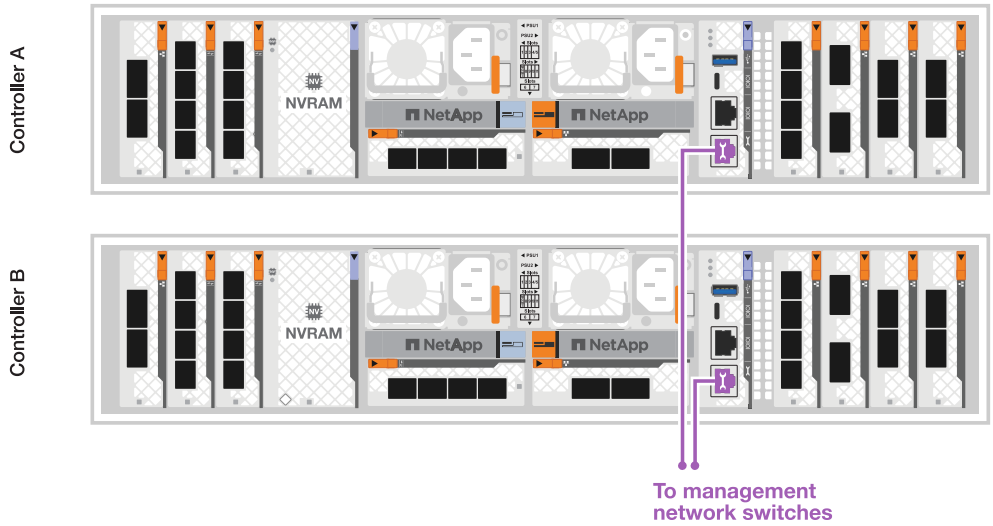


Step 3: Cable the management network connections

Use the 1000BASE-T RJ-45 cables to connect the management (wrench) ports on each controller to the management network switches.



1000BASE-T RJ-45 cables



Do not plug in the power cords yet.

Step 4: Cable the shelf connections

The following cabling procedures show how to connect your controllers to a storage shelf. Choose one of the following cabling options that matches your setup.

For the maximum number of shelves supported for your storage system and for all of your cabling options, see NetApp Hardware Universe.

For additional SAS shelf cabling guidance, see SAS cabling rules and concepts - shelves with IOM12/IOM12B modules.

About this task

The FAS70 and FAS90 storage systems supports DS212C, DS224C, DS460C, and NS224 shelves with either the NSM100 or NSM100B module.

The major differences between the NS224 modules are:

- NSM100 shelf modules use built-in ports e0a and e0b.
- NSM100B shelf modules use ports e1a and e1b in slot 1.

The following NS224 cabling example shows NSM100 modules in the NS224 shelves when referring to shelf module ports.

Option 1: One NS224 storage shelf

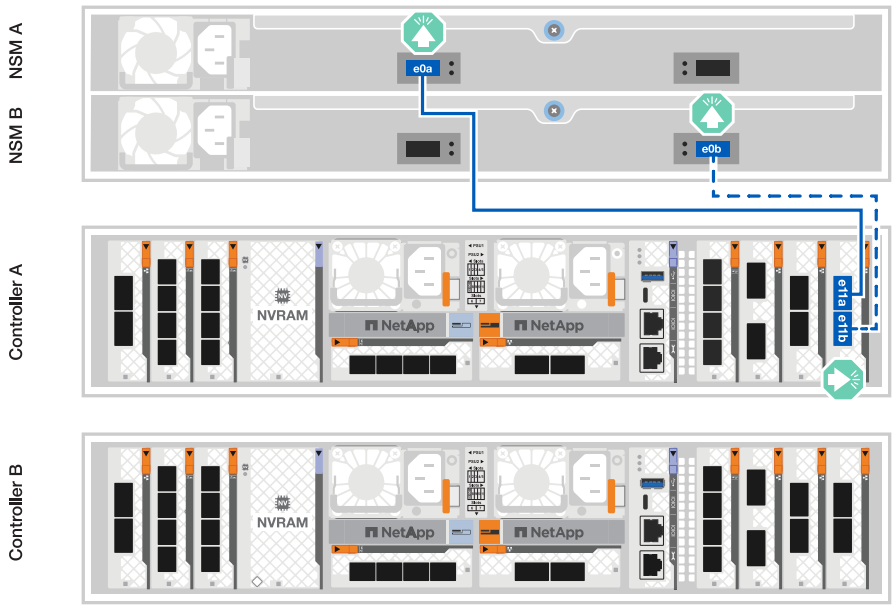
Connect each controller to the NSM modules on the NS224 shelf. The graphics show controller A cabling in blue and controller B cabling in yellow.

100 GbE QSFP28 copper cables

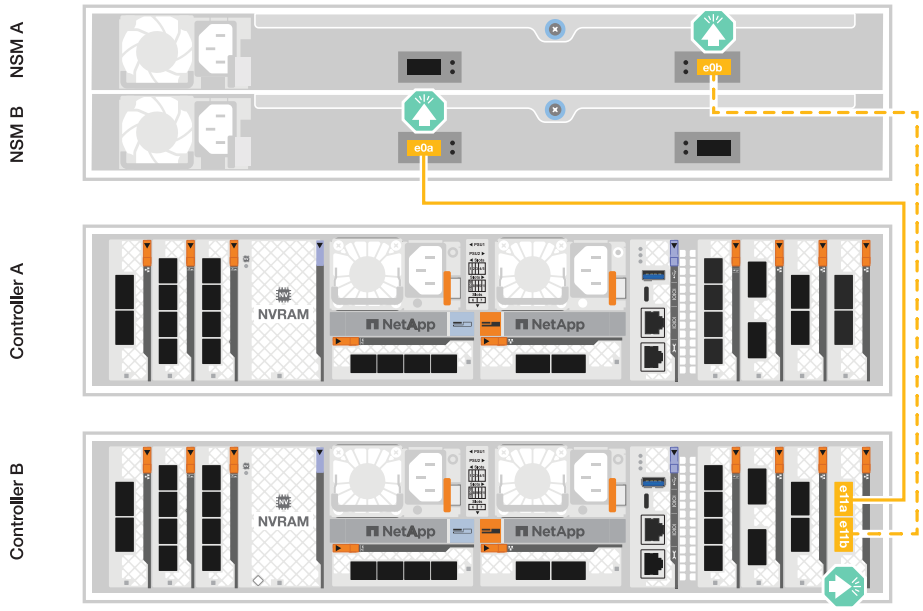


Steps

1. On controller A, connect the following ports:
   a. Connect port e11a to NSM A port e0a.
   b. Connect port e11b to port NSM B port e0b.



2. On controller B, connect the following ports:
   a. Connect port e11a to NSM B port e0a.
   b. Connect port e11b to NSM A port e0b.



Option 2: Two NS224 storage shelves

Cable each controller to the NSM modules on both NS224 shelves. The graphics show controller A cabling in blue and controller B cabling in yellow.

100 GbE QSFP28 copper cables



Steps

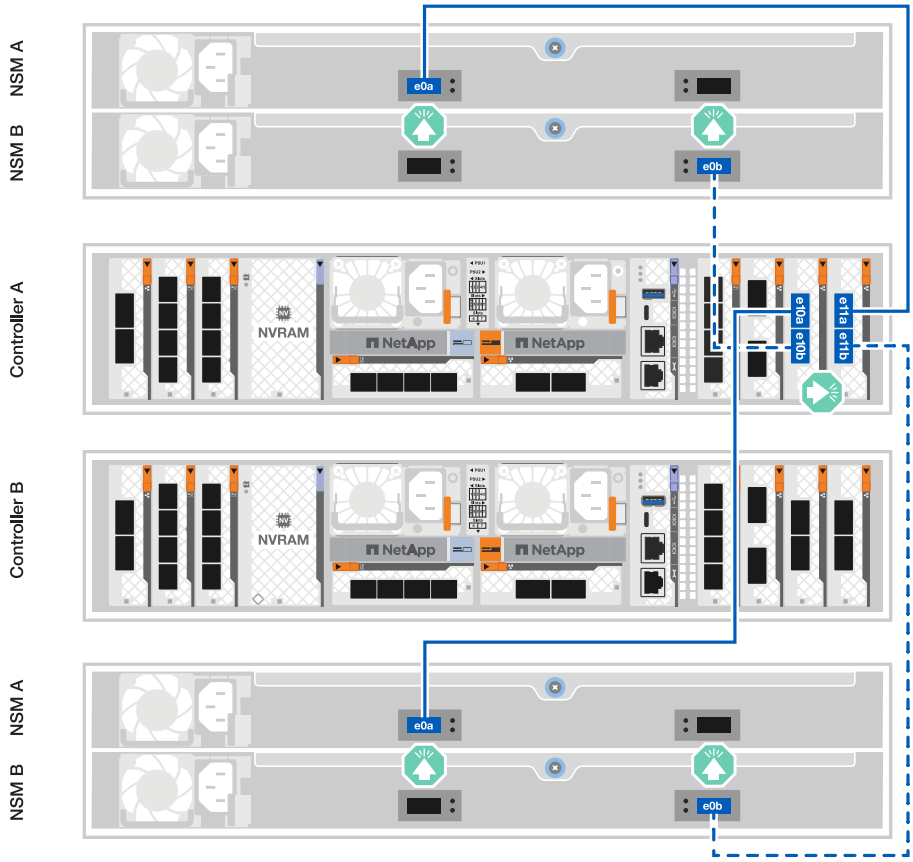
1. On controller A, connect the following ports:

    a. Connect port e11a to shelf 1 NSM A port e0a.

    b. Connect port e11b to shelf 2 NSM B port e0b.

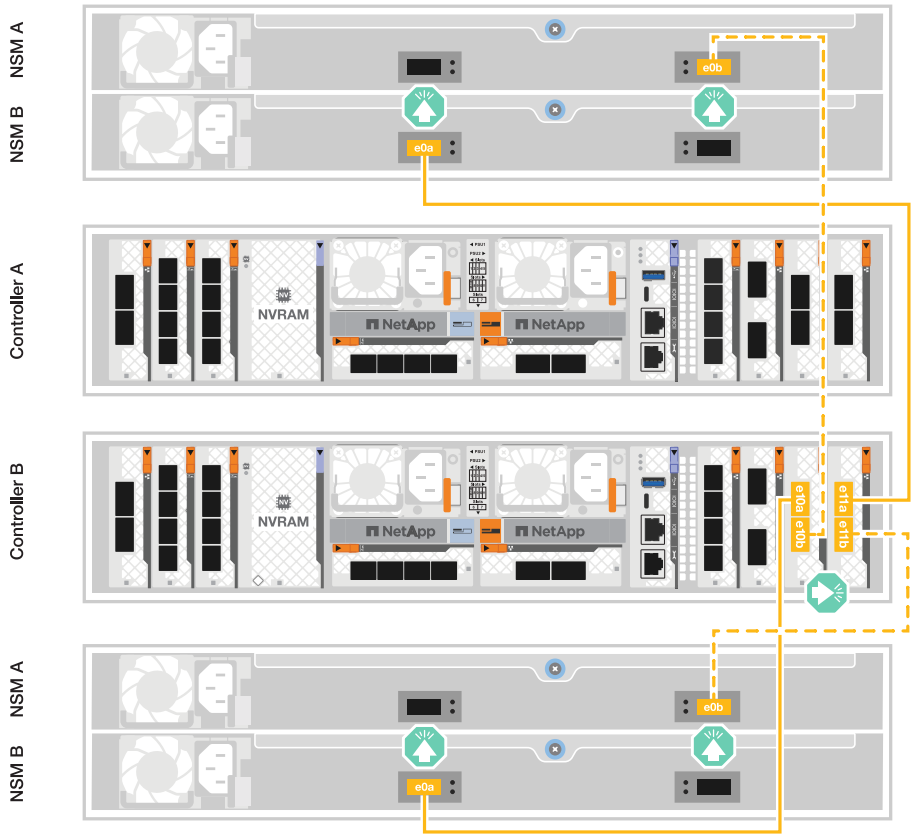
    c. Connect port e10a to shelf 2 NSM A port e0a.

    d. Connect port e10b to shelf 1 NSM A port e0b.



2. On controller B, connect the following ports:

   a. Connect port e11a to shelf 1 NSM B port e0a.
   b. Connect port e11b to shelf 2 NSM A port e0b.
   c. Connect port e10a to shelf 2 NSM B port e0a.
   d. Connect port e10b to shelf 1 NSM A port e0b.



Option 3: Two DS460C shelves

Cable each controller to the IOM modules on both DS460C shelves. The graphics show controller A cabling in blue and controller B cabling in yellow.

mini-SAS HD cable



Steps

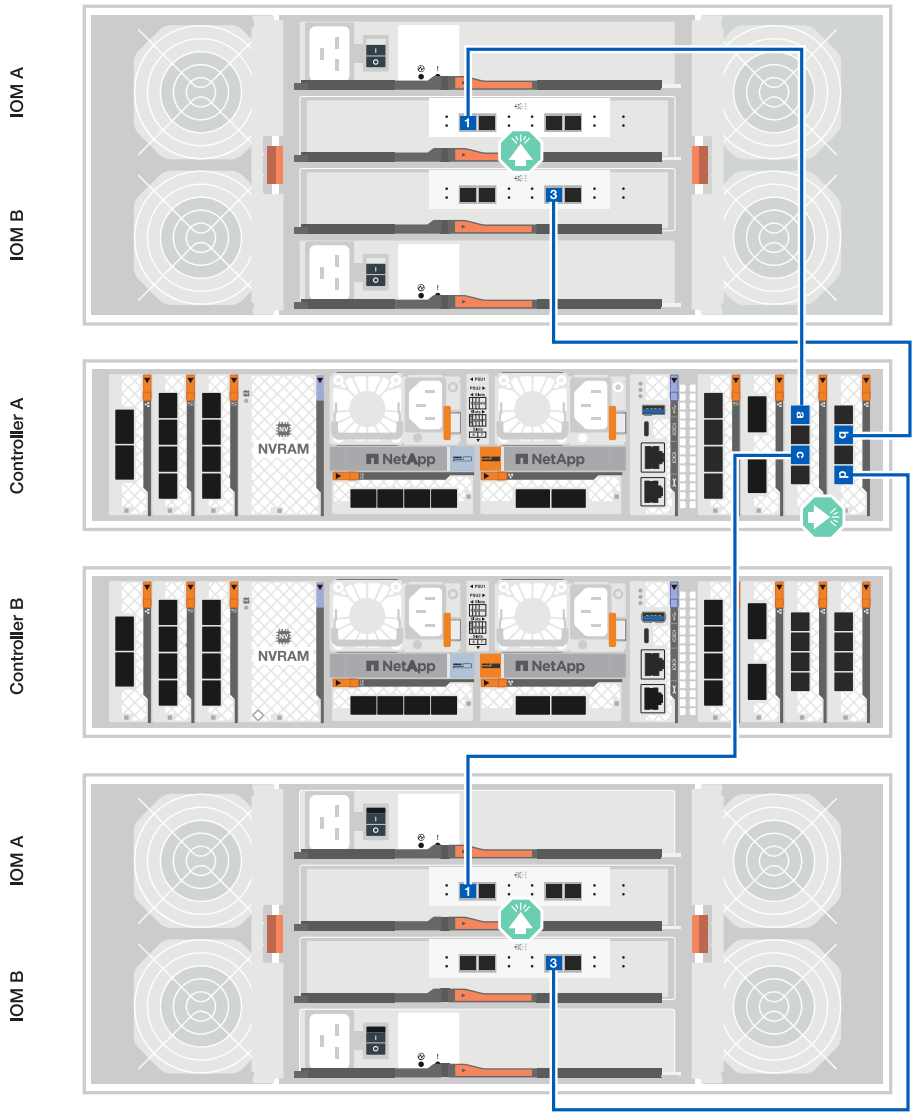
1. On controller A, cable the following connections:

    a. Connect port e10a to shelf 1 IOM A port 1.

    b. Connect port e10c to shelf 2 IOM A port 1

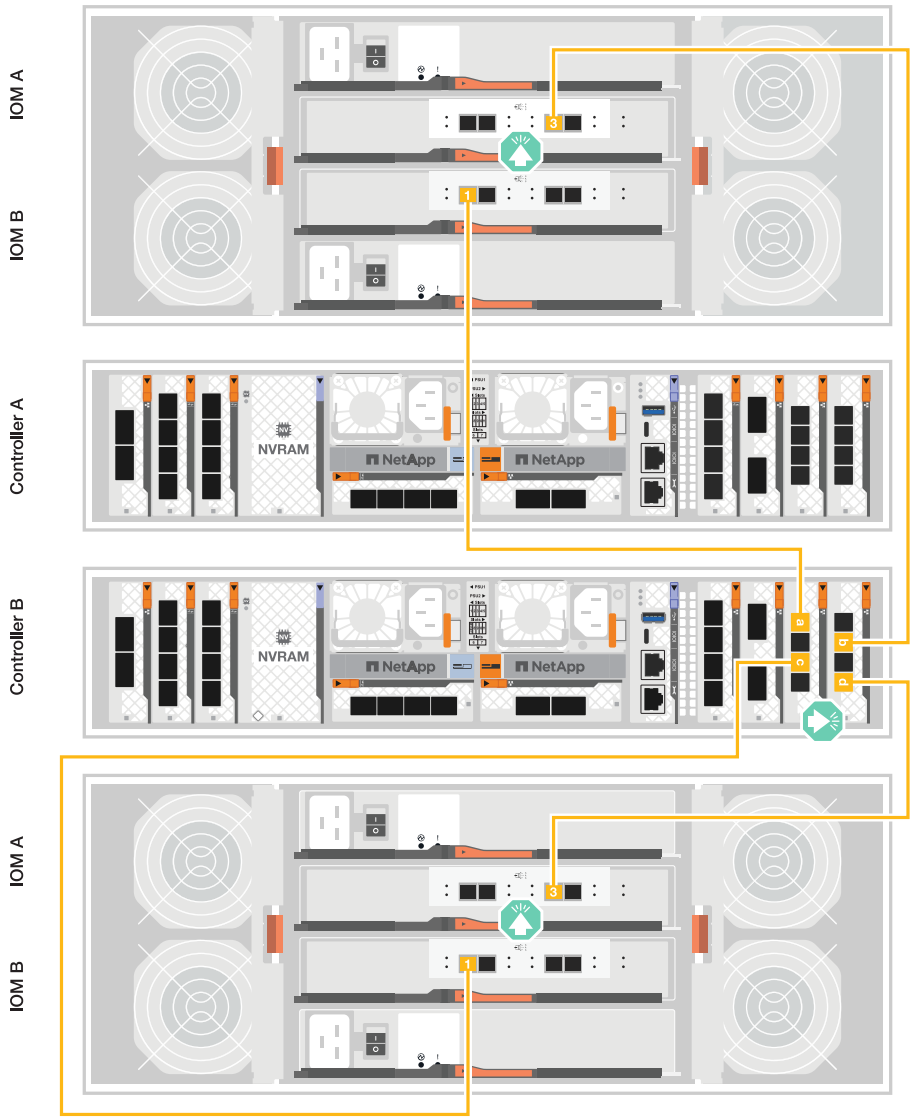
    c. Connect port e11b to shelf 1 IOM B port 3.

    d. Connect port e11d to shelf 2 IOM B port 3.



2. On controller B, cable the following connections:

   a. Connect port e10a to shelf 1 IOM B port 1.
   b. Connect port e10c to shelf 2 IOM B port 1.
   c. Connect port e11b to shelf 1 IOM A port 3.
   d. Connect port e11d to shelf 2 IOM A port 3.



What's next?

After you've cabled the hardware for your FAS70 or FAS90 system, you power on the FAS70 or FAS90 storage system.

Power on the storage system - FAS70 and FAS90

After you install the rack hardware for your FAS70 0r FAS90 storage system and install the cables for the controllers and storage shelves, you should power on your storage shelves and controllers.

Step 1: Power on the shelf and assign shelf ID

Option 1: NS224 shelves

Each shelf is distinguished by a unique shelf ID. This ID ensures that the shelf is distinct within your storage system setup.

Before you begin

Make sure you have a paperclip or narrow tipped ball point pen for setting NS224 storage shelf IDs.

About this task

- A valid shelf ID is 01 through 99.

  If you have internal shelves (storage), which are integrated within the controllers, they are assigned a fixed shelf ID of 00.

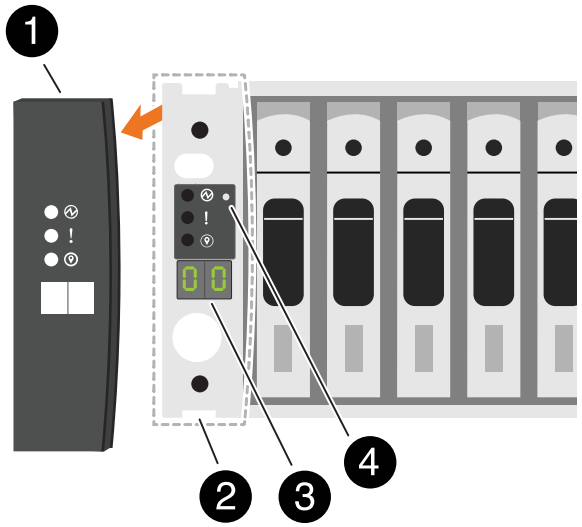
- You must power cycle a shelf (unplug both power cords, wait the appropriate amount of time, and then plug them back in) for the shelf ID to take effect.

Steps

1. Power on the shelf by connecting the power cords first to the shelf, securing them in place with the power cord retainer, and then connecting the power cords to power sources on different circuits.

   The shelf powers on and boots automatically when plugged into the power source.

2. Remove the left end cap to access the shelf ID button behind the faceplate.



1	Shelf end cap
2	Shelf faceplate
3	Shelf ID number
4	Shelf ID button

3. Change the first number of the shelf ID:

a. Insert the straightened end of a paperclip or narrow tipped ball point pen into the small hole to press the shelf ID button.

b. Press and hold the shelf ID button until the first number on the digital display blinks, and then release the button.

   It can take up to 15 seconds for the number to blink. This activates the shelf ID programming mode.

   > If the ID takes longer than 15 seconds to blink, press and hold the shelf ID button again, making sure to press it in all the way.

c. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

   Each press and release duration can be as short as one second.

   The first number continues to blink.

4. Change the second number of the shelf ID:

   a. Press and hold the button until the second number on the digital display blinks.

      It can take up to three seconds for the number to blink.

      The first number on the digital display stops blinking.

   b. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

      The second number continues to blink.

5. Lock in the desired number and exit the programming mode by pressing and holding the shelf ID button until the second number stops blinking.

   It can take up to three seconds for the number to stop blinking.

   Both numbers on the digital display start blinking and the amber LED illuminates after about five seconds, alerting you that the pending shelf ID has not yet taken effect.

6. Power-cycle the shelf for at least 10 seconds to make the shelf ID take effect.

   a. Unplug the power cord from both power supplies on the shelf.

   b. Wait 10 seconds.

   c. Plug the power cords back into the shelf power supplies to complete the power cycle.

      A power supply is powered on as soon as the power cord is plugged in. Its bicolored LED should illuminate green.

7. Replace the left end cap.

Option 2: SAS shelves

Each shelf is distinguished by a unique shelf ID. This ID ensures that the shelf is distinct within your storage system setup.

About this task

- A valid shelf ID is 01 through 99.

  If you have internal shelves (storage), which are integrated within the controllers, they are assigned a fixed shelf ID of 00.

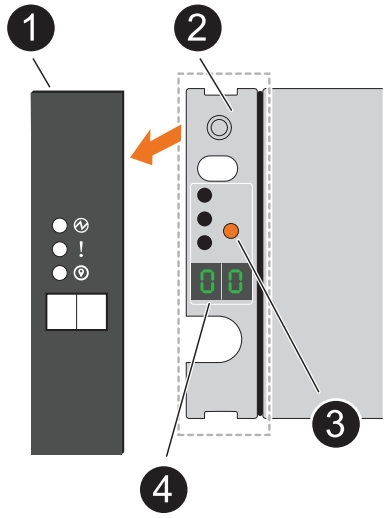
- You must power cycle a shelf (turn off the power switch on each of the power supplies of the SAS shelf, wait the appropriate amount of time, and then switch the power back on) for the shelf ID to take effect.

Steps

1. Power on the shelf by connecting the power cords first to the shelf, securing them in place with the power cord retainer, connecting the power cords to power sources on different circuits, and then turning on the power switch on each of the power supplies (at the rear of the shelf).

   The shelf powers on and boots automatically when powered on.

2. Remove the left end cap to access the orange shelf ID button on the faceplate.



1	Shelf end cap
2	Shelf faceplate
3	Shelf ID button
4	Shelf ID number

3. Change the first number of the shelf ID:
   a. Press and hold the shelf ID button until the first number on the digital display blinks, and then release the button.

      It can take up to 15 seconds for the number to blink. This activates the shelf ID programming mode.

If the ID takes longer than 15 seconds to blink, press and hold the shelf ID button again, making sure to press it in all the way.

b. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

   Each press and release duration can be as short as one second.

   The first number continues to blink.

4. Change the second number of the shelf ID:
   a. Press and hold the button until the second number on the digital display blinks.

      It can take up to three seconds for the number to blink.

      The first number on the digital display stops blinking.

   b. Press and release the shelf ID button to advance the number until you reach the desired number from 0 to 9.

      The second number continues to blink.

5. Lock in the desired number and exit the programming mode by pressing and holding the shelf ID button until the second number stops blinking.

   It can take up to three seconds for the number to stop blinking.

   Both numbers on the digital display start blinking and the amber LED illuminates after about five seconds, alerting you that the pending shelf ID has not yet taken effect.

6. Power-cycle the shelf for at least 10 seconds to make the shelf ID take effect.
   a. Turn off the power switch on each of the power supplies.
   b. Wait 10 seconds.
   c. Turn on the power switch on each of the power supplies to complete the power cycle.

      When a power supply is powered on, the bicolored LED should illuminate green.

7. Replace the left end cap.

Step 2: Power on the controllers

After you've powered on your shelves and assigned them unique IDs, power on the storage controllers.

Steps

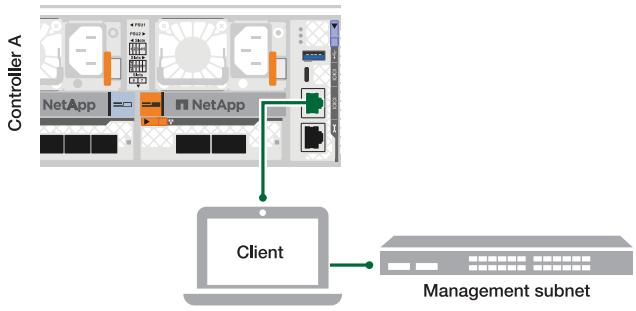
1. Connect your laptop to the serial console port. This will allow you to monitor the boot sequence when the controllers are powered on.
   a. Set the serial console port on the laptop to 115,200 baud with N-8-1.

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

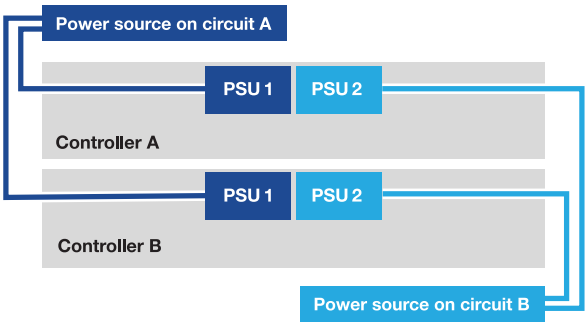
   b. Using the console cable provided with your storage system, connect one end of the console cable to

your laptop and the other end to the serial console port on controller A.

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



2. Assign a TCP/IP address to the laptop, using one that is on the management subnet.
3. Plug the two power cords into the controller power supplies, and then connect them to power sources on different circuits.



- The system begins to boot. Initial booting might take up to eight minutes.
- The LEDs flash on and the fans start, which indicates that the controllers are powering on.
- The fans might be very noisy when they first start up. The fan noise during start-up is normal.

4. Secure the power cords using the securing device on each power supply.

What's next?

After you've turned on your FAS70 or FAS90 storage system, you set up your cluster.

Maintain

Overview of the maintenance procedures - FAS70 and FAS90

Maintain the hardware of your FAS70 or FAS90 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the FAS70 or FAS90 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the FAS70 and FAS90 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot the image from a USB drive and restore the configuration from the partner node
Controller	A controller consists of a board, firmware, and software. It controls the drives and runs the ONTAP operating system software.
DIMM	A dual in-line memory module (DIMM) is a type of computer memory. They are installed to add system memory to a controller motherboard.
Fan	A fan cools the controller.
Flash Cache	Flash Cache speeds access to data through real-time intelligent caching of recently read user data and NetApp metadata. It's effective for random read-intensive workloads, including databases, email, and file services.
NVRAM	The NVRAM (Non-Volatile Random Access Memory) is a module that allows the controller to protect and save in-flight data if the system loses power. The system ID resides in the NVRAM module. When replaced, the controller assumes the new system ID from the replacement NVRAM module.
NV battery	The NV battery is responsible for providing power to the NVRAM module while data in-flight is being destaged to flash memory after a power loss.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real-time clock battery preserves system date and time information if the power is off.
System management module	The System management module provides the interface between the controller and a console or laptop for controller or system maintenance purposes. The System management module contains the boot media and stores the system serial number (SSN).

Boot media - automated recovery

Boot media automated recovery workflow - FAS70 and FAS90

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your FAS70 or FAS90 storage system.

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

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

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

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

3. **Replace the boot media**

Remove the failed boot media from the System Management module and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5. **Return the failed part to NetApp**

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

Requirements for automated boot media recovery - FAS70 and FAS90

Before replacing the boot media in your FAS70 or FAS90 storage system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming the cluster ports on the impaired controller are working properly, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

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

Review the following requirements.

- You must replace the failed boot media with a replacement boot media you received from NetApp.
- The cluster ports are used to communicate between the two controllers during the automated boot recovery process. Make sure that the cluster ports on the impaired controller are working properly.
- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg
  - /cfcard/kmip/certs/client.crt
  - /cfcard/kmip/certs/client.key
  - /cfcard/kmip/certs/CA.pem
- Understand the controller terminology used in this procedure:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - FAS70 and FAS90

Shut down the impaired controller in your FAS70 or FAS90 storage system to prevent data loss and maintain system stability during the automatic boot media recovery process.

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

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

About this task

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

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

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

Steps

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

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

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

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

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

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

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

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

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - FAS70 and FAS90

The boot media in your FAS70 or FAS90 storage system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media in the System Management module, and then reinstalling the System Management module.

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

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

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.

   If your storage system has DC power supplies, disconnect the power cable block from the power supply units (PSUs).

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

module.

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

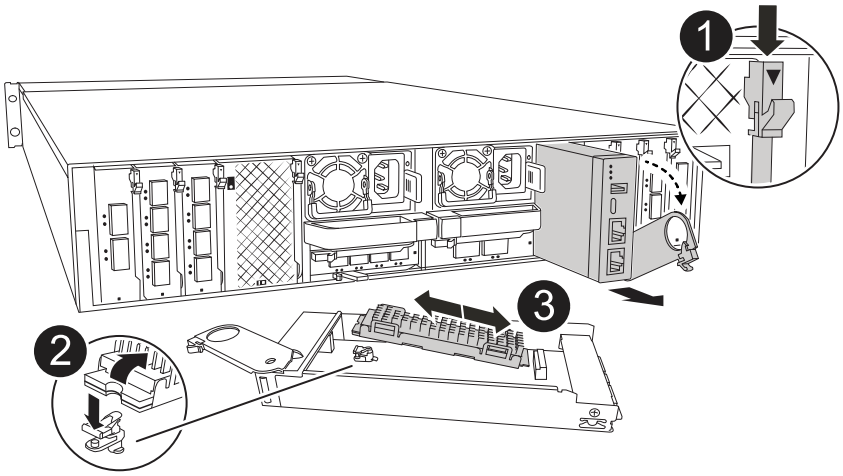
c. Depress the System Management cam button.

d. Rotate the cam latch down as far as it will go.

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

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

4. Remove the boot media from the management module:



1	System Management module cam latch
2	Boot media locking button
3	Boot media

a. Press the blue locking button.

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

5. Install the replacement boot media into the System Management module:

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

b. Rotate the boot media down toward the locking button.

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

6. Reinstall the System Management module:

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

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

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

a. Recable the System Management module.

8. Plug the power cables into the power supplies and reinstall the power cable retainer.

   The controller begins to boot as soon as power is reconnected to the system.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - FAS70 and FAS90

After installing the new boot media device in your FAS70 or FAS90 storage system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

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

Before you begin

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

Steps

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

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

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

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

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

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. a. Wait for the login prompt to display. b. Log into the node and give back the storage: `storage failover giveback -ofnode impaired_node_name` c. Go to re-enabling automatic giveback if it was disabled.
`key manager is configured.`	Encryption is installed. Go to restoring the key manager.

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

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

Onboard Key Manager (OKM)

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

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

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

a. Enter `y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the passphrase for onboard key management when prompted.

c. Enter the passphrase again when prompted to confirm.

d. Enter the backup data for onboard key manager when prompted.

Show example of passphrase and backup data prompts

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

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

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

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

Successfully recovered keymanager secrets.
```

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

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

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

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

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

Go to re-enabling automatic giveback if it was disabled.

External Key Manager (EKM)

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

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

a. Enter the EKM configuration settings when prompted:

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

Show example of client certificate contents

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

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

Show example of client key file contents

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

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

Show example of KMIP server file contents

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

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

Show example of server configuration file contents

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

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

Show example of ONTAP Cluster UUID prompt

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


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

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

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

Show example of temporary network setting prompts

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

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

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

b. Verify the key restoration status:

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

Show example of key recovery error and warning messages

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

Uptime: 11m32s
System halting...

LOADER-B>
```

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

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

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

Go to re-enabling automatic giveback if it was disabled.

5. If automatic giveback was disabled, reenable it:

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

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

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

What's next

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

Return the failed boot media part to NetApp - FAS70 and FAS90

If a component in your FAS70 or FAS90 storage system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - FAS70 and FAS90

The manual recovery of the boot image involves using a USB drive to reinstall ONTAP onto the FAS70 or FAS90 system's replacement boot media. You must download the appropriate ONTAP recovery image from the NetApp Support Site and copy it to a USB drive. This prepared USB drive is then used to perform the recovery and restore the system to operational status.

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

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



Review boot media replacement requirements

Review the requirements for replacing the boot media.



Check onboard encryption keys

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



Shut down the impaired controller

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



Replace the boot media

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



Boot the recovery image

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



Restore encryption

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



Return the failed part to NetApp

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

Requirements for manual boot media recovery - FAS70 and FAS90

Before replacing the boot media in your FAS70 or FAS90 system, ensure you meet the

necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

USB flash drive

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

File preparation

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

Component replacement

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

Controller identification

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

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

What's next?

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

Check encryption support for manual boot media recovery - FAS70 and FAS90

To ensure data security on your FAS70 or FAS90 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.

What's next?

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

Shut down the controller for manual boot media recovery - FAS70 and FAS90

Shut down the impaired controller in your FAS70 or FAS90 storage system to prevent data loss and maintain system stability during the manual boot media recovery process.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

What's next?

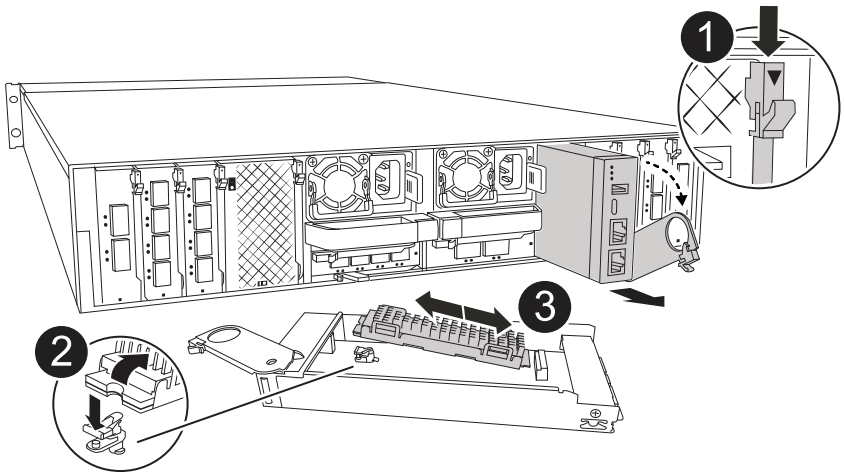
After shutting down the controller, you need to replace the boot media.

Replace the boot media and prepare for manual boot recovery - FAS70 and FAS90

The boot media in your FAS70 or FAS90 system stores essential firmware and configuration data. The replacement process involves removing the System Management module, removing the impaired boot media, installing the replacement boot media, and then manually transferring the ONTAP image to the replacement boot media using a USB flash drive.

Step 1: Replace the boot media

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



1	System Management module cam latch
2	Boot media locking button
3	Boot media

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.

   If your storage system has DC power supplies, disconnect the power cable block from the power supply units (PSUs).

3. Remove the System Management module:
   a. Remove any cables connected to the System Management module. Make sure that you label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
   b. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
   c. Depress the System Management cam button.
   d. Rotate the cam latch down as far as it will go.
   e. Remove the System Management module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
   f. Place the System Management module on an anti-static mat, so that the boot media is accessible.
4. Remove the boot media from the management module:
   a. Press the blue locking button.
   b. Rotate the boot media up, slide it out of the socket, and set it aside.
5. Install the replacement boot media into the System Management module:
   a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.
   b. Rotate the boot media down toward the locking button.
   c. Push the locking button, rotate the boot media all the way down and then release the locking button.
6. Reinstall the System Management module.
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
7. Rotate the cable management tray up to the closed position.
   a. Recable the System Management module.

Step 2: Transfer the ONTAP image to the boot media

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

Before you begin

- You must have an empty USB flash drive, formatted to FAT32, with at least 4GB capacity.
- Download a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site. Use the `version -v` command to display if your version of ONTAP supports NVE. If the command output displays

<10no- DARE>, your version of ONTAP does not support NVE.

- If NVE is supported by your version of ONTAP, download the image with NetApp Volume Encryption, as indicated in the download button.
- If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.

- If your system is an HA pair, you must have a network connection between the node management ports of the controllers (typically the e0M interfaces).

Steps

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

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

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

   c. Remove the USB flash drive from your laptop.
2. Insert the USB flash drive into the USB slot on the System Management module.

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

3. Plug the power cables back into the PSU.

   The system will begin to reboot, typically to the LOADER prompt.

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

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

What's next?

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

Manual boot media recovery from a USB drive - FAS70 and FAS90

After installing the new boot media device in your FAS70 or FAS90 system, you can boot the recovery image manually from a USB drive to restore the configuration from the partner node.

Before you begin

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

Steps

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

   ```
   boot_recovery
   ```

   The recovery image is downloaded from the USB flash drive.

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

ONTAP 9.16.0 or earlier

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

a. **On the impaired controller:** Press `Y` when you see `Do you want to restore the backup configuration now?`

b. **On the impaired controller:** If prompted, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. **On the partner controller:** Set the impaired controller to advanced privilege level:

```
set -privilege advanced
```

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

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

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

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

```
set -privilege admin
```

f. **On the impaired controller:** Press `Y` when you see `Was the restore backup procedure successful?`

g. **On the impaired controller:** Press `Y` when you see `...would you like to use this restored copy now?`

h. **On the impaired controller:** Press `Y` when prompted to reboot, then press `Ctrl-C` when you see the Boot Menu.

i. **On the impaired controller:** Do one of the following:

- If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.
- If the system uses encryption, go to Restore encryption.

ONTAP 9.16.1 or later

Complete the following steps on the impaired controller:

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

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

b. Press `Y` when prompted to confirm that the restore backup was successful.

c. Press `Y` when prompted to use the restored configuration.

d. Press `Y` when prompted to reboot the node.

e. Press `Y` when prompted to reboot again, then press `Ctrl-C` when you see the Boot Menu.

f. Do one of the following:

- If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.

- If the system uses encryption, go to Restore encryption.

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

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

6. If you disabled automatic giveback, reenable it:

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

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

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

What's next?

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

Restore encryption keys after manual boot recovery - FAS70 and FAS90

Restore encryption on the replacement boot media in your FAS70 or FAS90 system to ensure continued data protection. The replacement process involves verifying key availability, reapplying encryption settings, and confirming secure access to your data.

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---------------------------
012345678901234567890123456789012345678901234567890123456789012
23
123456789012345678901234567890123456789012345678901234567890123
34
234567890123456789012345678901234567890123456789012345678901234
45
345678901234567890123456789012345678901234567890123456789012345
56
456789012345678901234567890123456789012345678901234567890123456
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
```

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

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

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

   The recovery process completes and displays the following message:

   ```
   Successfully recovered keymanager secrets.
   ```

Show example prompt

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

Successfully recovered keymanager secrets.

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

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

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

Show example prompt

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


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

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

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

   On the partner controller:

8. Giveback the impaired controller:

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

   On the impaired controller:

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

   `security key-manager onboard sync`

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

Show example prompt

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

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

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

11. Verify that all keys are synced:

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

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

    On the partner controller:

12. Giveback the impaired controller:

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

13. Restore automatic giveback if you disabled it:

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

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

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

External Key Manager (EKM)

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

Before you begin

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

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

Steps

On the impaired controller:

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

   Show example boot menu

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

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

   Show example prompt

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

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

Show example

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

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

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

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

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

The recovery process completes and displays the following message:

```
Successfully recovered keymanager secrets.
```

Show example

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

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

Show example prompt

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

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

6. Restore automatic giveback if you disabled it:

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

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

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

What's next?

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

Return the failed part to NetApp - FAS70 and FAS90

If a component in your FAS70 or FAS90 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Chassis

Chassis replacement workflow - FAS70 and FAS90

Get started with replacing the chassis of your FAS70 or FAS90 storage system by reviewing the replacement requirements, shutting down the controller, replacing the

chassis, and verifying system operations.

1 Review the chassis replace requirements

Review the chassis replacement requirements.

2 Prepare for chassis replace

Prepare to replace the chassis by locating the system, gathering system credentials and necessary tools, verifying the replacement chassis was received, and labeling the system cables.

3 Shut down the controller

Shut down the controller so you can perform maintenance on the chassis.

4 Replace the chassis

Replace the chassis by moving the components from the impaired chassis to the replacement chassis.

5 Complete the chassis replacement

Complete the chassis replacement by bringing the controller up, giving back the controller, and returning the failed chassis to NetApp.

Requirements to replace the chassis - FAS70 and FAS90

Before replacing the chassis in your FAS70 or FAS90 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have local administrator credentials for ONTAP, the correct replacement chassis, and the necessary tools.

The chassis is the physical chassis housing all the system components such as the fans, controller/CPU unit, NVRAM12, system management module, I/O cards and blanking modules, and PSUs.

Review the following requirements.

- Make sure all other components in the system are functioning properly; if not, contact NetApp support for assistance.
- Obtain local administrator credentials for ONTAP if you don’t have them.
- Make sure that you have the necessary tools and equipment for the replacement.
- You can use the chassis replacement procedure with all versions of ONTAP supported by your system.
- The chassis replacement procedure is written with the assumption that you are moving the bezel, fans, controller module, NVRAM12, system management module, I/O cards and blanking modules, and PSUs to the new chassis, and that the replacement chassis is a new component from NetApp.

What’s next?

After you've reviewed the requirements to replace the chassis, you need to prepare to replace the chassis.

Prepare to replace the chassis - FAS70 and FAS90

Prepare to replace the impaired chassis in your FAS70 or FAS90 system by identifying the impaired chassis, verifying the replacement components, and labeling the cables and controller module.

Step 1: Locate and monitor your system

You should open a console session and save sessions logs for future reference, and also turn on the system location LED to find the impaired chassis.

Steps

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

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

      The Location LED remains lit for 30 minutes.

Step 2: Verify replacement components

You should verify that you received the necessary components, remove them from packaging, and save the packaging.

Steps

1. Before opening the packaging, you should look at the packaging label and verify:
   - Component part number.
   - Part description.
   - Quantity in the box.
2. Remove the contents from the packaging and use the packaging to returning the failed component to NetApp.

Step 3: Label the cables

You should label the cables before removing them from the I/O modules on the back of the system.

Steps

1. Label all the cables associated with the storage system. This aids recabling later in this procedure.
2. If you are not already properly grounded, ground yourself.

What's next?

After you've prepared to replace your FAS70 or FAS90 chassis hardware, you need to shut down the controller.

Shut down the controller in your FAS70 or FAS90 storage system to prevent data loss and ensure system stability when replacing the chassis.

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

About this task

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

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

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

Steps

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

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

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

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

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

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

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

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

What's next?

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

Replace the chassis - FAS70 and FAS90

Replace the chassis of your FAS70 or FAS90 system when a hardware failure requires it. The replacement process involves removing the controller, I/O cards, NVRAM12 module, system management module, and power supply units (PSUs), installing the replacement chassis, and reinstalling the chassis components.

Step 1: Remove the PSUs and cables

You need to remove the two power supply units (PSUs) before removing the controller.

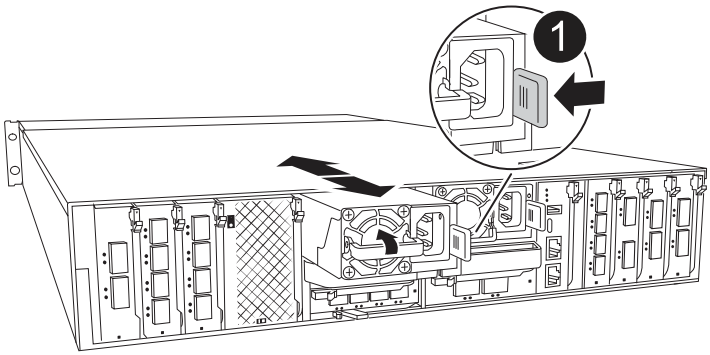
Steps

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

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

   c. Remove the two PSUs from the rear of the chassis by rotating the PSU handle up so that you can pull the PSU out, press the PSU locking tab, and then pull PSU out of the chassis.

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



1	Terracotta PSU locking tab

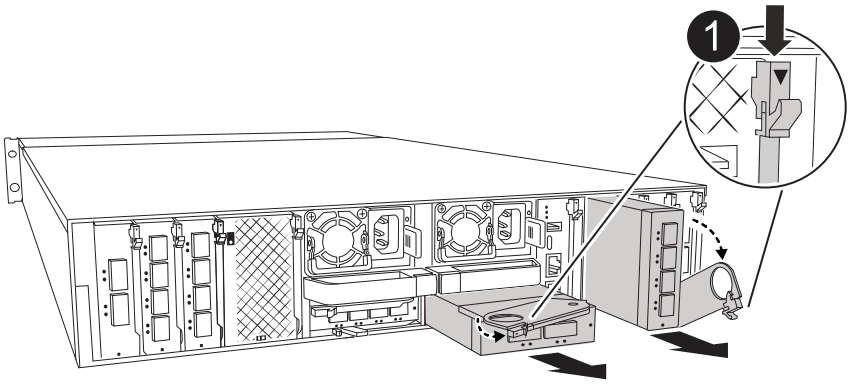
   d. Repeat these steps for the second PSU.
2. Remove the cables:
   a. Unplug the system cables and any SFP and QSFP modules (if needed) from the controller module, but leave them in the cable management device to keep them organized.

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

b. Remove the cable management device from the chassis and set it aside.

Step 2: Remove the I/O cards, NVRAM12, and system management module

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



1	I/O cam latch

a. Depress the cam button on the target module.

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

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

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

d. Set the I/O module aside and repeat these steps for any other I/O modules.

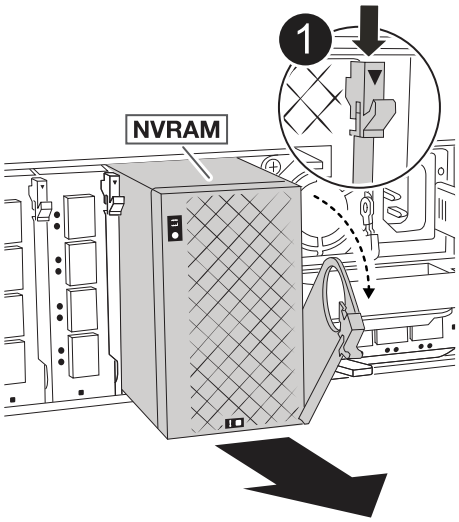
2. Remove the NVRAM12 module:

   a. Depress the locking cam button.

      The cam button moves away from the chassis.

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

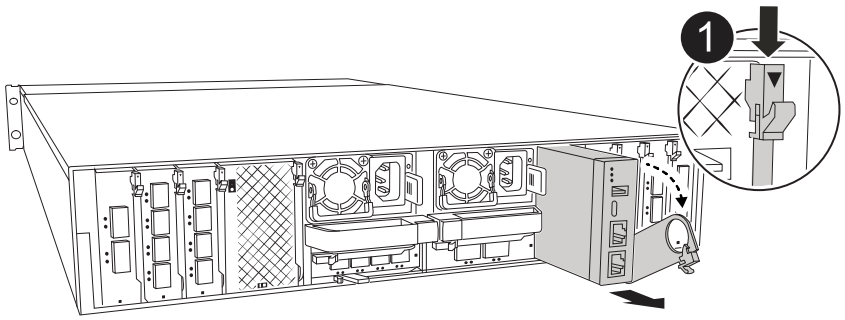
   c. Remove the NVRAM module from the chassis by hooking your finger into the cam lever opening and pulling the module out of the chassis.



1	NVRAM12 cam latch

d. Set the NVRAM module on a stable surface.

3. Remove the system management module:
   a. Depress the cam button on the System Management module.
   b. Rotate the cam lever down as far as it will go.
   c. Loop your finger into the hole on the cam lever and pull the module straight out of the system.

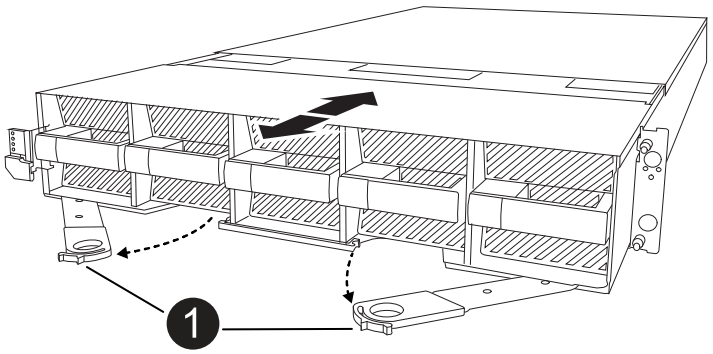


1	System Management module cam latch

Step 3: Remove the controller module

1. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the chassis.



1	Locking cam latches

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

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

Step 4: Replace the impaired chassis

Remove the impaired chassis and install the replacement chassis.

Steps

1. Remove the impaired chassis:
   a. Remove the screws from the chassis mount points.
   b. Slide the impaired chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
2. Install the replacement chassis:
   a. Install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
   b. Slide the chassis all the way into the equipment rack or system cabinet.
   c. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the impaired chassis.

Step 5: Install the chassis components

After the replacement chassis is installed, you need to install the controller module, recable the I/O modules and system management module, and then reinstall and plug in the PSUs.

Steps

1. Install the controller module:
   a. Align the end of the controller module with the opening in the front of the chassis, and then gently push the controller all the way into the chassis.
   b. Rotate the locking latches into the locked position.
2. Install the I/O cards at the rear of the chassis:
   a. Align the end of the I/O module with the same slot in the replacement chassis as in the impaired chassis, and then gently push the module all the way into the chassis.

b. Rotate the cam latch upward into the locked position.

c. Repeat these steps for any other I/O modules.

3. Install the system management module at the rear of the chassis:

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

   b. Rotate the cam latch upward into the locked position.

   c. If you have not already done so, reinstall the cable management device and reconnect the cables to the I/O cards and system management module.

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

      Make sure that the cables are connected according to the cable labels.

4. Install the NVRAM12 module in the back of the chassis at the rear of the chassis:

   a. Align the end of the NVRAM12 module with the opening in the chassis, and then gently push the module all the way into the chassis.

   b. Rotate the cam latch upward into the locked position.

5. Install the PSUs:

   a. Using both hands, support and align the edges of the PSU with the opening in the chassis.

   b. Gently push the PSU into the chassis until the locking tab clicks into place.

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

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

6. Reconnect the PSU power cables to both of the PSUs and secure each power cable to the PSU using the power cable retainer.

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

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

What's next?

After you've replaced the impaired FAS70 and FAS90 chassis and reinstalled the components into it, you need to complete the chassis replacement.

Complete the chassis replacement - FAS70 and FAS90

Reboot the controller, verify system health, and return the failed part to NetApp to complete the final step in the FAS70 and FAS90 chassis replacement procedure.

Step 1: Boot the controllers and verify system health

After the controllers reboot, boot ONTAP, give back the controllers, and verify the storage system health.

Steps

1. Check the console output:
   a. If the controller boots to the LOADER prompt, reboot the controller with the `boot_ontap` command.
   b. If the console displays `waiting for giveback` after the reboot, log into the partner controller and check that the replaced controller is ready for giveback with the `storage failover show` command.
2. Perform the giveback:
   a. Connect the console cable to the partner controller.
   b. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
   c. If automatic giveback was disabled, reenable it: `storage failover modify -node impaired_node_name -auto-giveback true`
   d. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`
3. After the giveback is complete, run Active IQ Config Advisor to verify the health of the storage system and correct any issues you encounter.

Step 2: Return the failed part to NetApp

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

Controller

Controller replacement workflow - FAS70 and FAS90

Get started with replacing the controller in your FAS70 or FAS90 storage system by shutting down the impaired controller, removing and replacing the controller, restoring the system configuration, and returning control of storage resources to the replacement controller.



Review controller replacement requirements

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



Shut down the impaired controller

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



Replace the controller

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



Restore and verify the system configuration

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

5 Give back the controller

Transfer the ownership of storage resources back to the replacement controller.

6 Complete controller replacement

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

Requirements to replace the controller - FAS70 and FAS90

Before replacing the controller in your FAS70 or FAS90 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying all other components in the system are functioning properly, verifying that you have the correct replacement controller, and saving the controller's console output to a text log file.

Review the requirements for replacing the controller.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the "impaired controller").
- Do not use this procedure for controller upgrades; instead, refer to the Choose your controller hardware upgrade procedure for guidance.
- 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 controller replacement procedure.
- You must replace the failed component with the field-replaceable unit (FRU) you received from NetApp.
- You must replace 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.
- Because the boot device is located on the System Management module that is installed in the back of the system, you do not need to move the boot device when replacing a 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 log file.

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

What's next?

After you've reviewed the requirements to replace your FAS70 or FAS90 controller, you need to shut down the impaired controller.

Shut down the impaired controller - FAS70 and FAS90

Shut down the controller in your FAS70 or FAS90 storage system to prevent data loss and ensure system stability when replacing the controller.

Option 1: Most systems

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

About this task

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

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

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

Steps

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

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

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

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

2. Disable automatic giveback:

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

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

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

3. Take the impaired controller to the LOADER prompt:

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

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

Option 2: Controller is in a MetroCluster

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

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

  `metrocluster node show`

Steps

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

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

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

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

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

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

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

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

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

What's next?

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

Replace the controller - FAS70 and FAS90

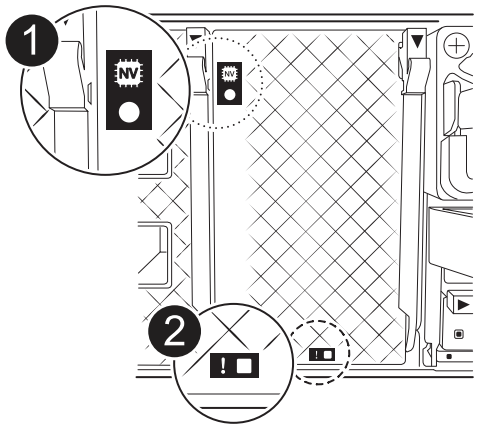
Replace the controller in your FAS70 or FAS90 system when a hardware failure requires it. The replacement process involves removing the impaired controller, moving the components to the replacement controller, installing the replacement controller, and rebooting it.

Step 1: Remove the controller module

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

Steps

1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:

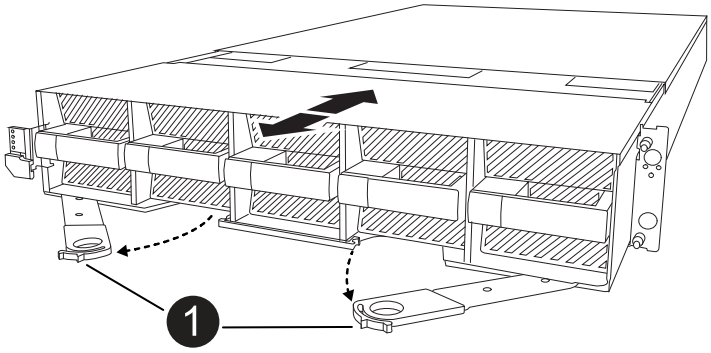


1	NVRAM status LED
2	NVRAM attention LED

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

2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.



1	Locking cam latches

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

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

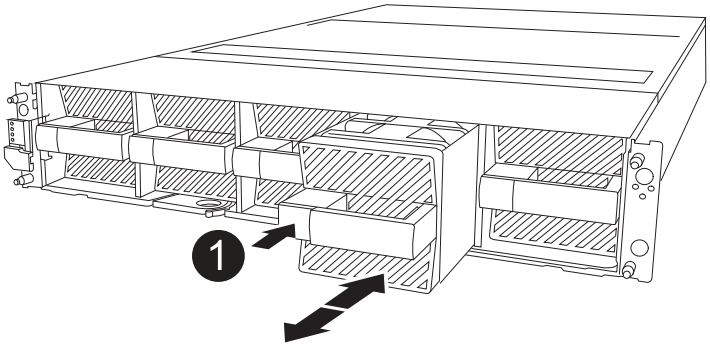
Step 2: Move the fans

You must remove the five fan modules from the impaired controller module to the replacement controller module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Press the gray locking button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

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



1	Black locking button

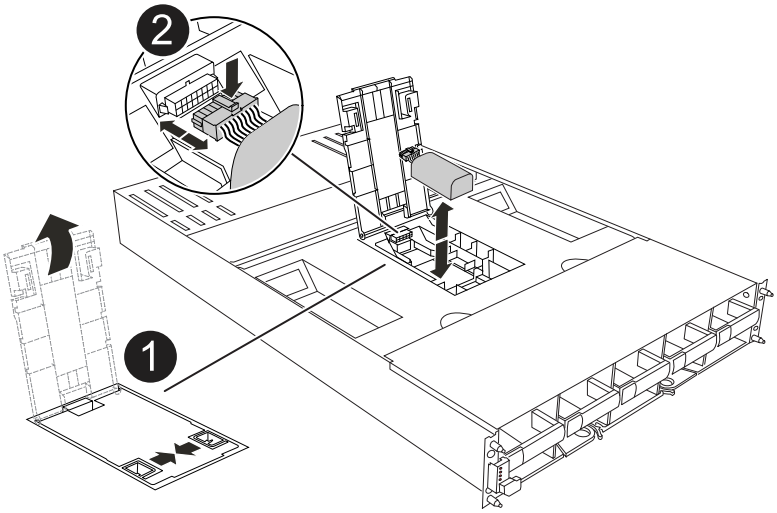
4. Install the fan in the replacement controller module:
   a. Align the edges of the fan housing with the opening in the front of the replacement controller module.
   b. Gently slide the fan module all the way into the replacement controller module until it locks in place.
5. Repeat the preceding steps for the remaining fan modules.

Step 3: Move the NV battery

Move the NV battery to the replacement controller.

Steps

1. Open the NV battery air duct cover and locate the NV battery.



1	NV battery air duct cover
2	NV battery plug
3	NV battery pack

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

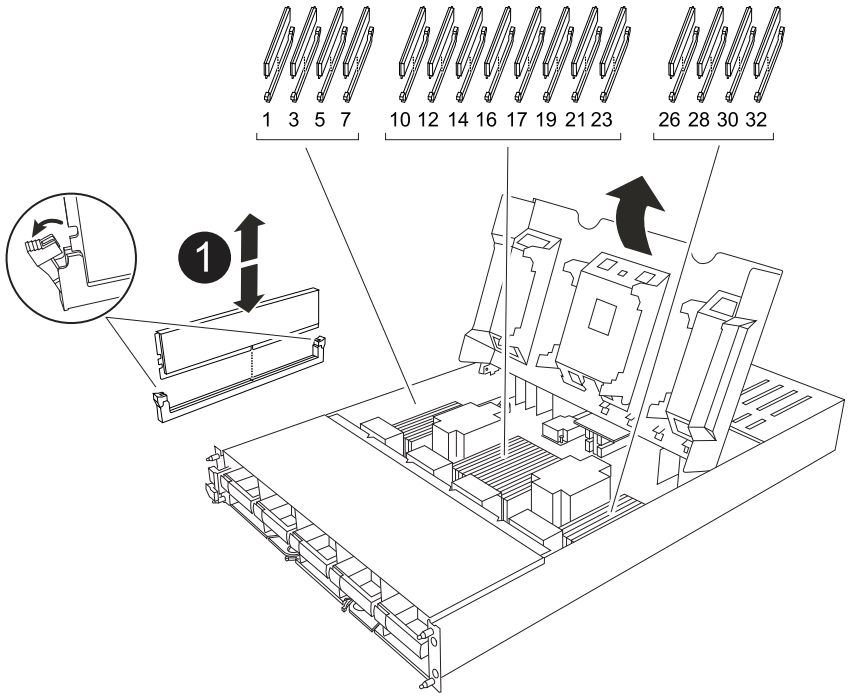
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module.
5. Move the battery pack to the replacement controller module and then install it in the NV battery air duct:
   a. Open the NV battery air duct in the replacement controller module.
   b. Plug the battery plug into the socket and make sure that the plug locks into place.
   c. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
   d. Close the air duct cover.

Step 4: Move system DIMMs

Move the DIMMs to the replacement controller module.

Steps

1. Open the motherboard air duct and locate the DIMMs.



1	System DIMM

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.

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

6. 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.
7. Repeat these steps for the remaining DIMMs.
   Close the motherboard air duct.

Step 5: Install the controller module

Reinstall the controller module and boot it.

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

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans

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

   The controller boots to the LOADER prompt as soon as it is fully seated.

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

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

5. Set the current time in GMT with the `set time hh:mm:ss` command. You can get the current GMT from the partner node with the `date -u`command.
6. Recable the storage system, as needed.

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

What's next?

After you've replaced the impaired FAS70 or FAS90 controller, you need to restore the system configuration.

Restore and verify the system configuration - FAS70 and FAS90

Verify that the controller's HA configuration is active and functioning correctly in your FAS70 or FAS90 storage system, and confirm that the system's adapters list all the paths

to the disks.

Step 1: Verify HA config settings

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

Steps

1. Boot to maintenance mode: `boot_ontap maint`
   a. Enter `y` when you see *Continue with boot?*.

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

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

   If you see *PERSONALITY MISMATCH* contact customer support.

3. From the `sysconfig -v` output, compare the adapter card information with the cards and locations in the replacement controller.
4. Verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

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

   The value for the HA state can be one of the following:

   - `ha`
   - `mcc` (not supported)
   - `mccip` (not supported in ASA systems)
   - `non-ha` (not supported)

6. Confirm that the setting has changed: `ha-config show`

Step 2: Verify disk list

Steps

1. Verify that the adapter lists the paths to all disks with the `storage show disk -p`.

   If you see any issues, check cabling and reseat cables.

2. Exit Maintenance mode: `halt`.

What's next?

After you've restored and verified the system configuration for your FAS70 or FAS90 system, you need to give back the controller.

Return control of storage resources to the replacement controller so your FAS70 or FAS90 system can resume normal operation. The give back procedure varies based on the encryption type used by your system: no encryption or Onboard Key Manager (OKM) encryption.

No encryption

Return the impaired controller to normal operation by giving back its storage.

Steps

1. From the LOADER prompt, enter `boot_ontap`.
2. Press <enter> when console messages stop.
   - If you see the *login* prompt, go to the next step at the end of this section.
   - If you see *Waiting for giveback*, press the <enter> key, log into the partner node, and then go to the next step at the end of this section.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

Onboard encryption (OKM)

Reset onboard encryption and return the controller to normal operation.

Steps

1. From the LOADER prompt, enter `boot_ontap maint`.
2. Boot to the ONTAP menu from the LOADER prompt `boot_ontap menu` and select option 10.
3. Enter the OKM passphrase.

   You are prompted twice for the passphrase.

4. Enter the backup key data when prompted.
5. At the boot menu, enter option `1` for normal boot.
6. Press <enter> when *Waiting for giveback* is displayed.
7. Move the console cable to the partner node and login as `admin`.
8. Give back only the CFO aggregates (the root aggregate): `storage failover giveback -fromnode local -only-cfo-aggregates true`
   - If you encounter errors, contact NetApp Support.
9. Wait 5 minutes after the giveback report completes, and check failover status and giveback status: `storage failover show` and `storage failover show-giveback`.
10. Synchronize and verify status of the keys:
    a. Move the console cable back to the replacement controller.
    b. Synchronize missing keys: `security key-manager onboard sync`

       You are prompted for the cluster-wide passphrase of OKM for the cluster.

    c. Verify status of the keys: `security key-manager key query -restored false`

The output should show no results when when properly synchronized.

If the output shows results (the key IDs of keys that are not present in the system's internal key table), contact NetApp Support.

11. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode *impaired_node_name*`
12. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback *true*`
13. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`

What's next?

After you've transferred the ownership of storage resources back to the replacement controller, you need to complete the controller replacement procedure.

Complete controller replacement - FAS70 and FAS90

To complete the controller replacement for your AFF A1K system, first restore the NetApp Storage Encryption configuration (if necessary). Next, confirm that the logical interfaces (LIFs) are reporting to their home ports and perform a cluster health check. Finally, return the failed part to NetApp.

Step 1: Verify LIFs and check cluster health

Before returning the replacement node to service, verify that the logical interfaces are on their home ports, check the cluster health, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports:

   ```
   network interface show -is-home false
   ```

   If any logical interfaces are listed as false, revert them to their home ports:

   ```
   network interface revert -vserver * -lif *
   ```

2. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article.
3. If automatic giveback was disabled, reenable it:

   ```
   storage failover modify -node local -auto-giveback true
   ```

Step 2: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - FAS70 and FAS90

Replace a DIMM in your FAS70 or FAS90 system if excessive correctable or uncorrectable memory errors are detected. Such errors can prevent the storage system from booting ONTAP. The replacement process involves shutting down the impaired controller, removing it, replacing the DIMM, reinstalling the controller, and then returning the failed part to NetApp.

Before you begin

- Make sure all other components in the system are functioning properly; if not, you must contact technical support.
- Make sure you replace the failed component with a replacement component you received from NetApp.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

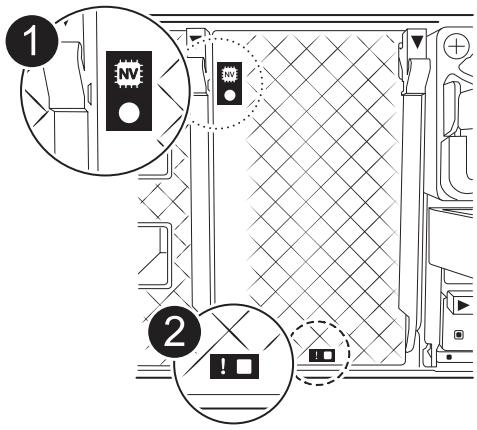
If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the enclosure when you replace the controller module or replace a component inside the controller module.

Steps

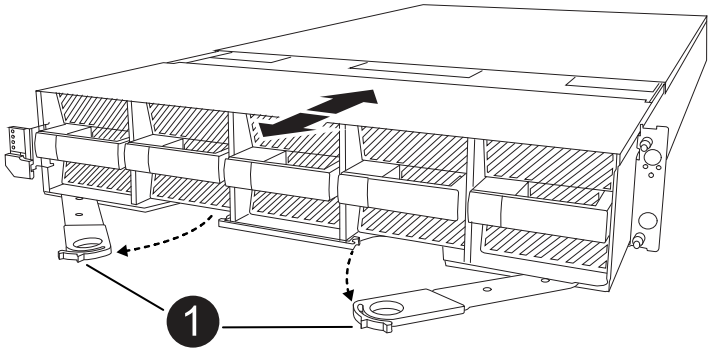
1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:



1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.
2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.



1	Locking cam latches

4. Slide the controller module out of the enclosure and place it on a flat, stable surface.

   Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

Step 3: Replace a DIMM

You must replace a DIMM when the system reports a permanent failure condition for that DIMM.

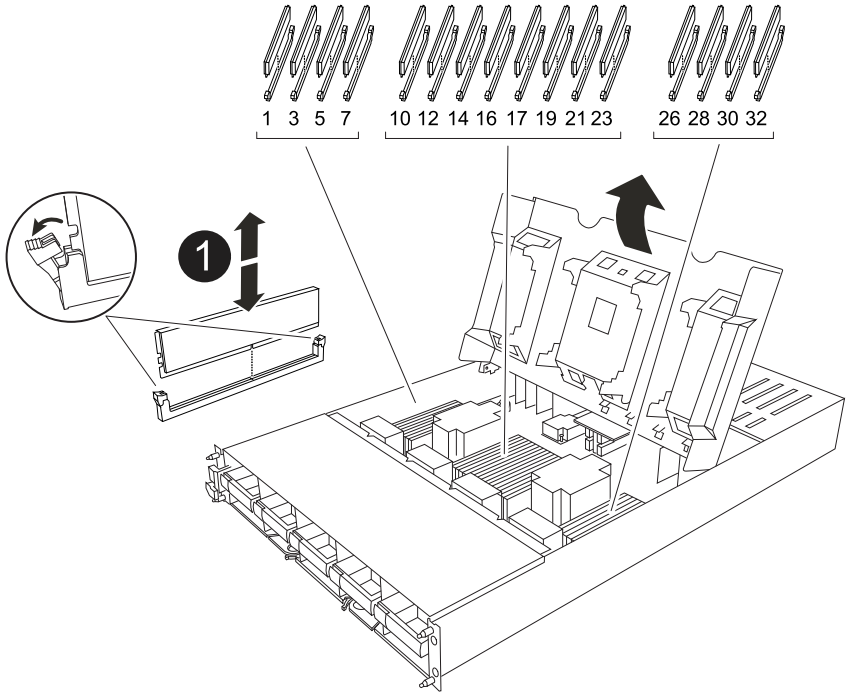
1. If you are not already grounded, properly ground yourself.
2. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
3. Locate the DIMMs on your controller module and identify the DIMM for replacement.

   DIMMs locations are dependent on the system model:

Model	DIMM slot location
FAS70	Slots 3, 10, 19, 26
FAS90	Slots 3, 7, 10, 14, 19, 23, 26, 30

4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.



1	DIMM and DIMM ejector tabs

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the controller air duct.

Step 4: Install the controller

Reinstall the controller module and boot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch

back under the fans

> Do not use excessive force when sliding the controller module into the enclosure to avoid damaging the connectors.

The controller module begins to boot as soon as it is fully seated in the enclosure.

4. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
6. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a fan - FAS70 and FAS90

Replace a fan module in your FAS70 or FAS90 system when a fan fails or is not operating efficiently, as this can affect system cooling and overall performance. The replacement process involves shutting down the controller, removing the controller, replacing the fan, reinstalling the controller, and returning the failed part to NetApp.

Steps

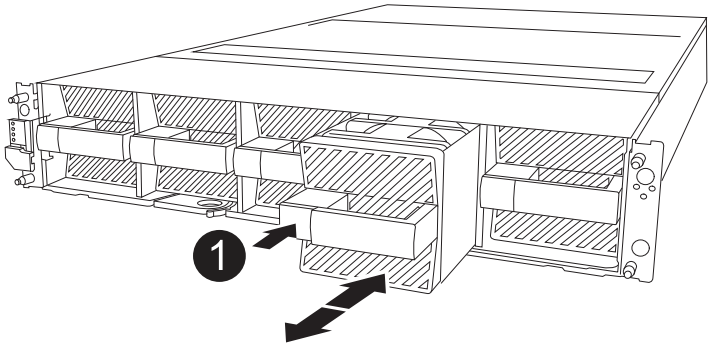
1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.

   Facing the controller module, fan modules are numbered 1 through 5, from left to right.

   > There is a single LED for each fan. It is green when the fan is functioning correctly and amber when not.

4. Press the black button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

   > The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.



1	Black release button

5. Set the fan module aside.
6. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

   When inserted into a live system, the amber Attention LED turns off once the fan is recognized by that system.

7. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the Flash Cache module carrier or a caching module - FAS70 and FAS90

The NVMe SSD Flash Cache module carrier in your FAS70 or FAS90 system contains one or two Flash Cache modules (caching modules) with a single SSD Flash Cache drive integrated into each caching module.

The FAS70 supports 2TB caching modules and FAS90 supports 4TB caching modules. You cannot mix caching modules of different capacity in the Flash Cache module carrier.

You can perform either of the following procedures depending on what component you need to replace: the entire Flash Cache module carrier or a caching module.

- Replace the Flash Cache module carrier
- Replace the caching module

Replace the Flash Cache module carrier

The Flash Cache module carrier is located in slot 6 and houses up to two Flash Cache modules. You cannot hot-swap the Flash Cache module carrier

Before you begin

- Ensure your storage system has the appropriate operating system for the replacement Flash Cache module carrier.
- Confirm all other components are functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired node

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

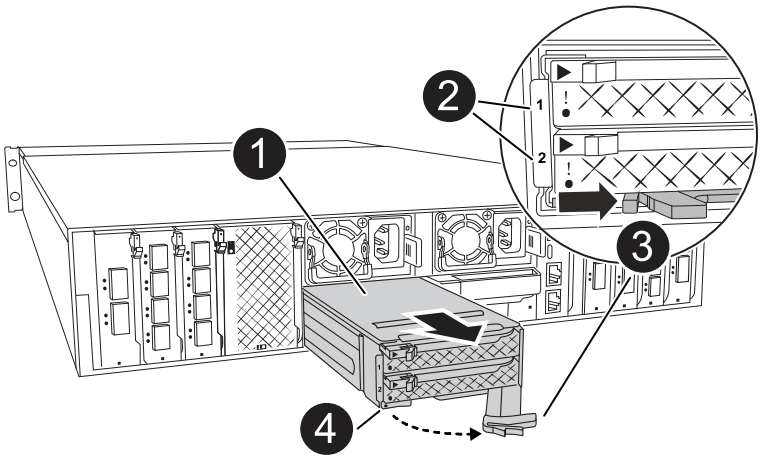
If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace the Flash Cache module carrier

Perform the following steps to replace the Flash Cache module carrier.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the failed Flash Cache module carrier, in slot 6, by the lit amber Attention LED on the front of the Flash Cache module carrier.



1	Flash Cache module carrier
2	Caching module slot numbers
3	Flash Cache module carrier cam handle
4	Flash Cache module carrier fault LED

3. Remove the failed Flash Cache module carrier:
   a. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.

b. Pinch the blue tab at the bottom of the Flash Cache module carrier.

c. Rotate the tab away from the module.

4. Pull the Flash Cache module carrier out of the controller module and set it on an antistatic mat.
5. Move the caching modules to the replacement Flash Cache module carrier:
   a. Pinch the Terra Cotta tab at the top of the caching module and rotate the cam handle away from the caching module.
   b. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the Flash Cache module carrier.
   c. Install the caching module into the same slot in the replacement Flash Cache module carrier and rotate the cam handle to the closed position on the caching module to lock it in place.
6. Repeat these steps if there is a second caching module.
7. Install the replacement Flash Cache module carrier into the system:
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
   c. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller

After you replace the Flash Cache module carrier, you must reboot the controller module.

Steps

1. From the LOADER prompt, reboot the node: *bye*

   This reinitializes the I/O cards and other components and reboots the node.

2. Return the node to normal operation: *storage failover giveback -ofnode impaired_node_name*
3. If automatic giveback was disabled, reenable it: *storage failover modify -node local -auto-giveback true*

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the caching module

The Flash Cache modules (caching modules) are located in slot 6-1 or in slot 6-2 or in both slot 6-1 and slot 6-2.

You can hot-swap the individual caching modules with caching modules of the same capacity from the same vendor or from a different supported vendor.

Before you begin

- Ensure the replacement caching module has the same capacity as the failed one, from the same vendor or from a different supported vendor.
- Confirm all other components are functioning properly; if not, you must contact technical support.
- The drives in the caching modules are not field replaceable units (FRU). You must replace the entire

caching module.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the failed caching module, in slot 6, by the lit amber Attention LED on the front of the caching module.
3. Prepare the caching module slot for replacement as follows:

   a. Record the caching module capacity, part number, and serial number on the target node: *system node run local sysconfig -av 6*

   b. In admin privilege level, prepare the target caching module slot for removal, responding `y` when prompted whether to continue: *system controller slot module remove -node node_name -slot slot_number* The following command prepares slot 6-1 on node1 for removal, and displays a message that it is safe to remove:

   ```
   ::> system controller slot module remove -node node1 -slot 6-1

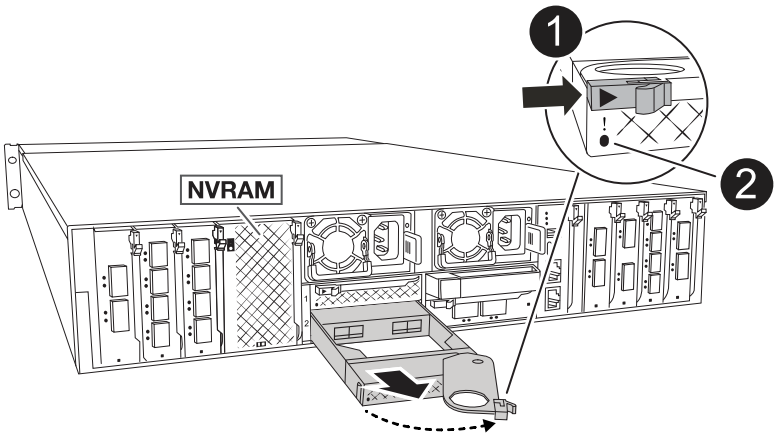
   Warning: SSD module in slot 6-1 of the node node1 will be powered off
   for removal.
   Do you want to continue? (y|n): _y_
   The module has been successfully removed from service and powered
   off. It can now be safely removed.
   ```

   c. Display the slot status with the `system controller slot module show` command.

   The caching module slot status displays `powered-off` in the screen output for the caching module that needs replacing.

   See the Command man pages for your version of ONTAP for more details.

4. Remove the caching module:



1	Caching module cam handle

2	Caching module fault LED

a. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.

b. Press the terra cotta release button on the front of the caching module.

c. Rotate the cam handle as far as it will go.

d. Remove the caching module module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the Flash Cache module carrier.

   Be sure to support the caching module as you remove it from the Flash Cache module carrier.

5. Install the replacement caching module:

   a. Align the edges of the caching module with the opening in the controller module.

   b. Gently push the caching module into the bay until the cam handle engages.

   c. Rotate the cam handle until it locks into place.

   d. Rotate the cable management tray up to the closed position.

6. Bring the replacement caching module online by using the `system controller slot module insert` command as follows:

   The following command prepares slot 6-1 on node1 for power-on, and displays a message that it is powered on:

```
::> system controller slot module insert -node node1 -slot 6-1

Warning: NVMe module in slot 6-1 of the node localhost will be powered
on and initialized.
Do you want to continue? (y|n): `y`

The module has been successfully powered on, initialized and placed into
service.
```

7. Verify the slot status using the `system controller slot module show` command.

   Make sure that command output reports status for the as `powered-on` and ready for operation.

8. Verify that the replacement caching module is online and recognized, and then visually confirm that the amber attention LED is not lit: `sysconfig -av slot_number`

   If you replace the caching module with a caching module from a different vendor, the new vendor name is displayed in the command output.

9. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace NVRAM - FAS70 and FAS90

Replace the NVRAM in your FAS70 or FAS90 system when the non-volatile memory becomes faulty or requires an upgrade. The replacement process involves shutting down the impaired controller, replacing the NVRAM module or the NVRAM DIMM, reassigning the disks, and returning the failed part to NetApp.

The NVRAM module consists of the NVRAM12 hardware and field-replaceable DIMMs. You can replace a failed NVRAM module or the DIMMs inside the NVRAM module.

Before you begin

- Make sure you have the replacement part available. You must replace the failed component with a replacement component you received from NetApp.
- Make sure all other components in the storage system are functioning properly; if not, contact NetApp support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Replace the NVRAM module or NVRAM DIMM

Replace the NVRAM module or NVRAM DIMMs using the appropriate following option.

Option 1: Replace the NVRAM module

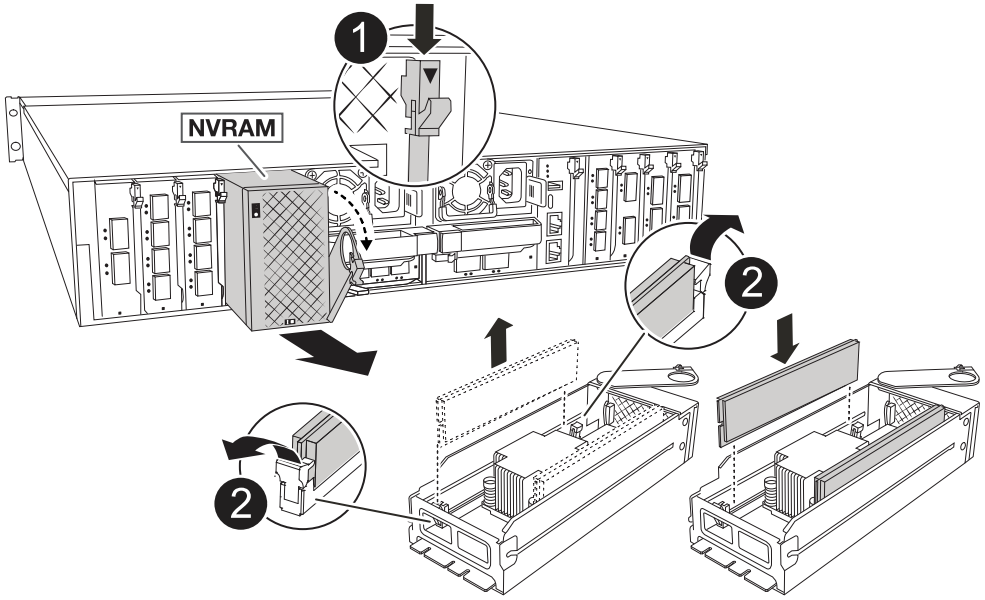
To replace the NVRAM module, locate it in slot 4/5 in the enclosure and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.
3. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.
4. Remove the impaired NVRAM module from the enclosure:
   a. Depress the locking cam button.

      The cam button moves away from the enclosure.

   b. Rotate the cam latch down as far as it will go.
   c. Remove the impaired NVRAM module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.



1	Cam locking button
2	DIMM locking tabs

5. Set the NVRAM module on a stable surface.
6. Remove the DIMMs, one at a time, from the impaired NVRAM module and install them in the replacement NVRAM module.
7. Install the replacement NVRAM module into the enclosure:
   a. Align the module with the edges of the enclosure opening in slot 4/5.

b. Gently slide the module into the slot all the way, and then rotate the cam latch all the way up to lock the module in place.

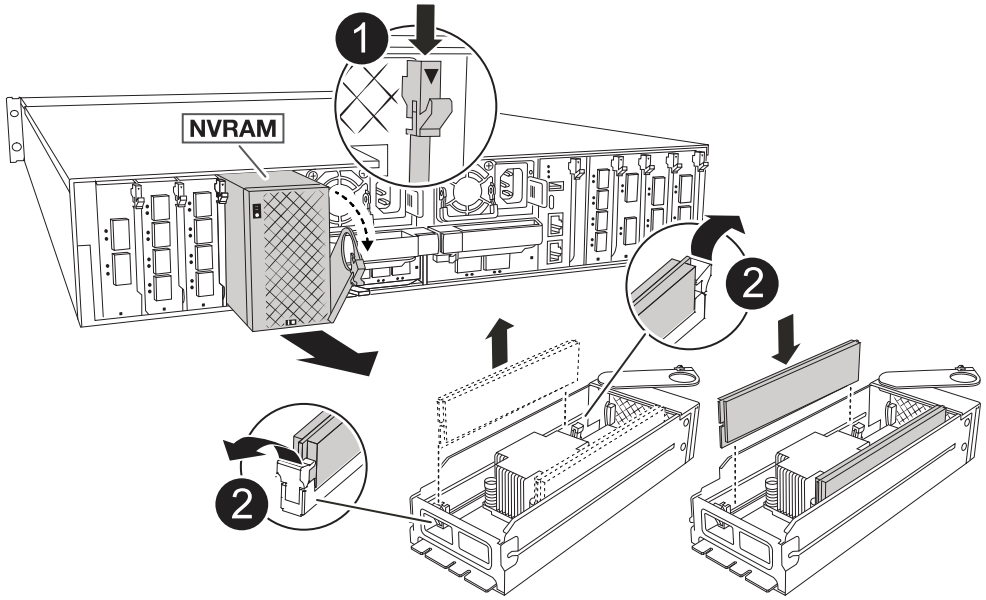
8. Recable the controller.
9. Rotate the cable management tray up to the closed position.

Option 2: Replace the NVRAM DIMM

To replace NVRAM DIMMs in the NVRAM module, you must remove the NVRAM module, and then replace the target DIMM.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the power supply cables from the PSUs.
3. Rotate the cable management tray down by gently pulling the pins on the ends of the tray and rotating the tray down.
4. Remove the target NVRAM module from the enclosure.



1	Cam locking button
2	DIMM locking tabs

5. Set the NVRAM module on a stable surface.
6. Locate the DIMM to be replaced inside the NVRAM module.

Consult the FRU map label on the side of the NVRAM module to determine the locations of DIMM slots 1 and 2.

7. Remove the DIMM by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.
8. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.

9. Install the NVRAM module into the enclosure:
    a. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.
10. Recable the controller.
11. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller

After you replace the FRU, you must reboot the controller module by plugging the power cables back into the PSU.

Steps

1. Plug the power cables back into the PSU.

   The system will begin to reboot, typically to the LOADER prompt.

2. Enter *bye* at the LOADER prompt.
3. Return the impaired controller to normal operation by giving back its storage: *`storage failover giveback -ofnode _impaired_node_name`*.
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 4: Reassign disks

You must confirm the system ID change when you boot the controller and then verify that the change was implemented.

> Disk reassignment is only needed when replacing the NVRAM module and does not apply to NVRAM DIMM replacement.

Steps

1. If the controller is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode and go to the LOADER prompt: *halt*
2. From the LOADER prompt on the controller, boot the controller and enter *y* when prompted to override the system ID due to a system ID mismatch.
3. Wait until the Waiting for giveback… message is displayed on the console of the controller with the replacement module and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: *storage failover show*

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, 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. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: *storage failover giveback -ofnode replacement_node_name*

      The controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter *y*.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      For more information, see the Manual giveback commands topic to override the veto.

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: *storage failover show*

      The output from the `storage failover show` command should not include the System ID changed on partner message.

5. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 151759706:

```
node1:> storage disk show -ownership

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        151759706  151759706  -
151759706 Pool0
1.0.1  aggr0_1  node1 node1           151759706  151759706  -
151759706 Pool0
.
.
.
```

6. If the system is in a MetroCluster configuration, monitor the status of the controller: *metrocluster node show*

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the impaired system ID until the MetroCluster configuration returns to a normal state.

7. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The controller is the current owner of the disks on the disaster site.

     See Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration for more information.

8. If your system is in a MetroCluster configuration, verify that each controller is configured: *metrocluster node show - fields configuration-state*

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
--------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

9. Verify that the expected volumes are present for each controller: `vol show -node node-name`
10. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`.
11. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
12. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NV battery - FAS70 and FAS90

Replace the NV battery in your FAS70 or FAS90 system when the battery begins to lose charge or fails, as it is responsible for preserving critical system data during power outages. The replacement process involves shutting down the impaired controller, removing the controller module, replacing the NV battery, reinstalling the controller module, and returning the failed part to NetApp.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

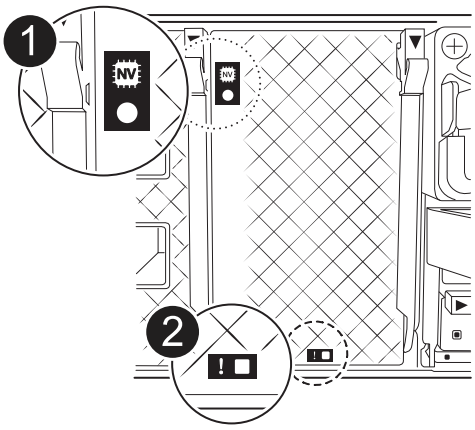
If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the enclosure when you replace the controller module or replace a component inside the controller module.

Steps

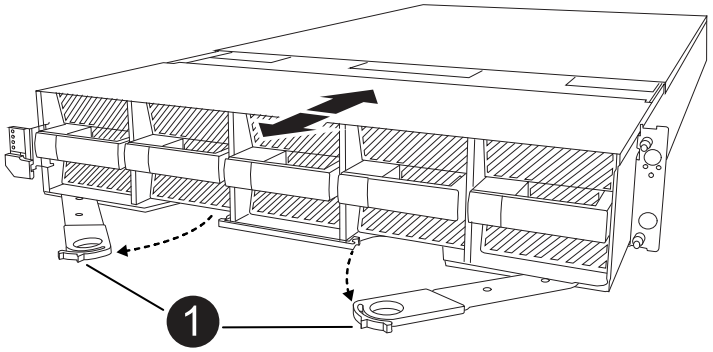
1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:



1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.
2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.



1	Locking cam latches

4. Slide the controller module out of the enclosure and place it on a flat, stable surface.

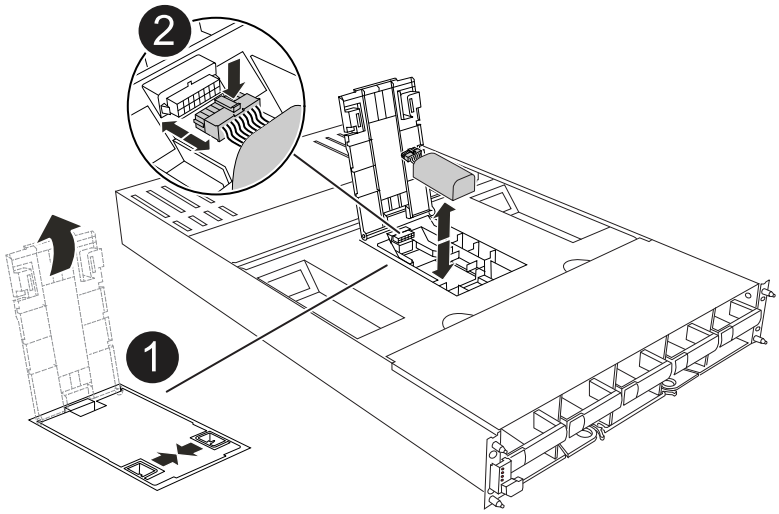
   Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

Step 3: Replace the NV battery

Remove the failed NV battery from the controller module and install the replacement NV battery.

Steps

1. Open the air duct cover and locate the NV battery.



1	NV battery air duct cover
2	NV battery plug

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module, and then set it aside.

5. Remove the replacement battery from its package.
6. Install the replacement battery pack into the controller:
   a. Plug the battery plug into the riser socket and make sure that the plug locks into place.
   b. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
7. Close the NV air duct cover.

   Make sure that the plug locks into the socket.

Step 4: Reinstall the controller module

Reinstall the controller module and boot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans

   > Do not use excessive force when sliding the controller module into the enclosure to avoid damaging the connectors.

   The controller module begins to boot as soon as it is fully seated in the enclosure.

4. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
6. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

I/O module

Overview of add and replace an I/O module - FAS70 and FAS90

The FAS70 or FAS90 system offers flexibility in expanding or replacing I/O modules to enhance network connectivity and performance. Adding or replacing an I/O module is essential when upgrading network capabilities or addressing a failed module.

You can replace a failed I/O module in your FAS70 or FAS90 storage system with the same type of I/O module,

or with a different kind of I/O module. You can also add an I/O module into a system with empty slots.

- Add an I/O module

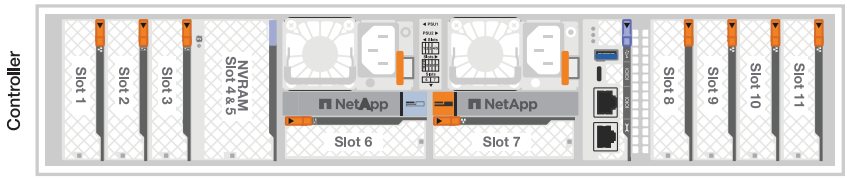
  Adding additional modules can improve redundancy, helping to ensure that the system remains operational even if one module fails.

- Replace an I/O module

  Replacing a failing I/O module can restore the system to its optimal operating state.

I/O slot numbering

The I/O slots on FAS70 and FAS90 controllers are numbered 1 through 11, as shown in the following illustration.



Add an I/O module - FAS70 and FAS90

Add an I/O module to yourFAS70 and FAS90 system to enhance network connectivity and expand your system's ability to handle data traffic.

You can add an I/O module to your FAS70 and FAS90 storage system when there are empty slots available or when all slots are fully populated.

Step 1: Shut down the impaired controller module

Shut down or take over the impaired controller module.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

Before you begin

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 command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the impaired controller:

   ```
   storage failover modify -node impaired-node -auto-giveback-of false
   ```

   When you see *Do you want to disable auto-giveback?*, enter `y`.

   a. If the impaired controller cannot be brought up or is already taken over, you must take the HA interconnect link down from the healthy controller before booting up the impaired controller. This will prevent the impaired controller from performing automatic giveback.

      ```
      system ha interconnect link off -node healthy-node -link 0

      system ha interconnect link off -node healthy-node -link 1
      ```

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 2: MetroCluster configuration

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next Step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Add the new I/O module

If the storage system has available slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your storage system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.

- Make sure that all other components are functioning properly.
- Make sure you have the replacement component you received from NetApp.

Add I/O module to an available slot

You can add a new I/O module into a storage system with available slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
3. Remove the target slot blanking module from the carrier:
   a. Depress the cam latch on the blanking module in the target slot.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.
4. Install the I/O module:
   a. Align the I/O module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
5. Cable the I/O module to the designated device.

   Make sure that any unused I/O slots have blanks installed to prevent possible thermal issues.

6. Rotate the cable management tray up to the closed position.
7. From the LOADER prompt, reboot the node:

   ```
   bye
   ```

   This reinitializes the I/O module and other components and reboots the node.

8. Give back the controller from the partner controller:

   ```
   storage failover giveback -ofnode target_node_name
   ```

9. Repeat these steps for controller B.
10. From the healthy node, restore automatic giveback if you disabled it:

    ```
    storage failover modify -node local -auto-giveback true
    ```

11. If AutoSupport is enabled, restore automatic case creation:

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Rotate the cable management tray down by pulling the buttons on the inside of the cable management tray and rotating it down.
4. Remove the target I/O module from the chassis:
   a. Depress the cam latch button.
   b. Rotate the cam latch away from the module as far as it will go.
   c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

      Make sure that you keep track of which slot the I/O module was in.
5. Install the I/O module into the target slot in the enclosure:
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
6. Cable the I/O module to the designated device.
7. Repeat the remove and install steps to replace additional modules for the controller.
8. Rotate the cable management tray up to the closed position.
9. Reboot the controller from the LOADER prompt:_bye_

   This reinitializes the PCIe cards and other components and reboots the node.

   If you encounter an issue during reboot, see BURT 1494308 - Environment shutdown might be triggered during I/O module replacement
10. Give back the controller from the partner controller:

    ```
    storage failover giveback -ofnode target_node_name
    ```
11. Enable automatic giveback if it was disabled:

```
storage failover modify -node local -auto-giveback true
```

12. Do one of the following:
    - If you removed a NIC I/O module and installed a new NIC I/O module, use the following network command for each port:

      ```
      storage port modify -node *<node name> -port *<port name> -mode network
      ```

    - If you removed a NIC I/O module and installed a storage I/O module, install and cable your NS224 shelves, as described in Hot-add workflow.
13. Repeat these steps for controller B.

Replace an I/O module - FAS70 and FAS90

Replace an I/O module in your FAS70 or FAS90 system when the module fails or requires an upgrade to support higher performance or additional features. The replacement process involves shutting down the controller, replacing the failed I/O module, rebooting the controller, and returning the failed part to NetApp.

You can use this procedure with all versions of ONTAP supported by your storage system.

Before you begin

- You must have the replacement part available.
- Make sure all other components in the storage system are functioning properly; if not, contact technical support.

Step 1: Shut down the impaired node

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next Step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Replace a failed I/O module

To replace an I/O module, locate it within the enclosure and follow the specific sequence of steps.

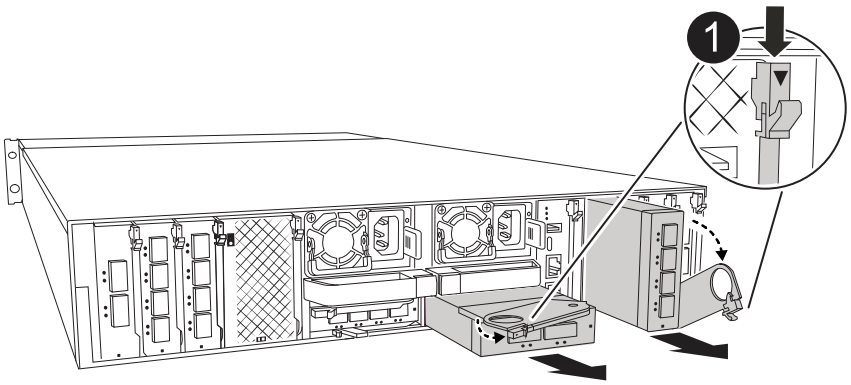
Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.

   Make sure that you label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.

3. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.

   This following illustration shows removing a horizontal and vertical I/O module. Typically, you will only remove one I/O module.



1	I/O cam latch

   Make sure that you label the cables so that you know where they came from.

4. Remove the target I/O module from the enclosure:
   a. Depress the cam button on the target module.
   b. Rotate the cam latch away from the module as far as it will go.

c. Remove the module from the enclosure by hooking your finger into the cam lever opening and pulling the module out of the enclosure.

   Make sure that you keep track of which slot the I/O module was in.

5. Set the I/O module aside.
6. Install the replacement I/O module into the enclosure:
   a. Align the module with the edges of the enclosure slot opening.
   b. Gently slide the module into the slot all the way into the enclosure, and then rotate the cam latch all the way up to lock the module in place.
7. Cable the I/O module.
8. Rotate the cable management tray up to the closed position.

Step 3: Reboot the controller

After you replace an I/O module, you must reboot the controller.

Steps

1. Reboot the controller from the LOADER prompt:

   ```
   bye
   ```

   Rebooting the impaired controller also reinitializes the I/O modules and other components.

2. Return the impaired controller to normal operation by giving back its storage:

   ```
   storage failover giveback -ofnode impaired_node_name
   ```

3. Restore automatic giveback from the console of the healthy controller:

   ```
   storage failover modify -node local -auto-giveback true
   ```

4. If AutoSupport is enabled, restore automatic case creation:

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Hot-swap a power supply - FAS70 and FAS90

Replace an AC or DC power supply unit (PSU) in your FAS70 or FAS90 system when it fails or becomes faulty, ensuring that your system continues to receive the required power for stable operation. The replacement process involves disconnecting the faulty PSU from the power source, unplugging the power cable, replacing the faulty PSU, and then reconnecting it to the power source.

The power supplies are redundant and hot-swappable. You do not have to shut down the controller to replace

a PSU.

About this task

- This procedure is written for replacing one PSU at a time.

  Do not mix PSUs with different efficiency ratings. Always replace like for like.

- Use the appropriate procedure for your type of PSU: AC or DC.

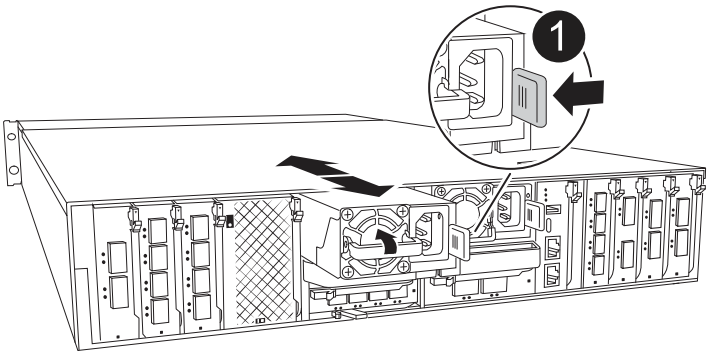
Option 1: Hot-swap an AC PSU

To replace an AC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Open the power cable retainer, and then unplug the power cable from the PSU.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull PSU out of the controller module.

   The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.



1	Terracotta PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one way.

      To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the PSU cabling:
   a. Reconnect the power cable to the PSU.
   b. Secure the power cable to the PSU using the power cable retainer.

   Once power is restored to the PSU, the status LED should be green.

7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

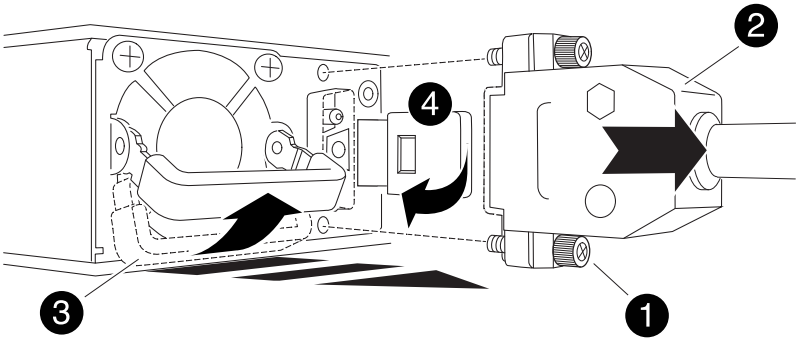
Option 2: Hot-swap a DC PSU

To replace a DC PSU, complete the following steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the PSU you want to replace, based on console error messages or through the red Fault LED on the PSU.
3. Disconnect the PSU:
   a. Unscrew the D-SUB DC cable connector using the thumb screws on the plug.
   b. Unplug the cable from the PSU and set it aside.
4. Remove the PSU by rotating the handle up, press the locking tab, and then pull the PSU out of the controller module.

   The PSU is short. Always use two hands to support it when removing it from the controller module so that it does not suddenly swing free from the controller module and injure you.



1	Thumb screws
2	D-SUB DC power PSU cable connector
3	Power supply handle
4	Blue PSU locking tab

5. Install the replacement PSU in the controller module:
   a. Using both hands, support and align the edges of the replacement PSU with the opening in the controller module.
   b. Gently push the PSU into the controller module until the locking tab clicks into place.

      The power supplies will only properly engage with the internal connector and lock in place one

way.

> To avoid damaging the internal connector, do not use excessive force when sliding the PSU into the system.

6. Reconnect the D-SUB DC power cable:
   a. Plug the power cable connector into the PSU.
   b. Secure the power cable to the PSU with the thumbscrews.

      Once power is restored to the PSU, the status LED should be green.
7. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - FAS70 and FAS90

Replace the real-time clock (RTC) battery, commonly known as a coin cell battery, in your FAS70 or FAS90 system to ensure that services and applications relying on accurate time synchronization remain operational.

Before you begin

- Understand that you can use this procedure with all versions of ONTAP supported by your system.
- Make sure all other components in the system are functioning properly; if not, you must contact technical support.

You must use an approved RTC battery.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

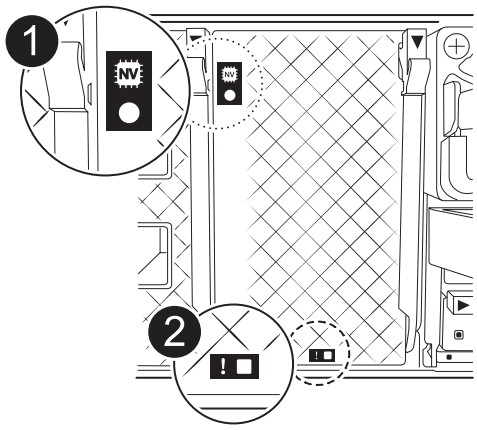
If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Step 2: Remove the controller module

You must remove the controller module from the enclosure when you replace the controller module or replace a component inside the controller module.

Steps

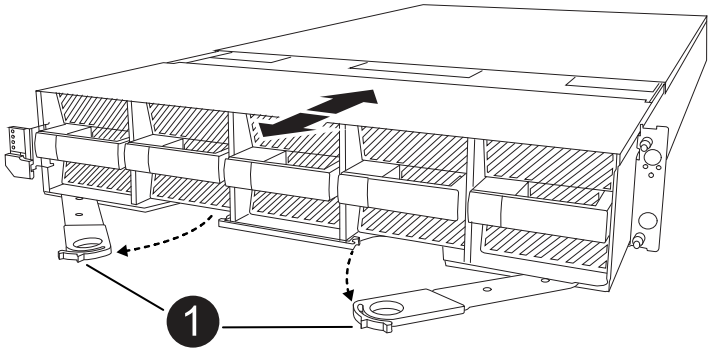
1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:



1	NVRAM status LED
2	NVRAM attention LED

   - If the NV LED is off, go to the next step.
   - If the NV LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.
2. If you are not already grounded, properly ground yourself.
3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

   The controller module moves slightly out of the enclosure.



1	Locking cam latches

4. Slide the controller module out of the enclosure and place it on a flat, stable surface.

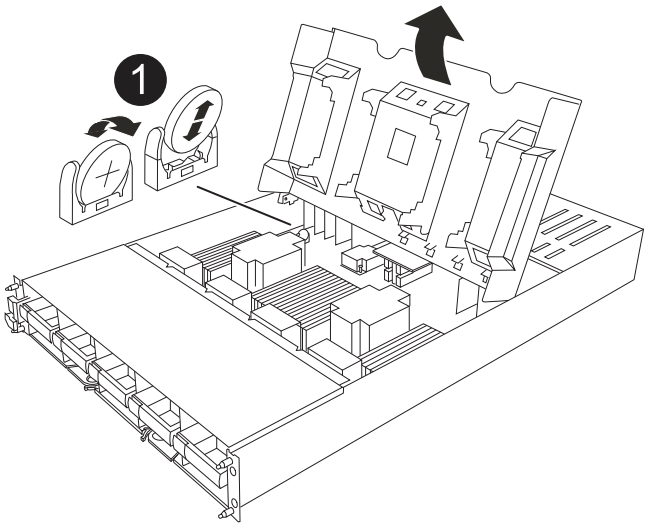
   Make sure that you support the bottom of the controller module as you slide it out of the enclosure.

Step 3: Replace the RTC battery

Remove failed RTC battery and install the replacement RTC battery.

Steps

1. Open the controller air duct on the top of the controller.
   a. Insert your fingers in the recesses at the far ends of the air duct.
   b. Lift the air duct and rotate it upward as far as it will go.
2. Locate the RTC battery under the air duct.



1	RTC battery and housing

3. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

Observe 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 indicates the correct orientation.

4. Remove the replacement battery from the antistatic shipping bag.
5. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
6. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

Step 4: Reinstall the controller module

Reinstall the controller module and boot it.

Steps

1. Ensure the air duct is completely closed by rotating it down as far as it will go.

   It must lie flush against the controller module sheet metal.

2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.
3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans

   Do not use excessive force when sliding the controller module into the enclosure to avoid damaging the connectors.

   The controller module begins to boot as soon as it is fully seated in the enclosure.

4. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
6. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 5: Reset the time and date on the controller

After replacing the RTC battery, inserting controller and powering on first BIOS reset, you will see the following error messages:
`RTC date/time error. Reset date/time to default`
`RTC power failure error`
These messages are excpected and you can continue with this procedure.

Steps

1. Check the date and time on the healthy controller with the *cluster date show* command.

If your system stops at the boot menu, select the option for `Reboot node` and respond *y* when prompted, then boot to LOADER by pressing *Ctrl-C*

a. At the LOADER prompt on the target controller, check the time and date with the `cluster date show` command.

b. If necessary, modify the date with the `set date mm/dd/yyyy` command.

c. If necessary, set the time, in GMT, using the `set time hh:mm:ss` command.

   1. Confirm the date and time on the target controller.
   2. At the LOADER prompt, enter *bye* to reinitialize the PCIe cards and other components and let the controller reboot.

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace system management module - FAS70 and FAS90

Replace the System Management module in your FAS70 or FAS90 system when it becomes defective or its firmware is corrupted. The replacement process involves shutting down the controller, replacing the failed System Management module, rebooting the controller, updating the license keys, and returning the failed part to NetApp.

The System Management module, located at the back of the controller in slot 8, contains onboard components for system management, as well as ports for external management. The target controller must be shut down to replace an impaired System Management module or replace the boot media.

The System Management module has the following onboard components:

- Boot media, allowing boot media replacement without removing the controller module.
- BMC
- Management switch

The System Management module also contains the following ports for external management:

- RJ45 Serial
- USB Serial (Type-C)
- USB Type-A (Boot recovery)
- e0M RJ45 Ethernet

Before you begin

- Make sure all other system components are working properly.
- Make sure that the partner controller is able to take over the impaired controller.
- Make sure you replace the failed component with a replacement component you received from NetApp.

About this task

This procedure uses the following terminology:

- The impaired controller is the controller on which you are performing maintenance.
- The healthy controller is the HA partner of the impaired controller.

Step 1: Shut down the impaired controller

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- You must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state:

  `metrocluster node show`

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:*> system node autosupport invoke -node * -type all -message
   MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node local -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next section.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

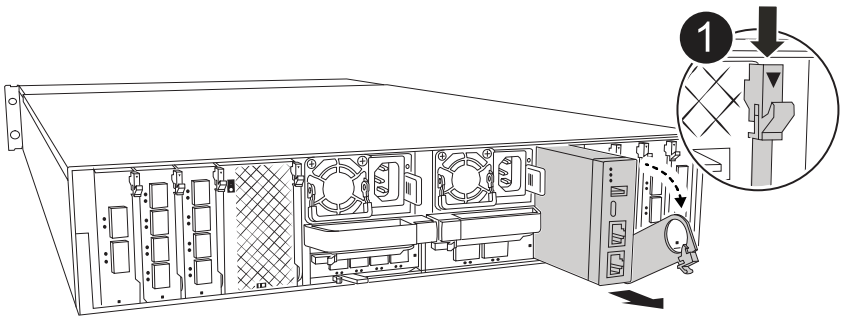
Step 2: Replace the impaired System Management module

Replace the impaired system management module.

Steps

1. Remove the System Management module:

> Make sure NVRAM destage has completed before proceeding. When the LED on the NV module is off, NVRAM is destaged. If the LED is flashing, wait for the flashing to stop. If flashing continues for longer than 5 minutes, contact Technical Support for assistance.



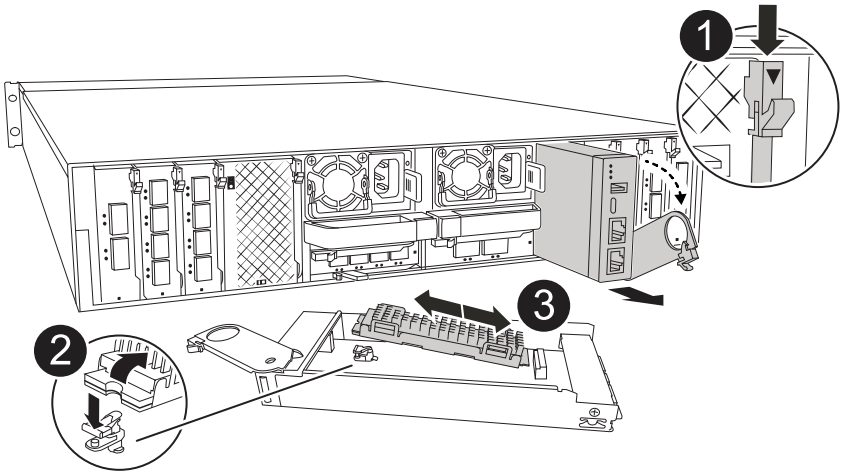
1	System Management module cam latch

   a. If you are not already grounded, properly ground yourself.
   b. Unplug the power supply cables from the PSUs.
2. Remove the System Manage module
   a. Remove any cables connected to the System Management module. Make sure that label where the cables were connected, so that you can connect them to the correct ports when you reinstall the module.
   b. Disconnect the power cords from the PSU for the impaired controller.
   c. Rotate the cable management tray down by pulling the buttons on both sides on the inside of the cable management tray and then rotate the tray down.
   d. Depress the cam button on the System Management module.
   e. Rotate the cam lever down as far as it will go.

f. Loop your finger into the hole on the cam lever and pull the module straight out of the system.

g. Place the System Management module on an anti-static mat, so that the boot media is accessible.

3. Move the boot media to the replacement System Management module:



1	System Management module cam latch
2	Boot media locking button
3	Boot media

a. Press the blue boot media locking button in the impaired System Management module.

b. Rotate the boot media up and slide it out of the socket.

4. Install the boot media in the replacement System Management module:

a. Align the edges of the boot media with the socket housing, and then gently push it squarely into the socket.

b. Rotate the boot media down until it touches the locking button.

c. Depress the blue locking and rotate the boot media all the way down and release the blue locking button.

5. Install the replacement System Management module into the enclosure:

a. Align the edges of the replacement System Management module with the system opening and gently push it into the controller module.

b. Gently slide the module into the slot until the cam latch begins to engage with the I/O cam pin, and then rotate the cam latch all the way up to lock the module in place.

6. Rotate the cable management arm up to the closed position.

7. Recable the System Management module.

Step 3: Reboot the controller module

Reboot the controller module.

Steps

1. Plug the power cables back into the PSU.

   The system will begin to reboot, typically to the LOADER prompt.

2. Enter *bye* at the LOADER prompt.
3. Return the impaired controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`.
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`.
5. If AutoSupport is enabled, restore/unsuppress automatic case creation: `system node autosupport invoke -node * -type all -message MAINT=END`.

Step 4: Install licenses and register serial number

You must install new licenses for the 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 node. However, if the 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 for the 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`

4. 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.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Key specifications

Key specifications for FAS70

The following are select specifications for the FAS70 storage system in a single high availability pair. Visit NetApp Hardware Universe (HWU) for the complete specifications for this storage system.

FAS70 specifications at a glance

- Platform Configuration: FAS70 Dual Chassis HA Pair
- Max Raw Capacity: 15.0000 PB
- Memory: 256.0000 GB
- Form Factor: 2U chassis with 1 HA controllers
- ONTAP Version: ONTAP: 9.16.1P2
- PCIe Expansion Slots: 18
- Minimum ONTAP Version: ONTAP 9.15.1

Scaleout maximums

- Type: NAS; HA Pairs: 12; Raw Capacity: 180.0 PB / 159.9 PiB; Max Memory: 3072 GB
- Type: SAN; HA Pairs: 6; Raw Capacity: 90.0 PB / 79.9 PiB; Max Memory: 1536 GB
- Type: HA Pair; Raw Capacity: 15.0 PB / 13.3 PiB; Max Memory: 256.0000

I/O

Onboard I/O

No onboard I/O data.

Total I/O

- Protocol: Ethernet 200 Gbps; Ports: 20
- Protocol: Ethernet 100 Gbps; Ports: 32
- Protocol: Ethernet 25 Gbps; Ports: 48
- Protocol: Ethernet 10 Gbps; Ports: 48
- Protocol: FC 64 Gbps; Ports: 48

- Protocol: NVMe/FC 64 Gbps; Ports: 48
- Ports: 0
- Protocol: SAS 12 Gbps; Ports: 48

Management ports

- Protocol: Ethernet 1 Gbps; Ports: 2
- Protocol: RS-232 115 Kbps; Ports: 4
- Protocol: USB 600 Mbps; Ports: 2

Storage networking supported

- CIFS
- FC
- iSCSI
- NFS v3
- NFS v4.0
- NFS v4.1
- NFS v4.2
- NVMe/FC
- 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: 3392 BTU/hr
- Worst-case Power: 5136 BTU/hr
- Weight: 59.5 lb, 27.0 kg
- Height: 2U
- Width: 19" IEC rack-compliant (17.7" 44.9 cm)
- Depth: 30.0" (35.2" with cable management bracket)
- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°F) at up to 3048m (10000 ft) elevation; 8% to 80% 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): 8.5; Sound pressure (LpAm) (bystander positions): 67.7 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-55032,
  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 15598-1,
  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 for external shelves

Key specifications for FAS90

The following are select specifications for the FAS90 storage system in a single high availability pair. Visit NetApp Hardware Universe (HWU) for the complete specifications for this storage system.

FAS90 specifications at a glance

- Platform Configuration: FAS90 Dual Chassis HA Pair
- Max Raw Capacity: 15.0000 PB
- Memory: 1024.0000 GB
- Form Factor: 2U chassis with 1 HA controllers
- ONTAP Version: ONTAP: 9.16.1P2
- PCIe Expansion Slots: 18
- Minimum ONTAP Version: ONTAP 9.15.1

Scaleout maximums

- Type: NAS; HA Pairs: 12; Raw Capacity: 180.0 PB / 159.9 PiB; Max Memory: 12288 GB
- Type: SAN; HA Pairs: 6; Raw Capacity: 90.0 PB / 79.9 PiB; Max Memory: 6144 GB
- Type: HA Pair; Raw Capacity: 15.0 PB / 13.3 PiB; Max Memory: 1024.0000

I/O

Onboard I/O

No onboard I/O data.

Total I/O

- Protocol: Ethernet 200 Gbps; Ports: 20
- Protocol: Ethernet 100 Gbps; Ports: 32
- Protocol: Ethernet 25 Gbps; Ports: 48
- Protocol: Ethernet 10 Gbps; Ports: 48
- Protocol: FC 64 Gbps; Ports: 48
- Protocol: NVMe/FC 64 Gbps; Ports: 48
- Ports: 0
- Protocol: SAS 12 Gbps; Ports: 48

Management ports

- Protocol: Ethernet 1 Gbps; Ports: 2
- Protocol: RS-232 115 Kbps; Ports: 4
- Protocol: USB 600 Mbps; Ports: 2

Storage networking supported

- CIFS
- FC
- iSCSI
- NFS v3
- NFS v4.0
- NFS v4.1
- NFS v4.2
- NVMe/FC
- 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: 3392 BTU/hr
- Worst-case Power: 5136 BTU/hr
- Weight: 59.5 lb, 27.0 kg
- Height: 2U
- Width: 19" IEC rack-compliant (17.7" 44.9 cm)
- Depth: 30.0" (35.2" with cable management bracket)
- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°F) at up to 3048m (10000 ft) elevation; 8% to 80% 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): 8.5; Sound pressure (LpAm) (bystander positions): 67.7 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-55032,
  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 15598-1,
  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 for external shelves

FAS2700 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.

Quick guide - FAS2700

This page gives 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.

Access the *Installation and Setup Instructions* PDF poster:

AFF A220/FAS2700 Systems Installation and Setup Instructions

Video steps - FAS2700

The following video shows how to install and cable your new system.

[] | *https://img.youtube.com/vi/FUtG1Je5D1g?/maxresdefault.jpg*

Detailed guide - FAS2700

This page gives detailed step-by-step instructions for installing a typical NetApp system. Use this guide if you want more detailed installation instructions.

Step 1: Prepare for installation

To install your FAS2700 system, you need to create an account on the NetApp Support Site, register your system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

You need to have access to the Hardware Universe for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

NetApp Hardware Universe

Find the Release Notes for your version of ONTAP 9

You need to provide the following 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.

      NetApp Product Registration

4. Download and install Config Advisor on your laptop.

   NetApp Downloads: Config Advisor

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 the Hardware Universe to locate the cable and identify its use.

   NetApp Hardware Universe

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

Type of cable…	Part number and length	Connector type	For…
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*.

   Cluster Configuration Worksheet

Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

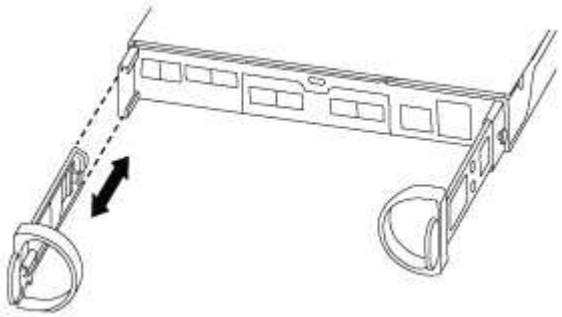
Steps

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 your network

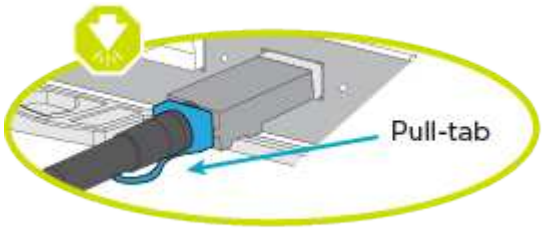
You can cable the controllers to your network by using the two-node switchless cluster method or by using the cluster interconnect network.

Option 1: Cable a two-node switchless cluster, unified network configuration

Management network, UTA2 data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers.

You must have contacted your network administrator for information about connecting the system to the switches.

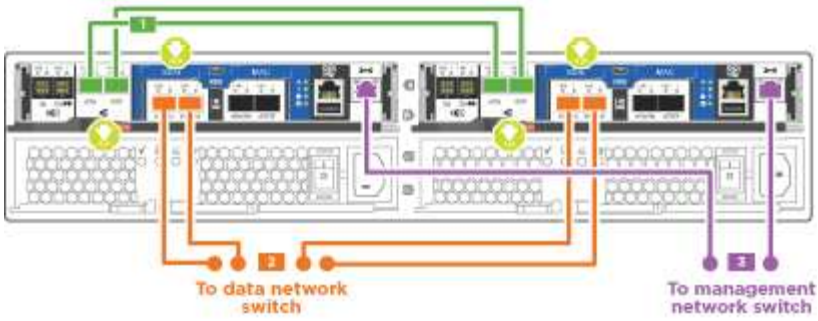
Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

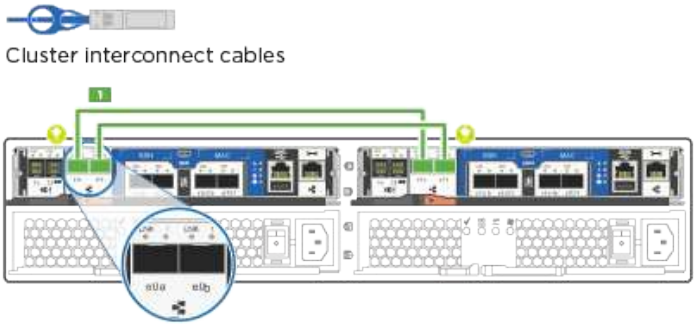


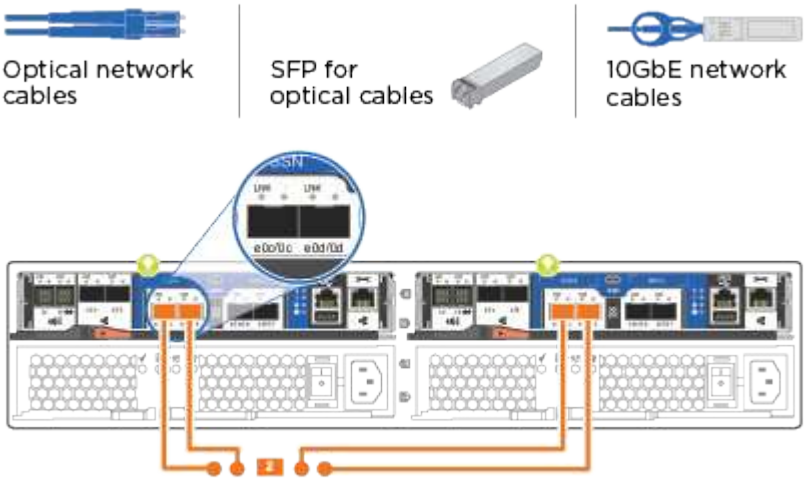
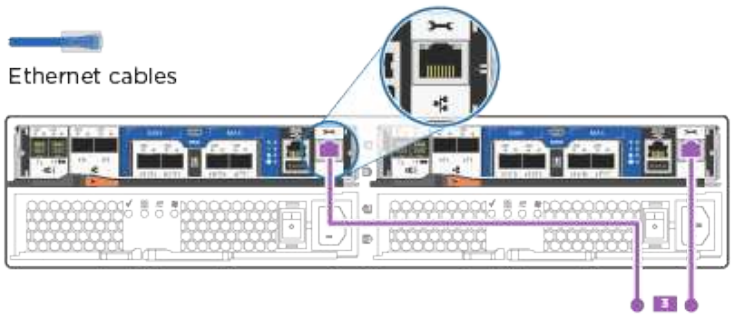
As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. You can use the graphic or the step-by step instructions to complete the cabling between the controllers and to the switches:



Step	Perform on each controller
1	Cable the cluster interconnect ports to each other with the cluster interconnect cable: • e0a to e0a • e0b to e0b 

Step	Perform on each controller
2	Use one of the following cable types to cable the UTA2 data ports to your host network: An FC host • 0c and 0d • **or** 0e and 0f A 10GbE • 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. 
3	Cable the e0M ports to the management network switches with the RJ45 cables: 
	DO NOT plug in the power cords at this point.

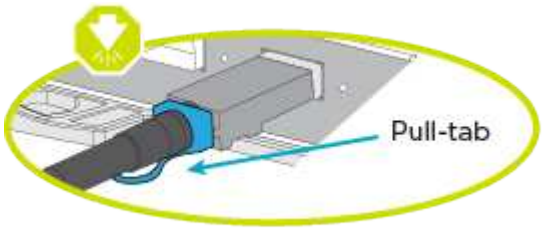
2. To cable your storage, see Step 4: Cable controllers to drive shelves

Option 2: Cable a switched cluster, unified network configuration

Management network, UTA2 data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled to the cluster interconnect switches.

You must have contacted your network administrator for information about connecting the system to the switches.

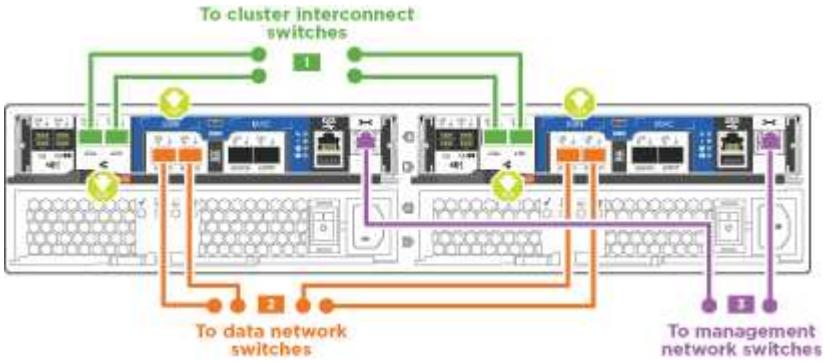
Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

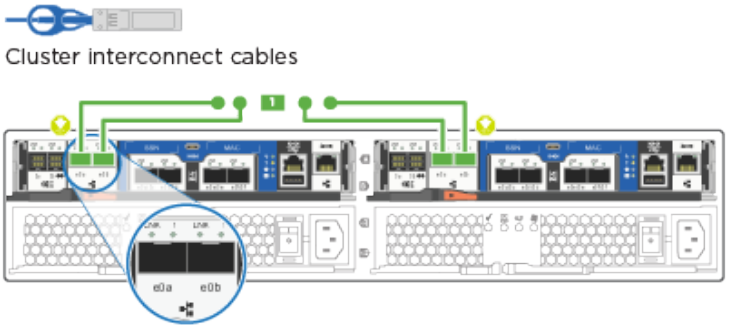
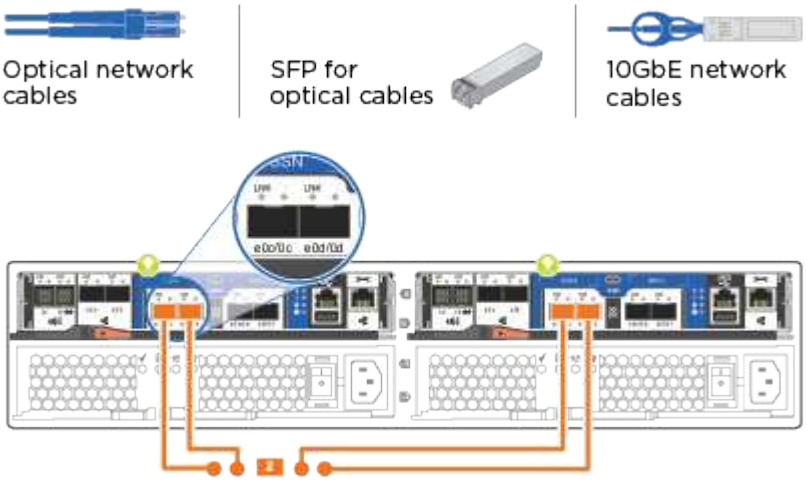


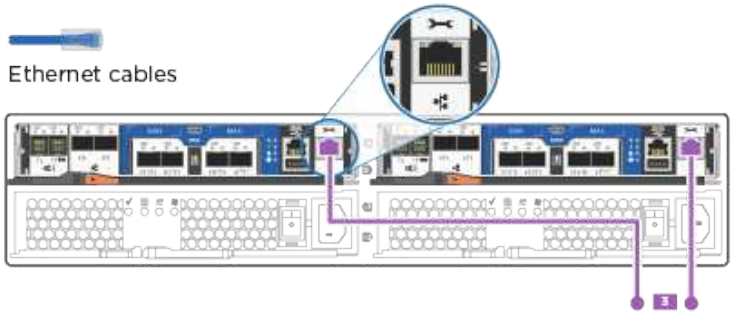
As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. You can use the graphic or the step-by step instructions to complete the cabling between the controllers and the switches:



Step	Perform on each controller module
1	Cable e0a and e0b to the cluster interconnect switches with the cluster interconnect cable: 
2	Use one of the following cable types to cable the UTA2 data ports to your host network: An FC host • 0c and 0d • **or** 0e and 0f A 10GbE • 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. 

Step	Perform on each controller module
3	Cable the e0M ports to the management network switches with the RJ45 cables: 
	DO NOT plug in the power cords at this point.

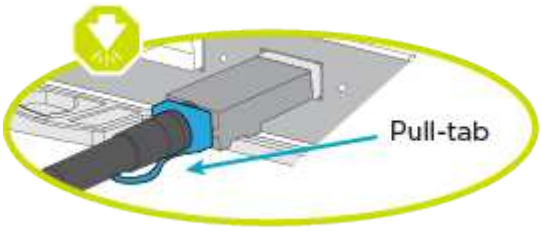
2. To cable your storage, see Step 4: Cable controllers to drive shelves

Option 3: Cable a two-node switchless cluster, Ethernet network configuration

Management network, Ethernet data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers.

You must have contacted your network administrator for information about connecting the system to the switches.

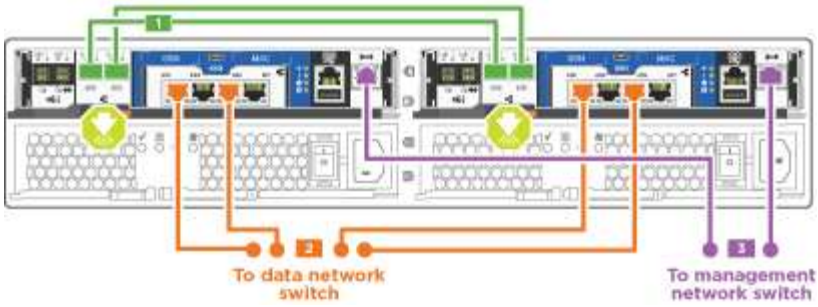
Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

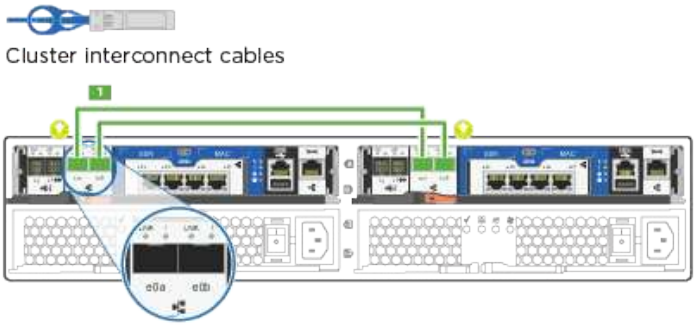
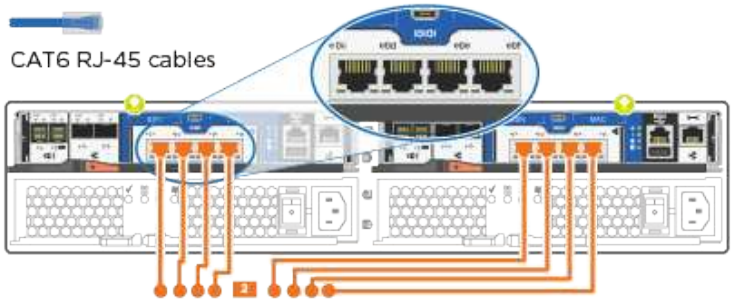


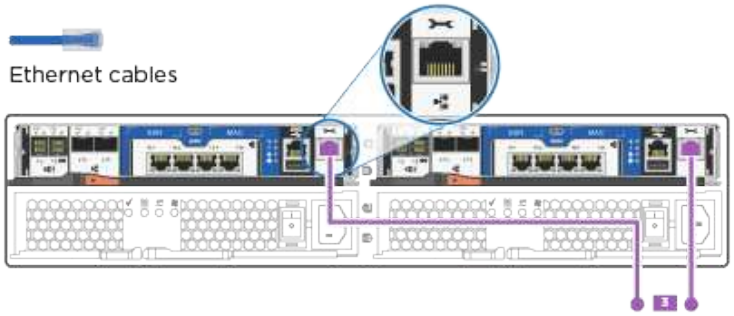
As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. You can use the graphic or the step-by step instructions to complete the cabling between the controllers and to the switches:



Step	Perform on each controller
1	Cable the cluster interconnect ports to each other with the cluster interconnect cable: • e0a to e0a • e0b to e0b 
2	Use the Cat 6 RJ45 cable to cable the e0c through e0f ports to your host network: 

Step	Perform on each controller
3	Cable the e0M ports to the management network switches with the RJ45 cables: 
!	DO NOT plug in the power cords at this point.

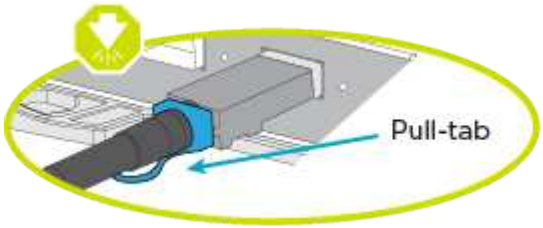
2. To cable your storage, see Step 4: Cable controllers to drive shelves

Option 4: Cable a switched cluster, Ethernet network configuration

Management network, Ethernet data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled to the cluster interconnect switches.

You must have contacted your network administrator for information about connecting the system to the switches.

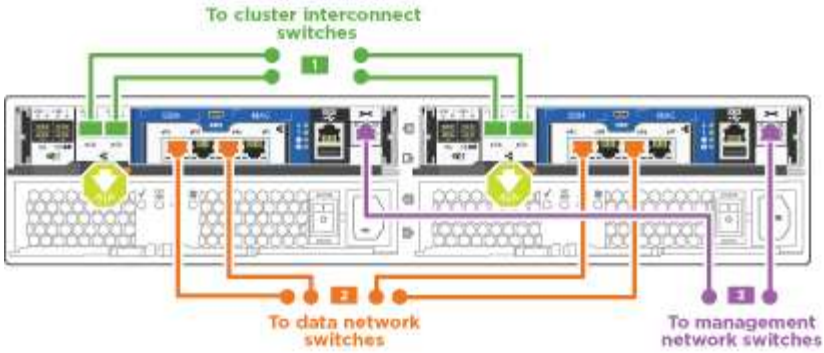
Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.

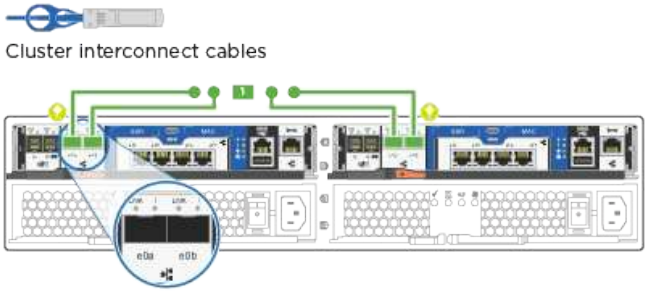
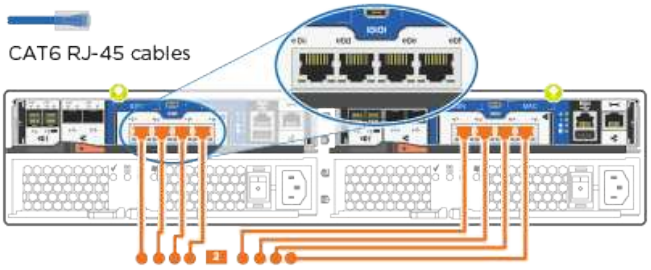


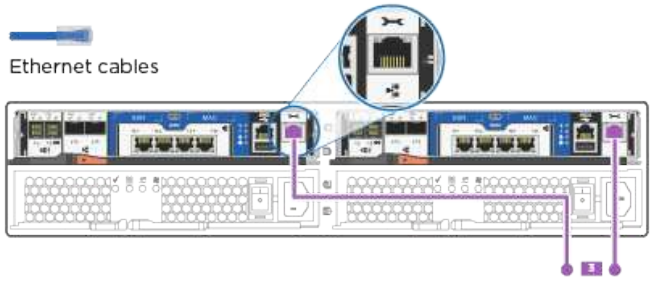
As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. You can use the graphic or the step-by step instructions to complete the cabling between the controllers and the switches:



Step	Perform on each controller module
1	Cable e0a and e0b to the cluster interconnect switches with the cluster interconnect cable: 
2	Use the Cat 6 RJ45 cable to cable the e0c through e0f ports to your host network: 

Step	Perform on each controller module
3	Cable the e0M ports to the management network switches with the RJ45 cables: 
	DO NOT plug in the power cords at this point.

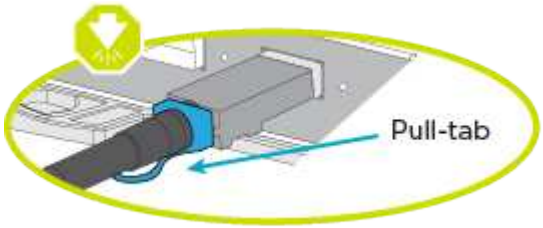
2. To cable your storage, see Step 4: Cable controllers to drive shelves

Step 4: Cable controllers to drive shelves

You must cable the controllers to your shelves using the onboard storage ports. NetApp recommends MP-HA cabling for systems with external storage. 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.

Before you begin

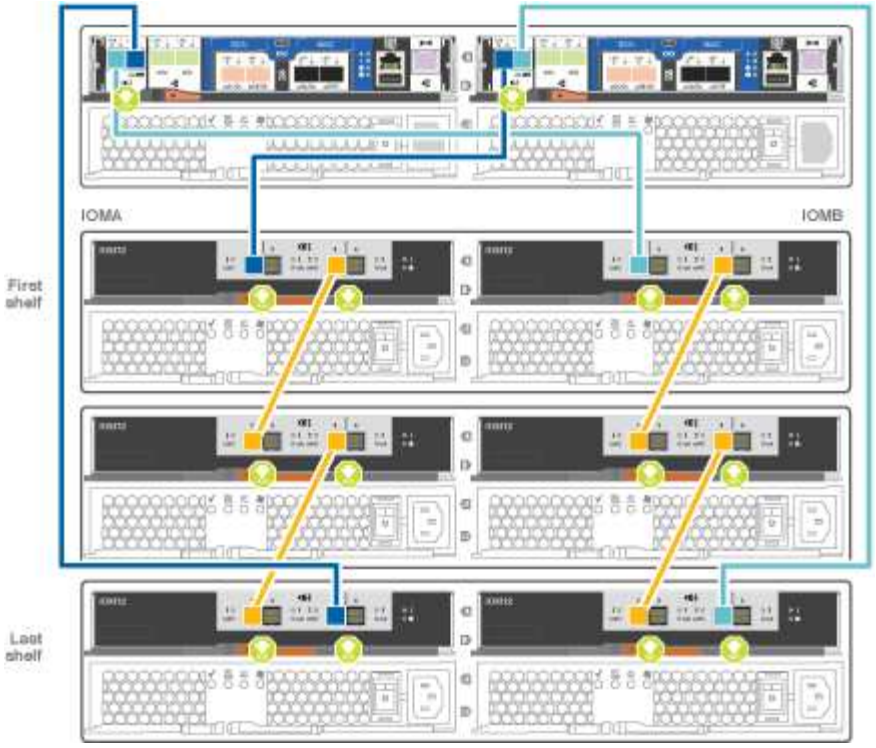
- You must cable the shelf-to-shelf connections, and then cable both controllers to the drive shelves.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.



Steps

1. Cable the HA pair with external drive shelves:

The example uses DS224C. Cabling is similar with other supported drive shelves.



Step	Perform on each controller
1	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
2	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
3	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 more SAS cabling information and worksheets, see SAS cabling rules, worksheets, and examples overview - shelves with IOM12 modules

1. To complete setting up your system, see Step 5: Complete system setup and configuration

Step 5: Complete system setup and configuration

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: Complete system setup 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. Use the following animation to connect your laptop to the Management switch.

   Animation - Connect your laptop to the Management switch

6. Select an ONTAP icon listed to discover:

   drw_autodiscovery_controler_select_ieops-1849.svg[Select an ONTAP icon]

   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. Use System Manager guided setup to configure your system using the data you collected in the *NetApp ONTAP Configuration Guide*.

ONTAP Configuration Guide

8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Completing system setup and configuration if network discovery is not enabled

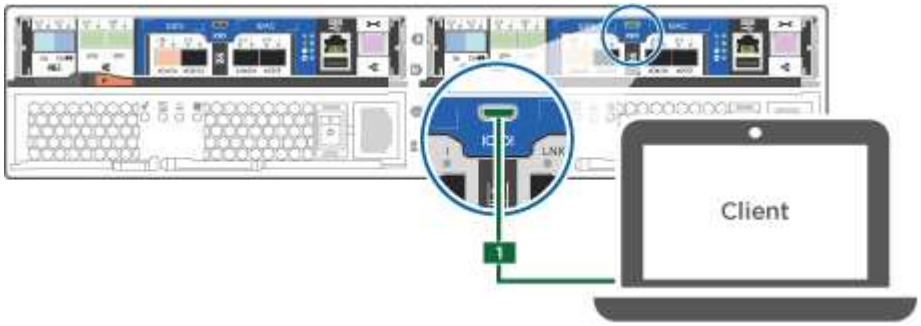
If network discovery is not enabled on your laptop, you must complete the configuration and setup using this task.

Steps

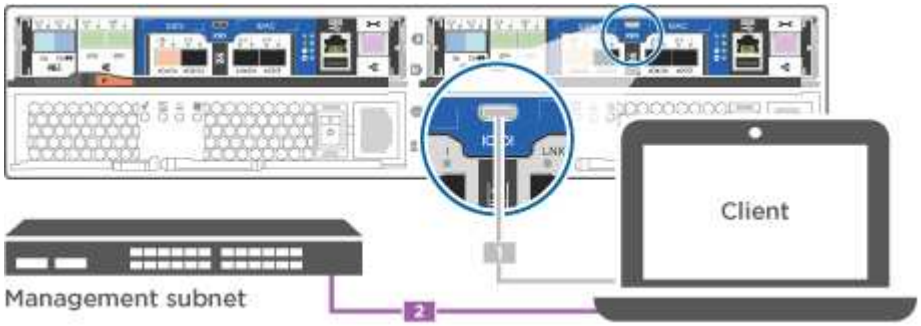
1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for 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 *NetApp ONTAP Configuration guide*.
      ONTAP Configuration Guide
7. Verify the health of your system by running Config Advisor.
8. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain FAS2700 hardware

For the FAS2700 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.

Caching module

You must replace the controller's caching module when your system registers a single AutoSupport (ASUP) message that the module has gone offline.

Chassis

The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.

Controller

A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.

DIMM

You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.

Drive

A drive is a device that provides the physical storage media for data.

NVMEM battery

A battery is included with the 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 - FAS2700

Learn about boot media replacement on a FAS2700 system and understand the different replacement methods.

The boot media stores primary and secondary boot image files that the system uses during startup. Depending on your network configuration, you can perform either a nondisruptive replacement (HA pair connected to network) or a disruptive replacement (requires two reboots).

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 - FAS2700

Verify encryption key support and status before shutting down the impaired controller on a FAS2700 system. This procedure includes checking ONTAP version compatibility with NetApp Volume Encryption (NVE), verifying the key manager configuration, and backing up encryption information to ensure data security during boot media recovery.

The FAS2700 system supports only manual boot media recovery procedures. Automated boot media recovery is not supported.

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 - FAS2700

Shut down or take over the impaired controller using the appropriate procedure for your configuration. The FAS2700 system supports only manual boot media recovery procedures. Automated boot media recovery is not supported.

Option 1: Most configurations

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

Option 2: Controller is in a MetroCluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `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 - FAS2700

Replace the failed boot media on a FAS2700 controller module. This procedure includes removing the controller module from the chassis, physically replacing the boot media component, transferring the boot image to the replacement media using a USB flash drive, and restoring the system to normal operation.

The FAS2700 system supports only manual boot media recovery procedures. Automated boot media recovery is not supported.

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.

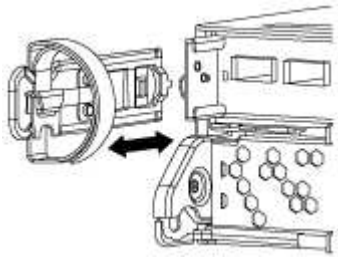
Step 1: Remove the controller module

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.

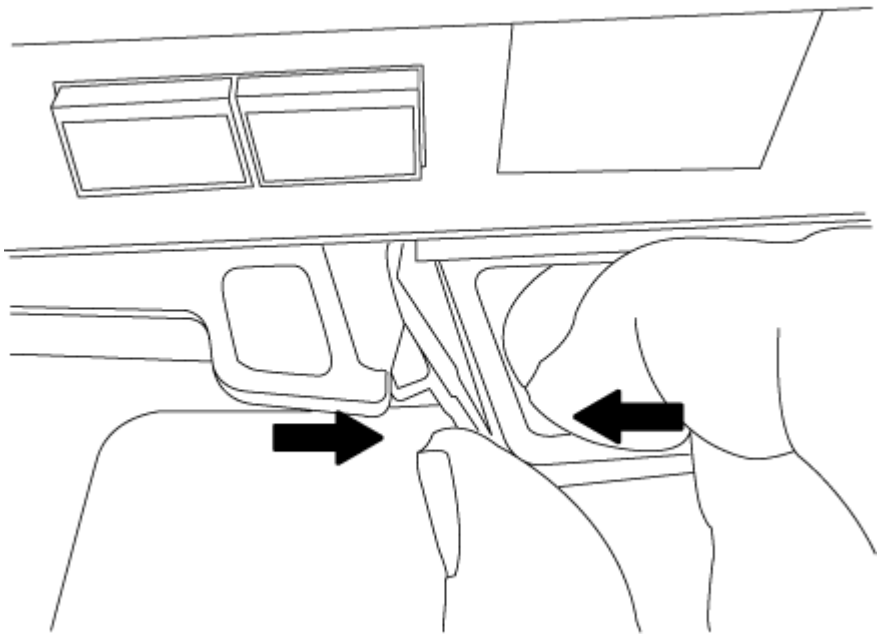
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

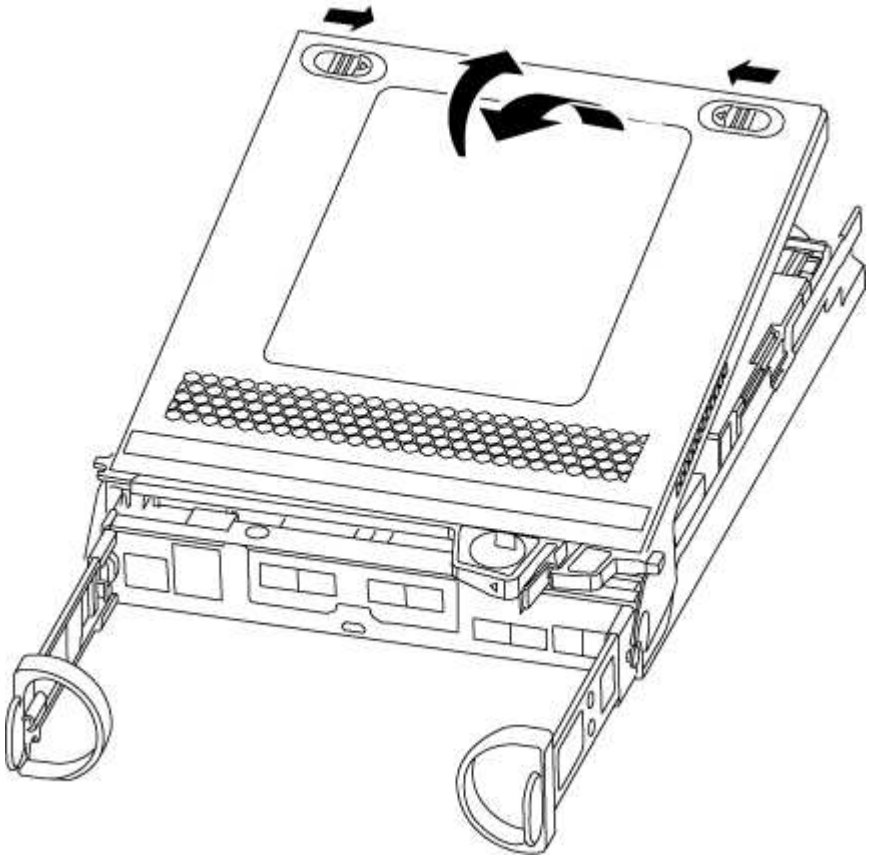
3. Remove and set aside the cable management devices from the left and right sides of the controller module.



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.

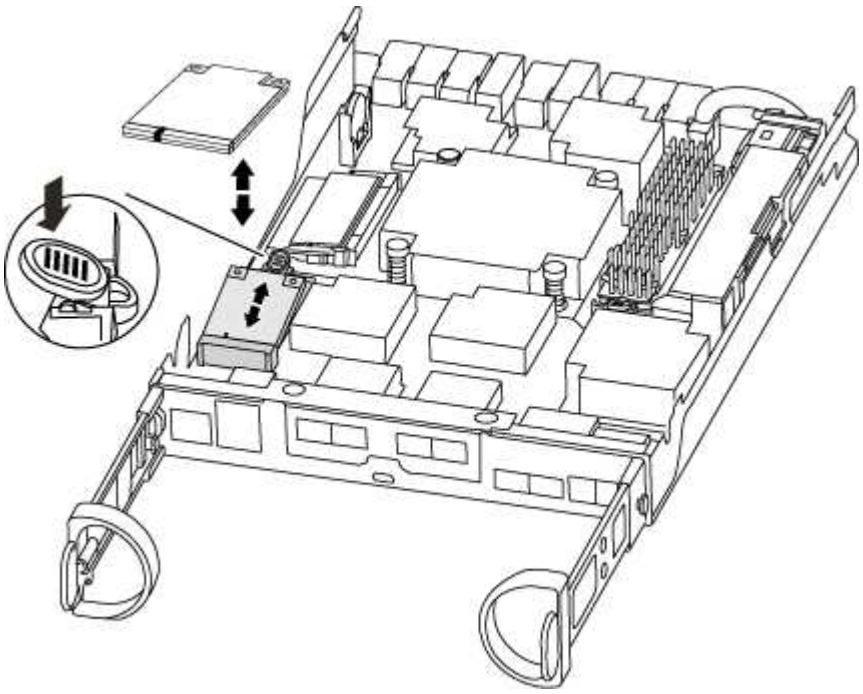


Step 2: Replace the boot media

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 - FAS2700

Boot the ONTAP recovery image from the USB drive on a FAS2700 system to restore the boot media. This procedure includes booting from the USB flash drive, restoring the file system, verifying environmental variables, and returning the controller to normal operation after boot media replacement.

The FAS2700 system supports only manual boot media recovery procedures. Automated boot media recovery is not supported.

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.

If your system has…	Then…
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 - FAS2700

Restore encryption on the replacement boot media. The FAS2700 system supports only manual boot media recovery procedures. Automated boot media recovery is not supported.

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---------------------------
012345678901234567890123456789012345678901234567890123456789012
23
123456789012345678901234567890123456789012345678901234567890123
34
234567890123456789012345678901234567890123456789012345678901234
45
345678901234567890123456789012345678901234567890123456789012345
56
456789012345678901234567890123456789012345678901234567890123456
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
```

```
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
012345678901234567890123456789012345678901234567890123456789012345678901
23
123456789012345678901234567890123456789012345678901234567890123456789012
34
234567890123456789012345678901234567890123456789012345678901234567890123
45
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
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 A220 and FAS2700

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.

The FAS2700 system supports only manual boot media recovery procedures. Automated boot media recovery is not supported.

Replace the caching module - FAS2700

You must replace the caching module in the controller module when your system registers a single AutoSupport (ASUP) message that the module has gone offline; failure to do so results in performance degradation.

- You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller.

Synchronize a node with the cluster

You might want to erase the contents of your caching module before replacing it.

Steps

1. Although data on the caching module is encrypted, you might want to erase any data from the impaired caching module and verify that the caching module has no data:
   a. Erase the data on the caching module: `system controller flash-cache secure-erase run -node node name localhost -device-id device_number`

      Run the `system controller flash-cache show` command if you don't know the Flash Cache device ID.

   b. Verify that the data has been erased from the caching module: `system controller flash-cache secure-erase show`
2. If the impaired controller is part of an HA pair, 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: • For an HA pair, take over the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`. • For a stand-alone system: `system node halt impaired_node_name`

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

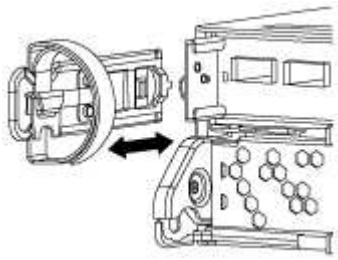
To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

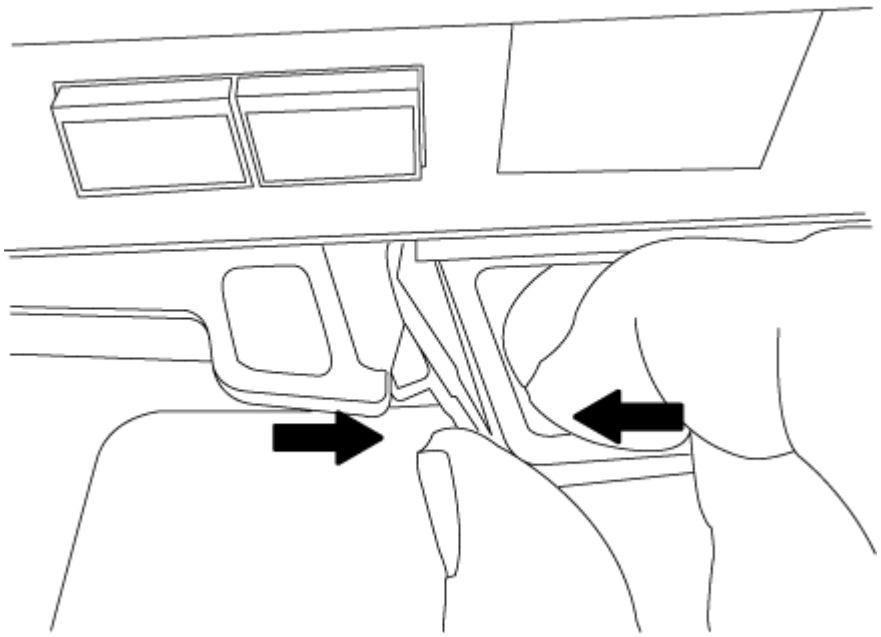
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

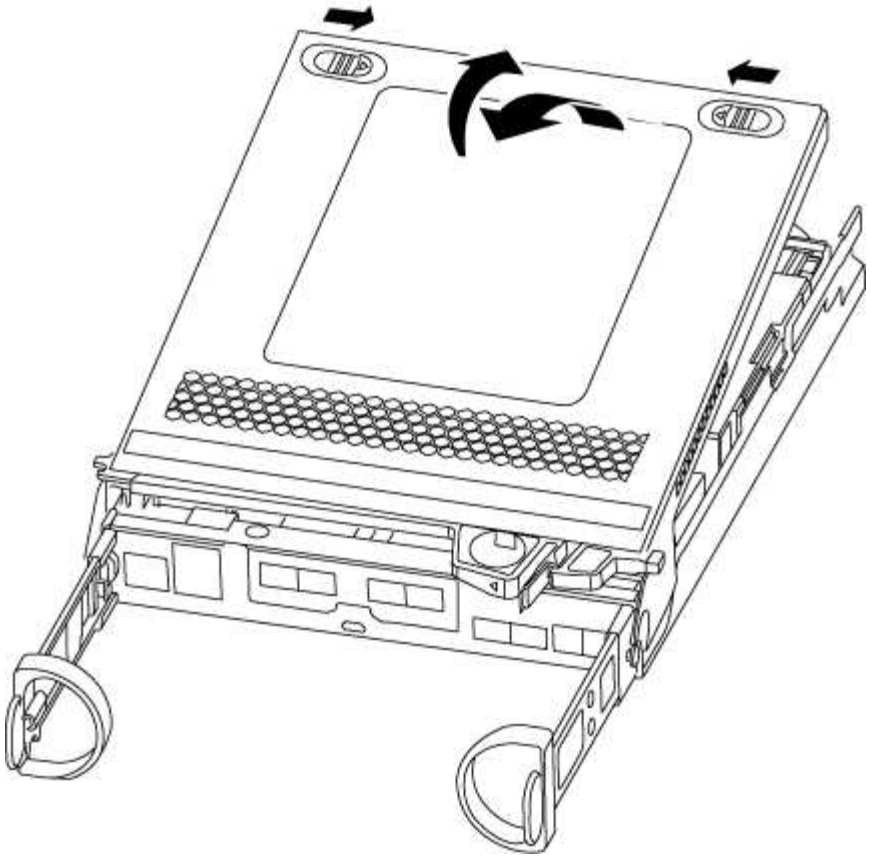
3. Remove and set aside the cable management devices from the left and right sides of the controller module.



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 a caching module

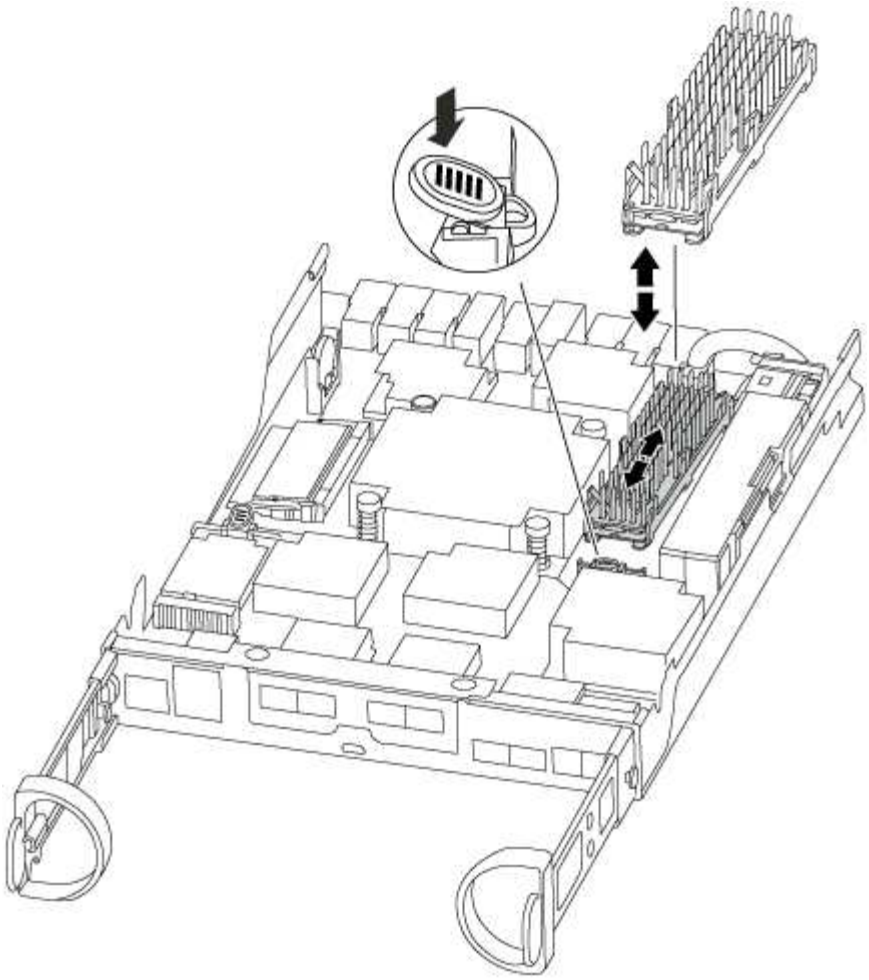
To replace a caching module referred to as the M.2 PCIe card on the label on your controller, locate the slot inside the controller and follow the specific sequence of steps.

Your storage system must meet certain criteria depending on your situation:

- It must have the appropriate operating system for the caching module you are installing.
- It must support the caching capacity.
- All other components in the storage system must be functioning properly; if not, you must contact technical support.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the caching module at the rear of the controller module and remove it.
    a. Press the release tab.
    b. Remove the heatsink.



3. Gently pull the caching module straight out of the housing.
4. Align the edges of the caching module with the socket in the housing, and then gently push it into the socket.
5. 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.

6. Reseat and push the heatsink down to engage the locking button on the caching module housing.
7. Close the controller module cover, as needed.

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, and 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.

Chassis

Overview of chassis replacement - FAS2700

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 - FAS2700

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 accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.

3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Controller is in a MetroCluster configuration

> Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message`

`MAINT=number_of_hours_downh`

The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Move and replace hardware - AFF A220 and FAS2700

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.

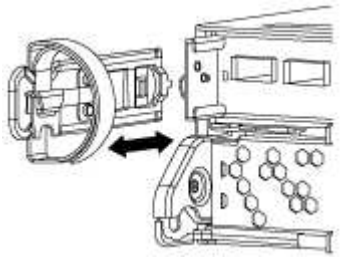
Step 2: Remove the controller module

Remove the controller module or modules from the old chassis.

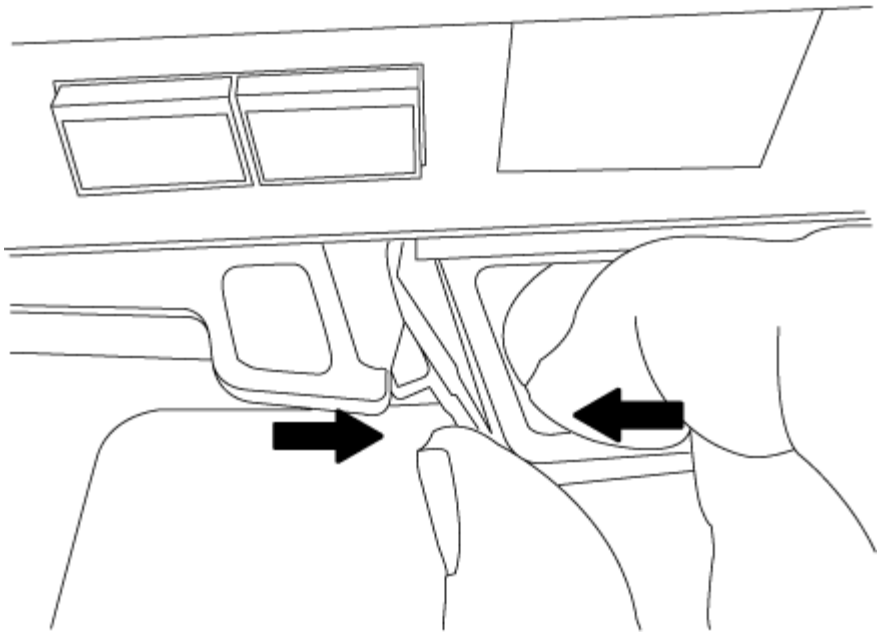
1. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   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 the 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.

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.

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. 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 - FAS2700

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. The next step depends on your system configuration.
5. 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 - FAS2700

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 - FAS2700

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller module hardware - FAS2700

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.

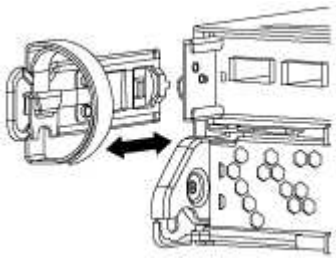
Step 1: Remove controller module

To replace the controller module, you must first remove the old controller module from the chassis.

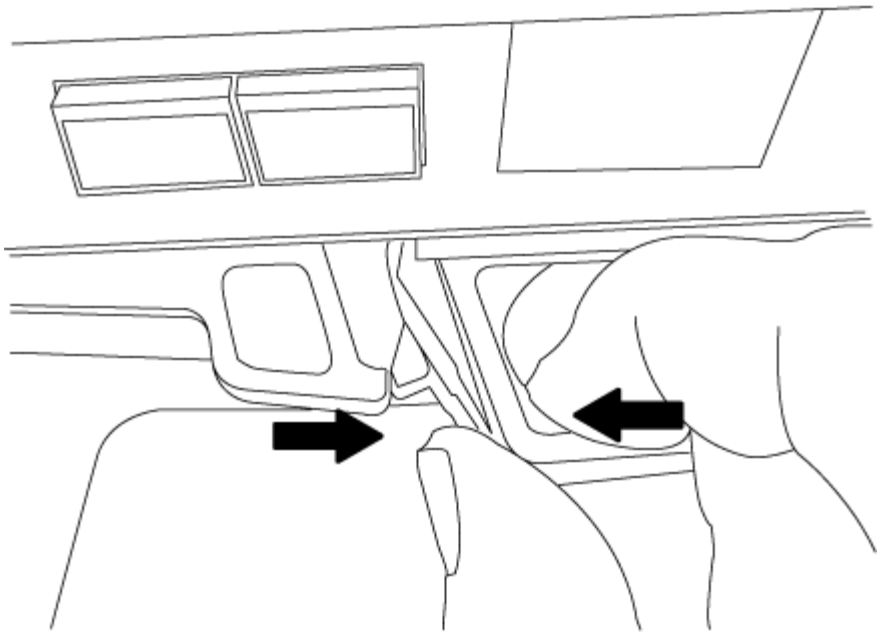
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

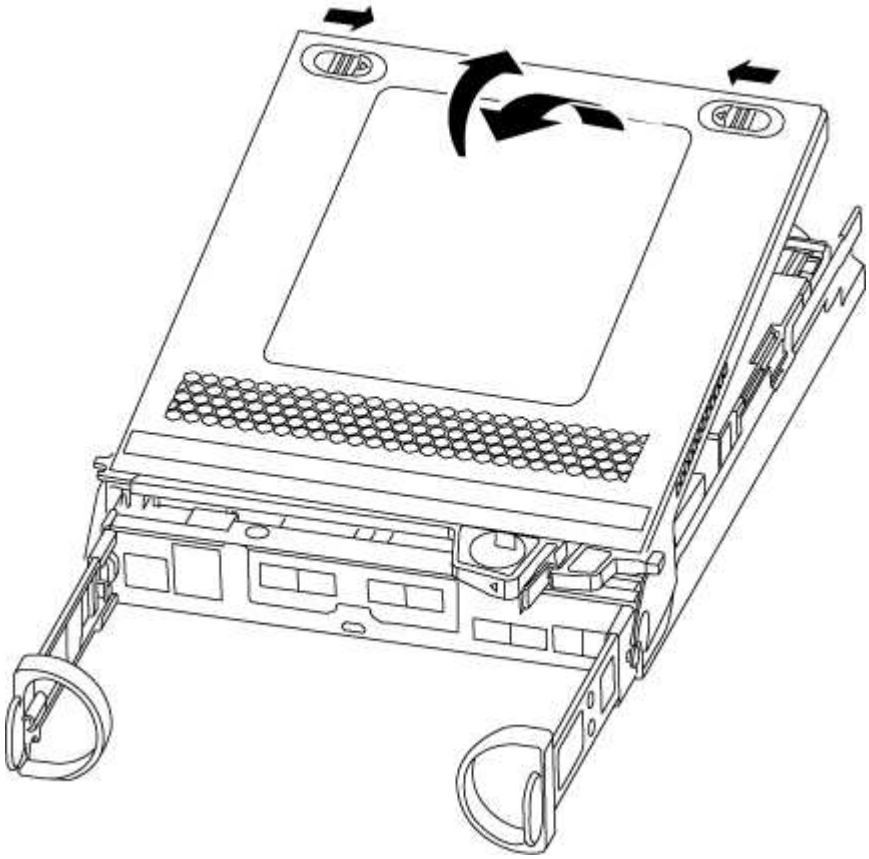
3. Remove and set aside the cable management devices from the left and right sides of the controller module.



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.

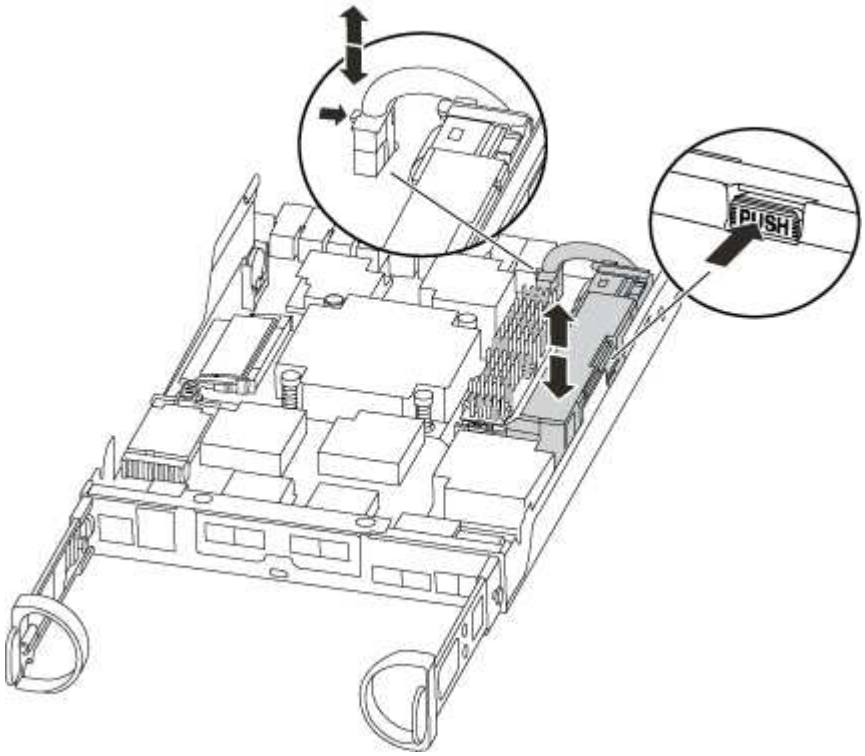


The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - 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.



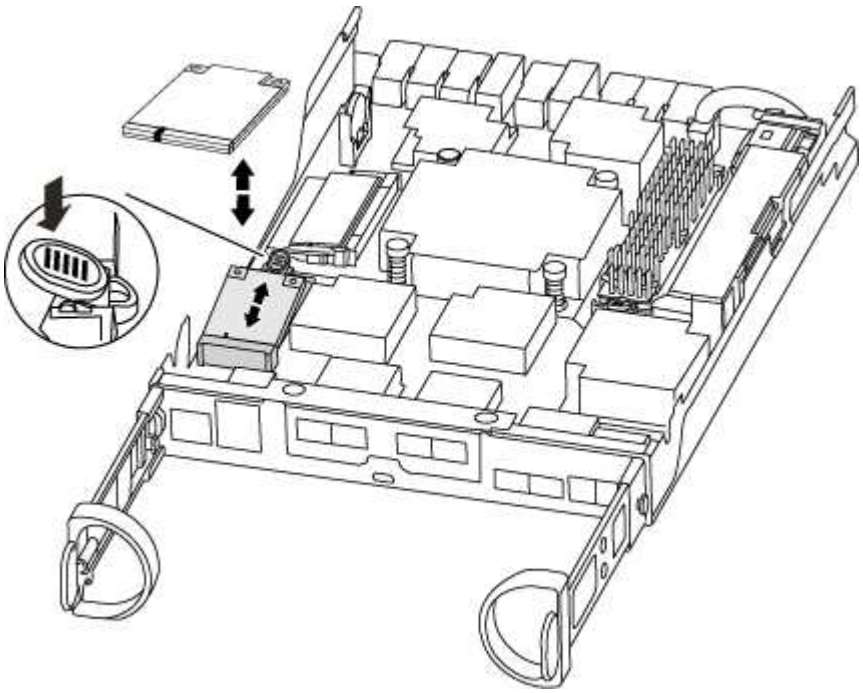
3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
5. 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

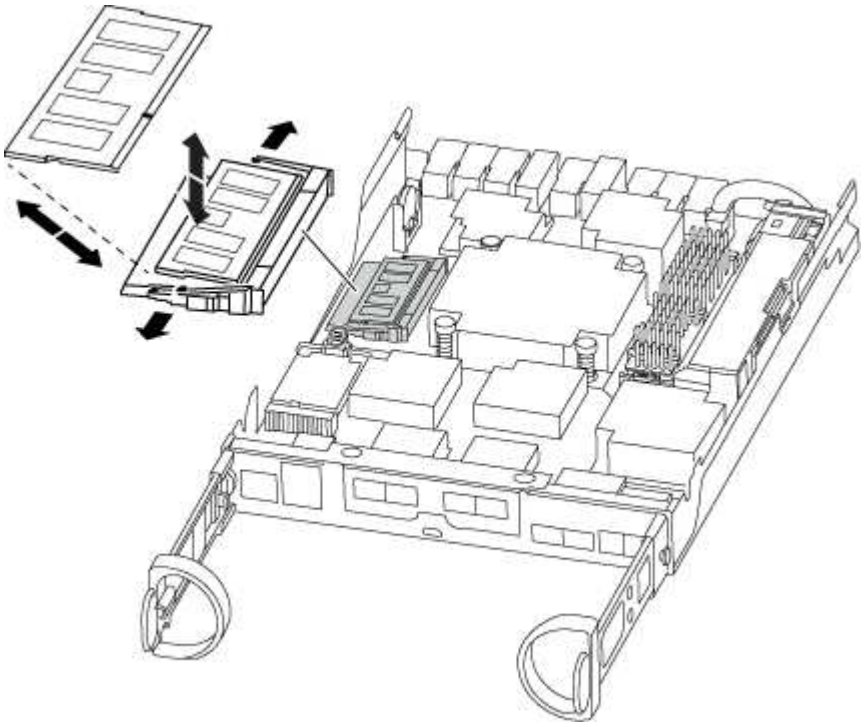
controller module to the corresponding slots in the replacement controller module.

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

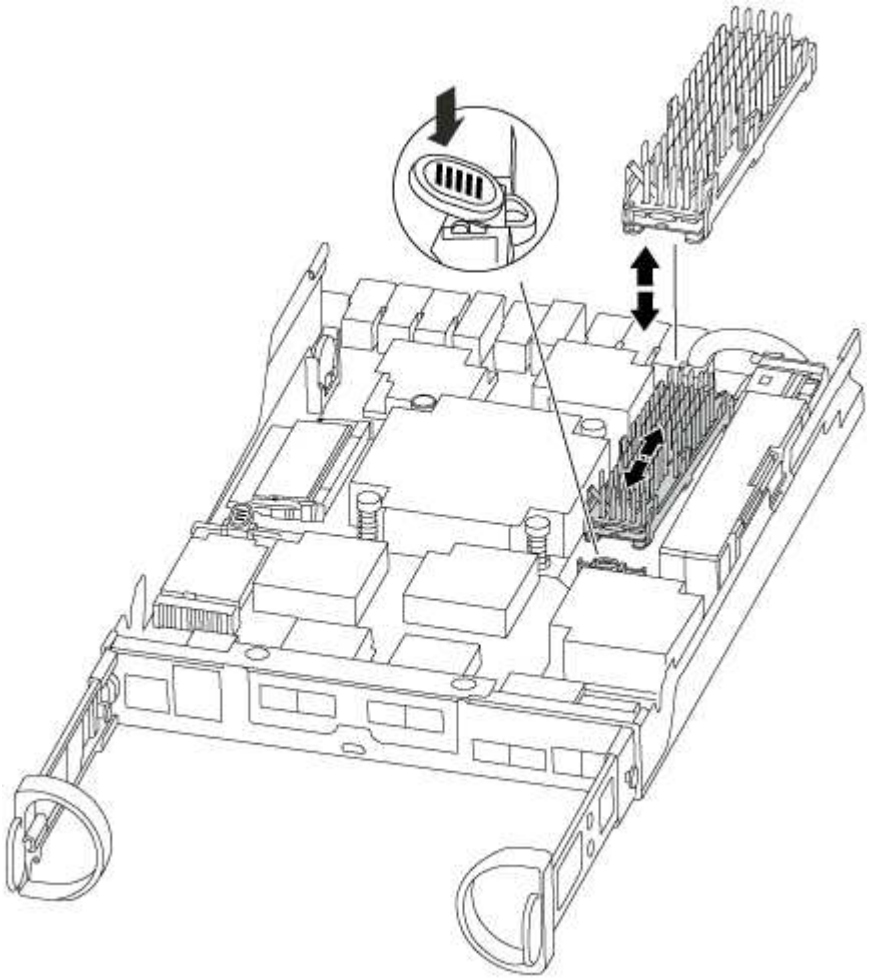
If your AFF A220 or FAS2700 system has a caching module, you need to move the caching module from the old controller module to the replacement controller module. The caching module is referred to as the “M.2 PCIe card” on the controller module label.

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.

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date` *`mm/dd/yyyy`*
5. If necessary, set the time in GMT on the replacement node: `set time` *`hh:mm:ss`*
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the controller's HA state

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 - FAS2700

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
Two-node MetroCluster configuration	Option 3: Manually reassign the system ID on systems in a two-node MetroCluster configuration

Option 1: Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:`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`

Option 3: Manually reassign the system ID on systems in a two-node MetroCluster configuration

In a two-node MetroCluster configuration running ONTAP, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems in a two-node MetroCluster configuration running ONTAP.

You must be sure to issue the commands in this procedure on the correct node:

- The *impaired* node is the node on which you are performing maintenance.
- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the DR partner of the impaired node.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by entering `Ctrl-C`, and then select the option to boot to Maintenance mode from the displayed menu.

   You must enter `Y` when prompted to override the system ID due to a system ID mismatch.

2. View the old system IDs from the healthy node: `` `metrocluster node show -fields node-systemid`,dr-partner-systemid` ``

   In this example, the Node_B_1 is the old node, with the old system ID of 118073209:

   ```
   dr-group-id cluster         node                 node-systemid dr-
   partner-systemid
    ----------- -------------------- ------------------- -------------
   ------------------
    1           Cluster_A            Node_A_1             536872914
   118073209
    1           Cluster_B            Node_B_1             118073209
   536872914
    2 entries were displayed.
   ```

3. View the new system ID at the Maintenance mode prompt on the impaired node: `disk show`

   In this example, the new system ID is 118065481:

   ```
   Local System ID: 118065481
       ...
       ...
   ```

4. Reassign disk ownership (for FAS systems), by using the system ID information obtained from the disk show command: `disk reassign -s old system ID`

In the case of the preceding example, the command is: `disk reassign -s 118073209`

You can respond `Y` when prompted to continue.

5. Verify that the disks were assigned correctly: `disk show -a`

   Verify that the disks belonging to the *replacement* node show the new system ID for the *replacement* node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK     OWNER                  POOL   SERIAL NUMBER  HOME
-------    -------------          -----  -------------  -------------
disk_name   system-1  (118065481) Pool0  J8Y0TDZC       system-1
(118065481)
disk_name   system-1  (118065481) Pool0  J8Y09DXC       system-1
(118065481)
.
.
.
```

6. From the healthy node, verify that any coredumps are saved:

   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Verify that the coredumps are saved: `system node run -node local-node-name partner savecore`

      If the command output indicates that savecore is in progress, wait for savecore to complete before issuing the giveback. You can monitor the progress of the savecore using the `system node run -node local-node-name partner savecore -s command`.</info>.

   c. Return to the admin privilege level: `set -privilege admin`

7. If the *replacement* node is in Maintenance mode (showing the *> prompt), exit Maintenance mode and go to the LOADER prompt: `halt`

8. Boot the *replacement* node: `boot_ontap`

9. After the *replacement* node has fully booted, perform a switchback: `metrocluster switchback`

10. Verify the MetroCluster configuration: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify the operation of the MetroCluster configuration in Data ONTAP:
    a. Check for any health alerts on both clusters: `system health alert show`
    b. Confirm that the MetroCluster is configured and in normal mode: `metrocluster show`
    c. Perform a MetroCluster check: `metrocluster check run`
    d. Display the results of the MetroCluster check: `metrocluster check show`
    e. Run Config Advisor. Go to the Config Advisor page on the NetApp Support Site at support.netapp.com/NOW/download/tools/config_advisor/.

       After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.
12. Simulate a switchover operation:
    a. From any node's prompt, change to the advanced privilege level: `set -privilege advanced`

       You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (*>).

    b. Perform the switchback operation with the -simulate parameter: `metrocluster switchover -simulate`
    c. Return to the admin privilege level: `set -privilege admin`

Complete system restoration - FAS2700

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 - FAS2700

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

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

```
cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

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

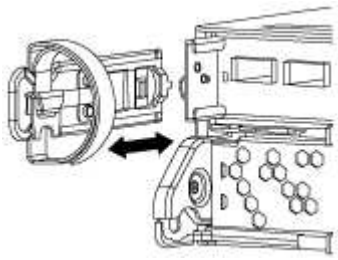
To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

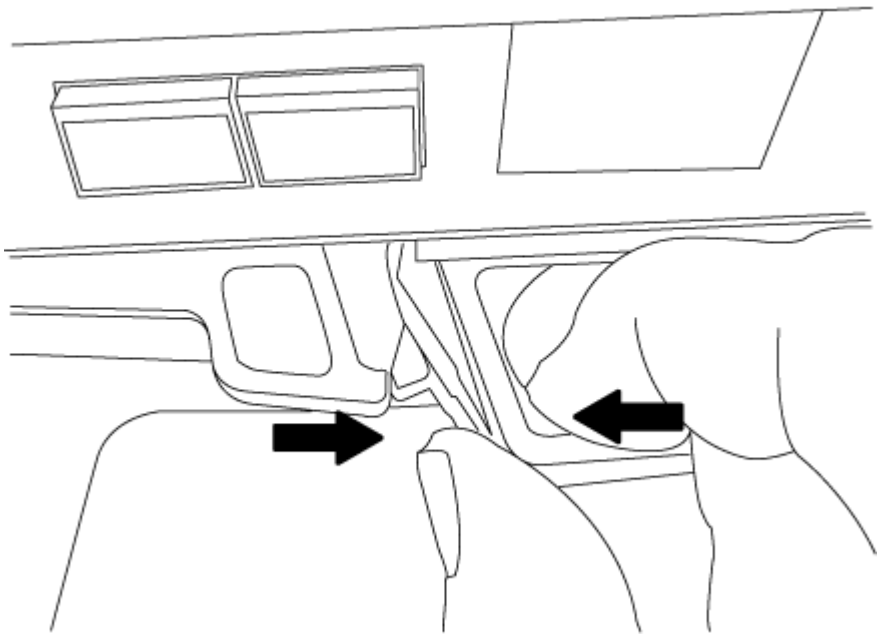
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

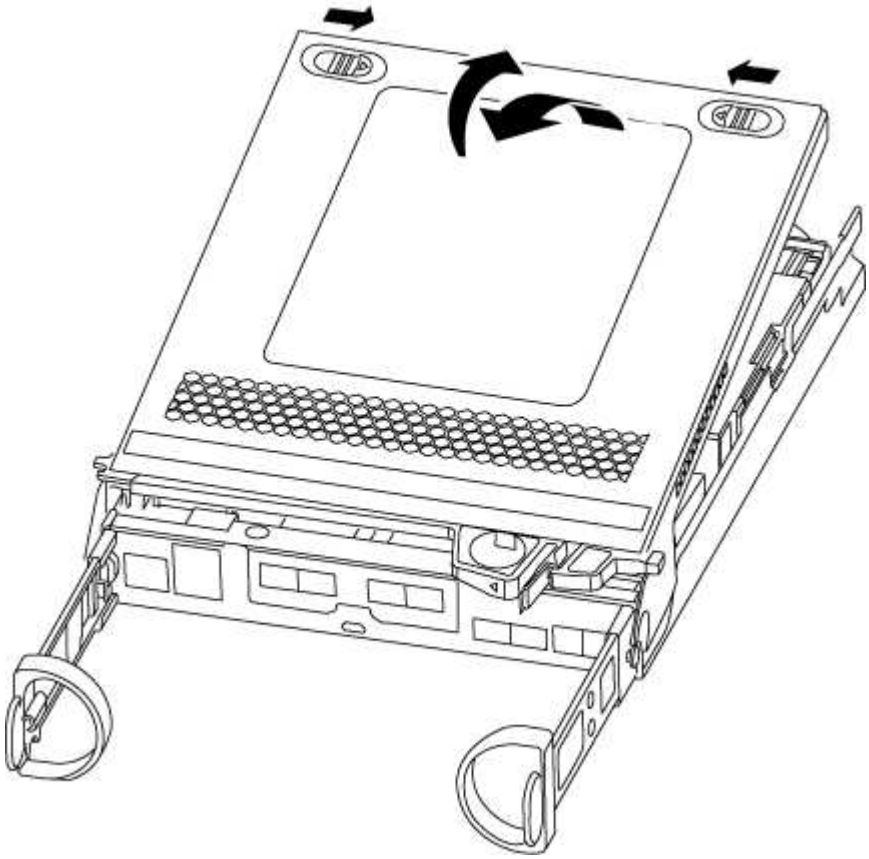
3. Remove and set aside the cable management devices from the left and right sides of the controller module.



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.

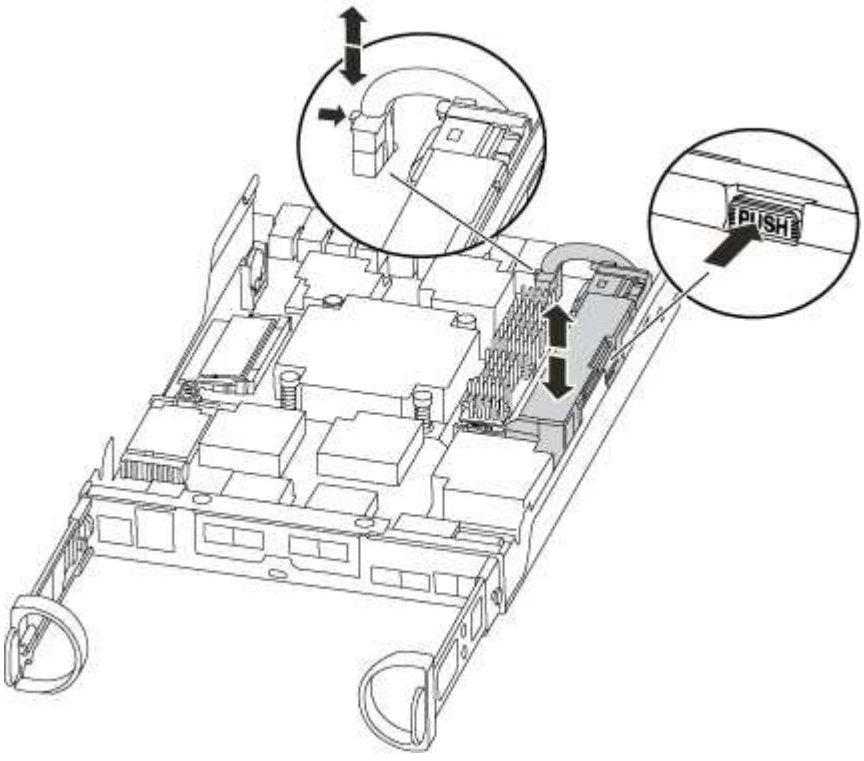
Steps

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.

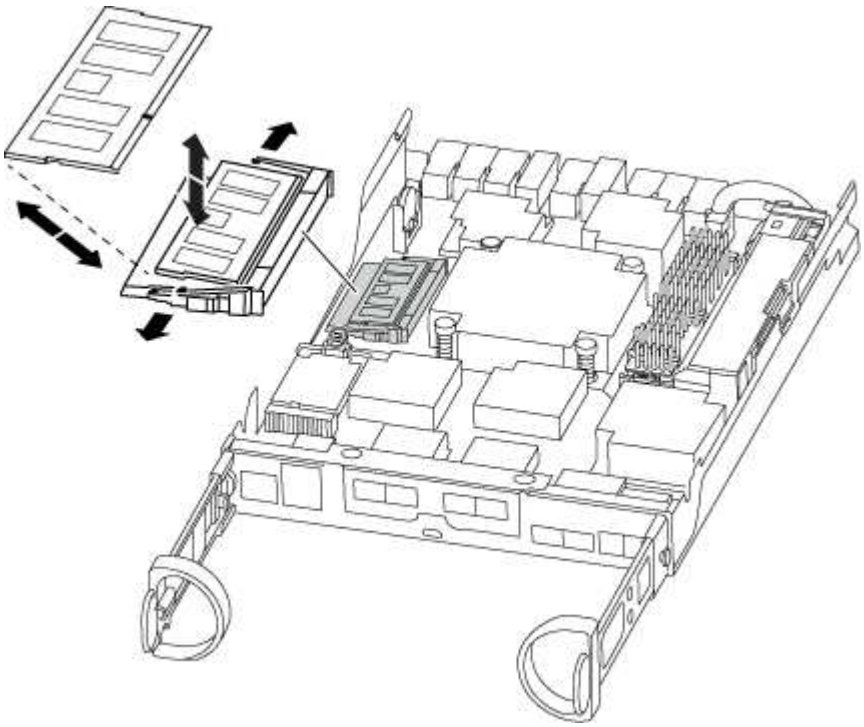
   b. Confirm that the NVMEM LED is no longer lit.
   c. Reconnect the battery connector.

5. Return to Step 3: Replace the DIMMs in 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.

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.

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 A220 and FAS2700

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 - FAS2700

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

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

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

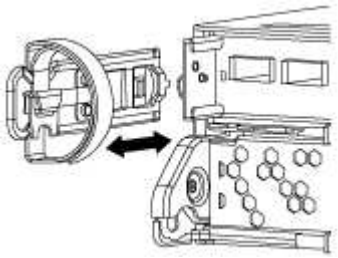
To access components inside the controller, you must first remove the controller module from the system and then remove the cover on the controller module.

Steps

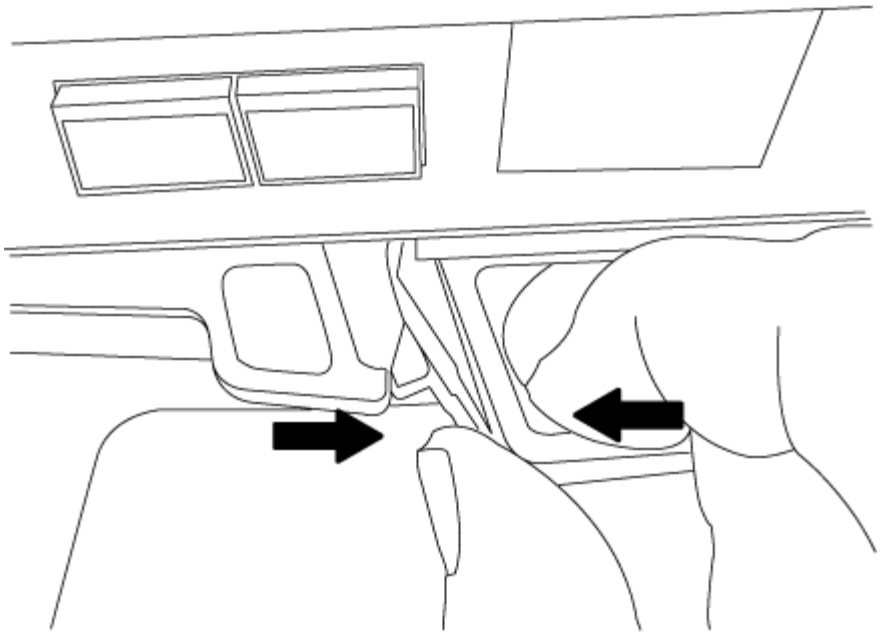
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   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.

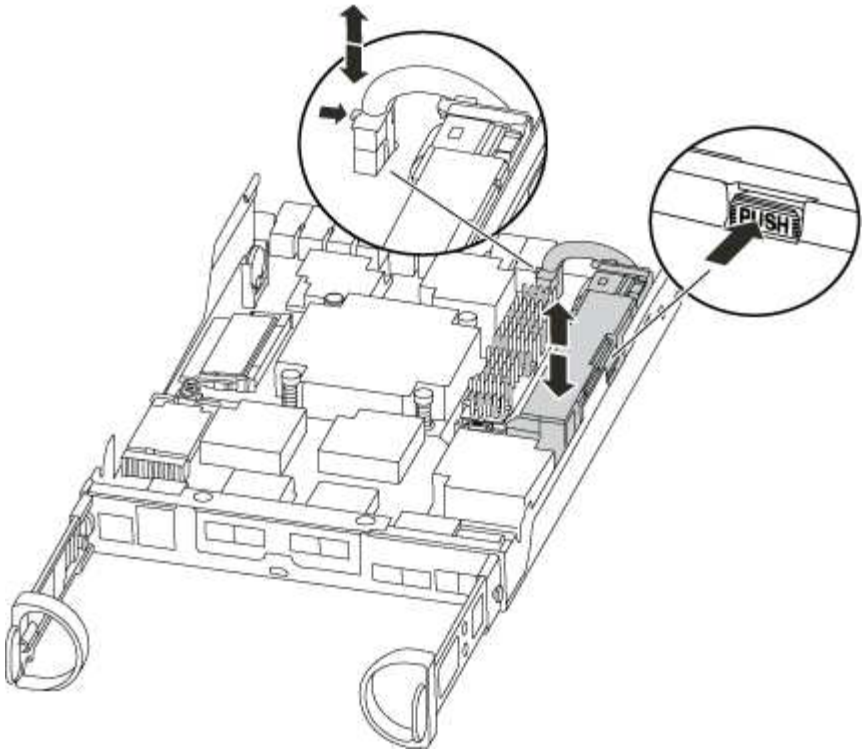


The NVRAM LED blinks while destaging contents to the flash memory when you halt the system. After the destage is complete, the LED turns off.

   - 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, and 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.

Hot-swap a power supply - FAS2700

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.

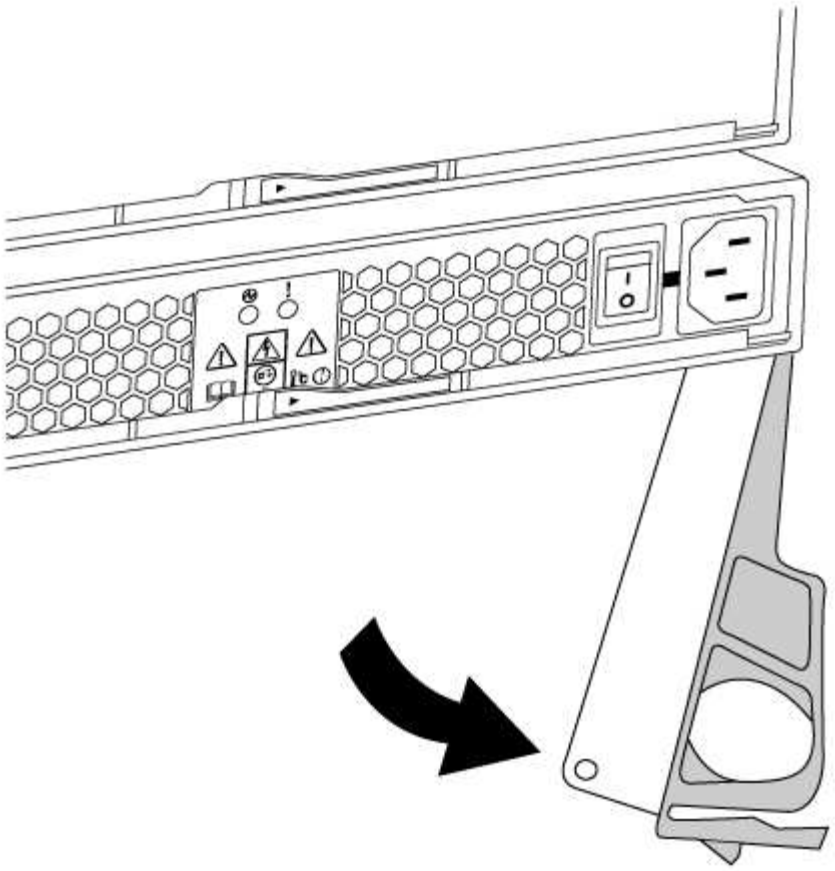
- The power supplies are redundant and hot-swappable. You do not have to shut down the controller to replace a PSU.
- This procedure is written for replacing one power supply at a time.

> 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 - FAS2700

You must use an approved RTC battery.

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

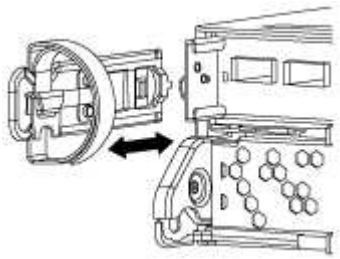
Step 2: Remove controller module

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.

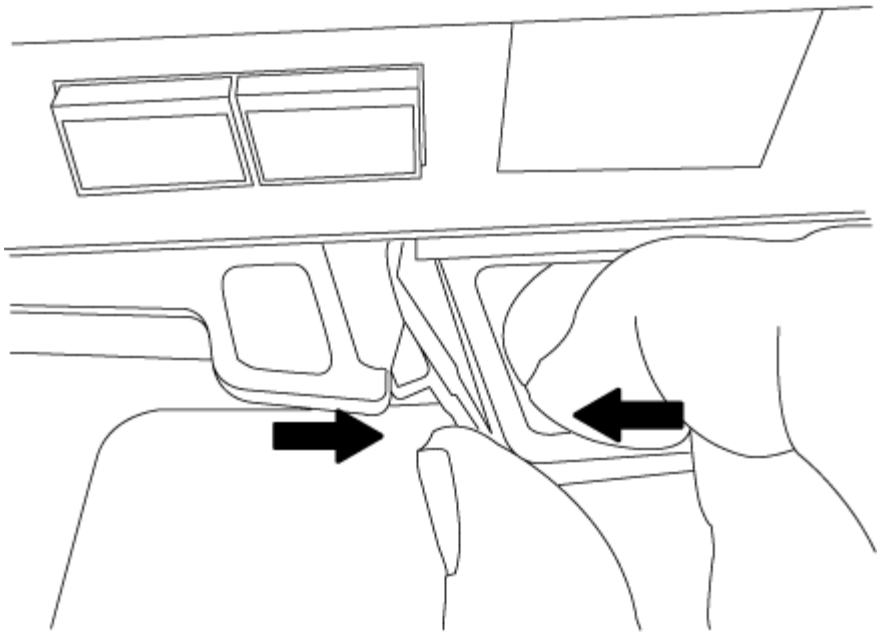
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

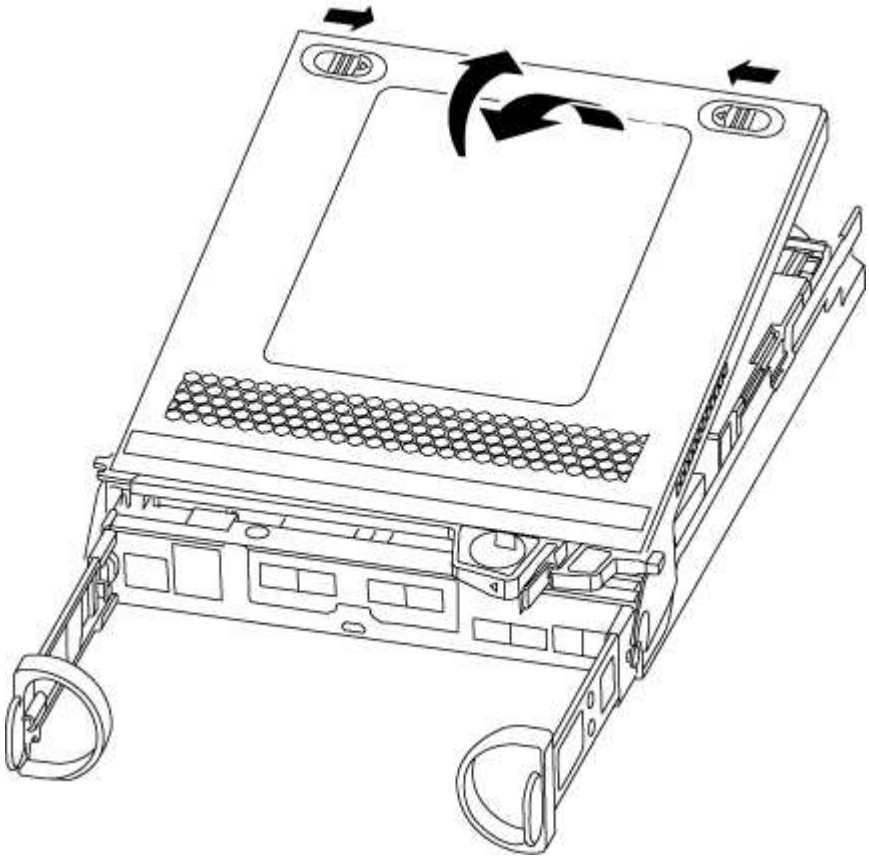
3. Remove and set aside the cable management devices from the left and right sides of the controller module.



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.



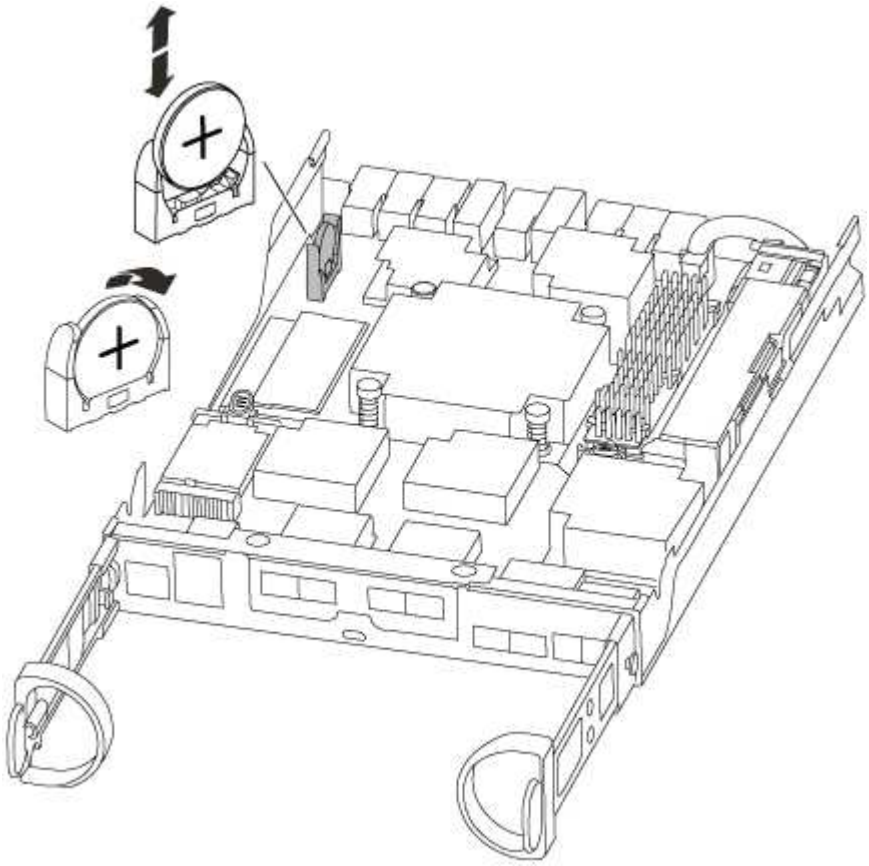
5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by sliding in the blue tabs to release the cover, and then swing the cover up and open.



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 FAS2750

The following are select specifications for the FAS2750 storage system in a single high availability pair. Visit NetApp Hardware Universe for the complete specifications for this storage system.

Key specifications for FAS2750

- Platform Configuration: FAS2750 UTA2 Single Chassis HA Pair
- Max Raw Capacity: 2.6832 PB
- Memory: 64.0000 GB
- Form Factor: 2U chassis with 2 HA controllers and 24 drive slots
- ONTAP Version: ONTAP: 9.16.1P2
- PCIe Expansion Slots: 2
- Minimum ONTAP Version: ONTAP 9.4RC1

Scaleout maximums

- Type: NAS; HA Pairs: 12; Raw Capacity: 32.2 PB / 28.6 PiB; Max Memory: 768 GB
- Type: SAN; HA Pairs: 6; Raw Capacity: 16.1 PB / 14.3 PiB; Max Memory: 384 GB
- Type: HA Pair; Raw Capacity: 2.7 PB / 2.4 PiB; Max Memory: 64.0000

I/O

Onboard I/O

- Protocol: Ethernet 10 Gbps; Ports: 12
- Protocol: FC 16 Gbps; Ports: 8
- Protocol: SAS 12 Gbps; Ports: 4

Total I/O

- Protocol: Ethernet 10 Gbps; Ports: 12
- Protocol: FC 16 Gbps; Ports: 8
- Protocol: SAS 12 Gbps; Ports: 4

Management ports

- Protocol: Ethernet 1 Gbps; Ports: 2
- Protocol: RS-232 115 Kbps; Ports: 4
- Protocol: USB 12 Mbps; Ports: 2

Storage networking supported

- CIFS
- FC
- FCoE
- iSCSI
- NFS v3
- NFS v4.0
- NFS v4.1
- 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 for external shelves

FAS2820 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.

Quick guide - FAS2820

This guide gives 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.

Access the *Installation and Setup Instructions* PDF poster:

FAS2820 Systems Installation and Setup Instructions

Video steps - FAS2820

The following video shows how to install and cable your new system.

Animation - FAS2820 Installation and setup instructions

Detailed steps - FAS2820

This procedure gives detailed step-by-step instructions for installing a typical NetApp storage system. Use this procedure if you want more detailed installation instructions.

Step 1: Prepare for installation

Before you begin

You need to provide the following at your site:

- Rack space for the storage system in either a telco rack or system cabinet.
  - 2U for the storage system
  - 2U or 4U for each drive shelf in your system

- Phillips #2 screwdriver
- Additional networking cables to connect your storage 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
  - Access to the NetApp Hardware Universe for information about site requirements as well as additional information on your configured storage system.
  - You might also want to have access to the Release Notes for your version of ONTAP 9 for your version of ONTAP for more information about this storage system.

Steps

1. Unpack all boxes and inventory the contents.

   Customers with specific power requirements must check NetApp Hardware Universe for their configuration options.

2. Access the Configure ONTAP on a new cluster with System Manager
   a. Review the requirements amd procedure steps.
   b. Gather information about your storage system by completing the setup worksheet^ (need the URL to the worksheet).
   c. Record the storage system serial number from the controllers.



The following table identifies the types of cables you might receive. If you receive a cable not listed in the table, see the NetApp Hardware Universe to locate the cable and identify its use.

Type of cable…	Part number and length	Connector type	For…
10 GbE, SFP28 cable (order dependent)	X6566B-05-R6, .5, X6566B-2-R6, 2m		Network cable
25Gb Ethernet, SFP28	X66240A-05, .5m X66240-2, 2m X66240A-5, 5m		Network cable
32Gb Fiber Channel, SFP+ (target/initiator)	X66250-2, 2m X66250-5, 5m X66250-15, 15m		FC network

Type of cable…	Part number and length	Connector type	For…
Cat 6, RJ-45 (order dependent)	X6561-R6 X6562-R6		Management network and Ethernet data
Storage	X66030A, 0.5m X66031A, 1m X66032A, 2m		Storage
USB-C console cable	No part number label		Console connection during software setup on non-Windows or Mac laptop/console
Power cables	No part number label		Powering up the storage system
Optional FC cable	Optional FC cable		Additional FC network cable

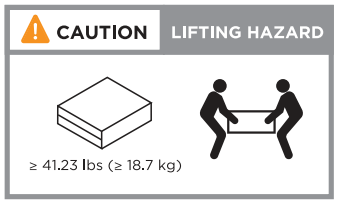
Step 2: Install the hardware

Install your storage system in a telco rack or NetApp storage system cabinet, as applicable.

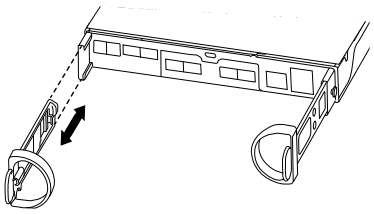
Steps

1. Install the rail kits, as needed.
2. Install and secure your storage system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the storage system.



3. Attach cable management devices (as shown).



4. Place the bezel on the front of the storage system.

Step 3: Cable controllers to your network

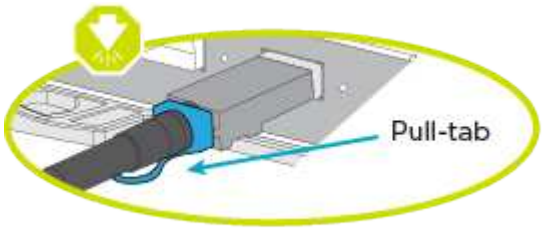
Cable the controllers to your network as either a two-node switchless cluster or a switched cluster.

The following table identifies the cable type with the call out number and cable color in the illustrations for both two-node switchless cluster and switched cluster.

Cabling	Connection type
1	Cluster interconnect
2	Management network switch
3	Host network switches

Before you begin

- Contact your network administrator for information about connecting the storage system to the switches.
- 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 the cable head over and try again.
  - If connecting to an optical switch, insert the SFP into the controller port before cabling to the port.



You can use either the applicable animation or detailed steps in the table to cable your controllers to your network.

Animation - Cabling a two-node switchless cluster cabling

Animation - Switched cluster cabling

Option 1: Cable a two-node switchless cluster

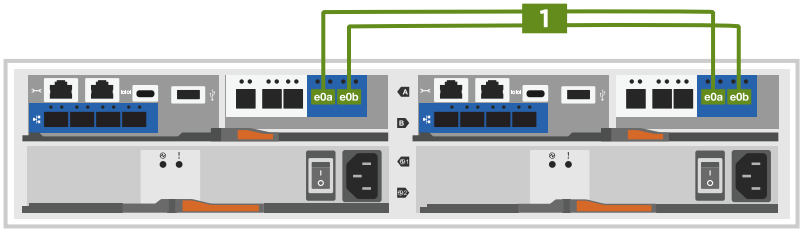
Cable your network connections and your cluster interconnect ports for a two-node switchless cluster.

Steps

1. Cable the cluster interconnect ports e0a to e0a and e0b to e0b with the cluster interconnect cable:



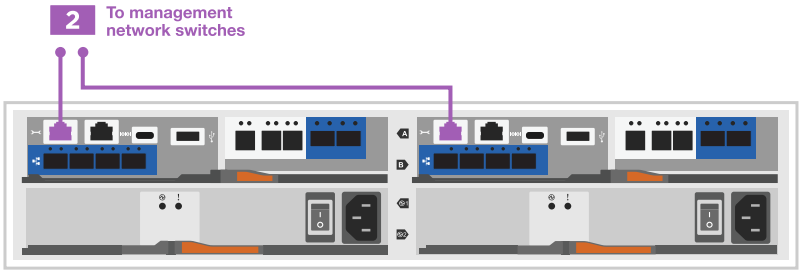
Cluster interconnect cables



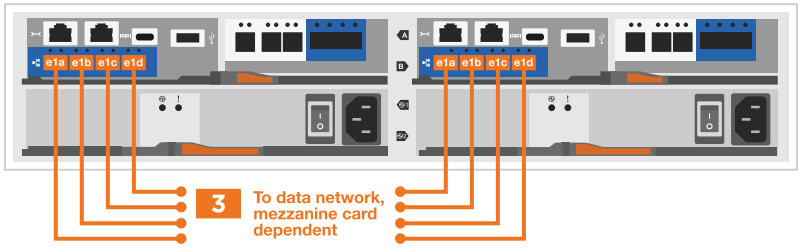
2. Cable the e0M ports to the management network switches with the RJ45 cables:



RJ45 cables

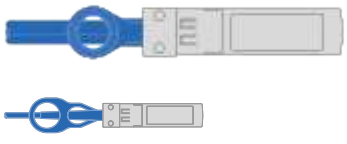


3. Cable the mezzanine card ports to your host network.



   a. If you have a 4-port Ethernet data network, cable ports e1a through e1d to your Ethernet data network.

      - 4-ports, 10/25Gb Ethernet, SFP28



- 4-ports, 10GBASE-T, RJ45



b. If you have a 4-port Fiber Channel data network, cable ports 1a through 1d for your FC network.

- 4-ports, 32Gb Fiber Channel, SFP+ (target only)



- 4-ports, 32Gb Fiber Channel, SFP+ (initiator/target)



c. If you have a 2+2 card (2 ports with Ethernet connections and 2 ports with Fiber Channel connections), cable ports e1a and e1b to your FC data network and ports e1c and e1d to your Ethernet data network.

- 2-ports, 10/25Gb Ethernet (SFP28) + 2-ports 32Gb FC (SFP+)



DO NOT plug in the power cords.

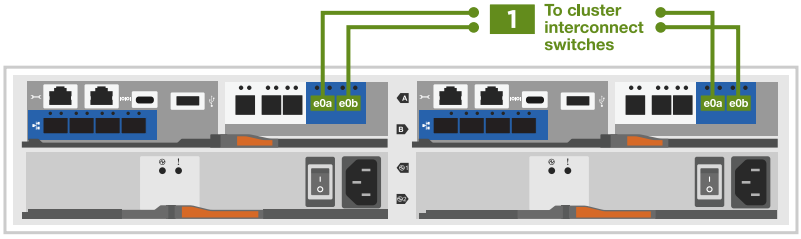
Option 2: Cable a switched cluster

Cable your network connections and your cluster interconnect ports for a switched cluster.

Steps

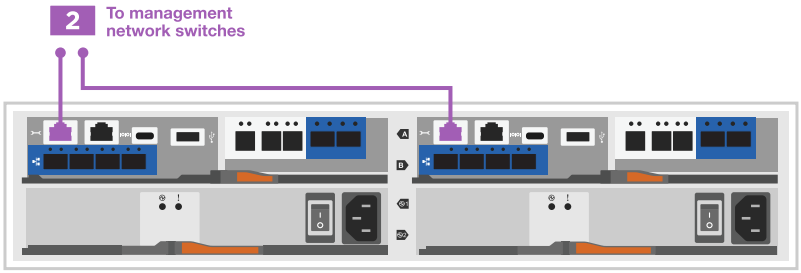
1. Cable the cluster interconnect ports e0a to e0a and e0b to e0b with the cluster interconnect cable:



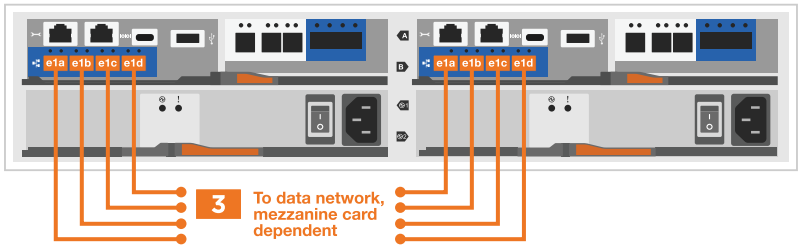


2. Cable the e0M ports to the management network switches with the RJ45 cables:





3. Cable the mezzanine card ports to your host network.



a. If you have a 4-port Ethernet data network, cable ports e1a through e1d to your Ethernet data network.

   - 4-ports, 10/25Gb Ethernet, SFP28



   - 4-ports, 10GBASE-T, RJ45



b. If you have a 4-port Fiber Channel data network, cable ports 1a through 1d for your FC network.

   - 4-ports, 32Gb Fiber Channel, SFP+ (target only)



   - 4-ports, 32Gb Fiber Channel, SFP+ (initiator/target)



c. If you have a 2+2 card (2 ports with Ethernet connections and 2 ports with Fiber Channel connections), cable ports e1a and e1b to your FC data network and ports e1c and e1d to your Ethernet data network.

   - 2-ports, 10/25Gb Ethernet (SFP28) + 2-ports 32Gb FC (SFP+)



DO NOT plug in the power cords.

Step 4: Cable controllers to drive shelves

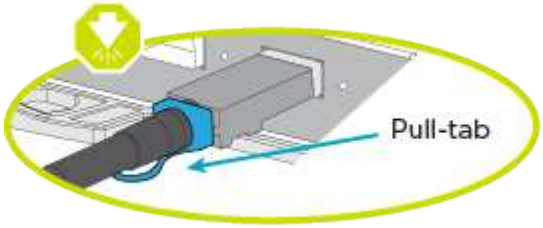
Cable your controllers to external storage.

The following table identifies the cable type with the call out number and cable color in the illustrations for cabling your drive shelves to your storage system.

The example uses DS224C. Cabling is similar with other supported drive shelves. See Install and cable shelves for a new system installation - shelves with IOM12/IOM12B modules for more information.

Cabling	Connection type
1	Shelf-to-shelf cabling
2	Controller A to the drive shelves
3	Controller B to the drive shelves

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation.



About this task

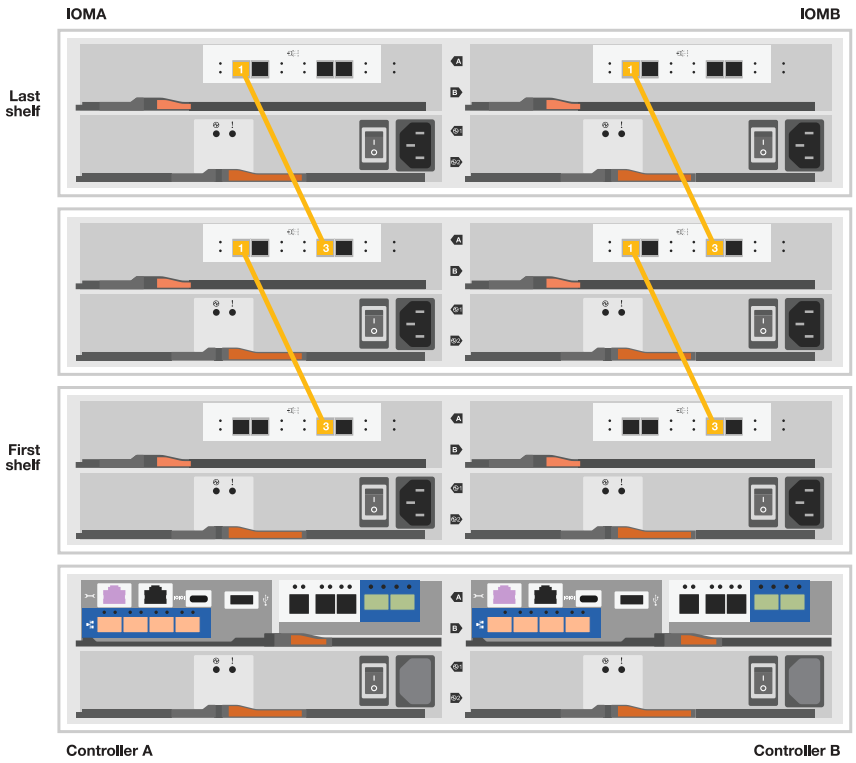
Use the animation or the step-by step instructions to complete the cabling between the controllers and to the drive shelves.

Do not use port 0b2 on a FAS2820. This SAS port is not used by ONTAP and is always disabled. See Install a shelf in a new storage system for more information.

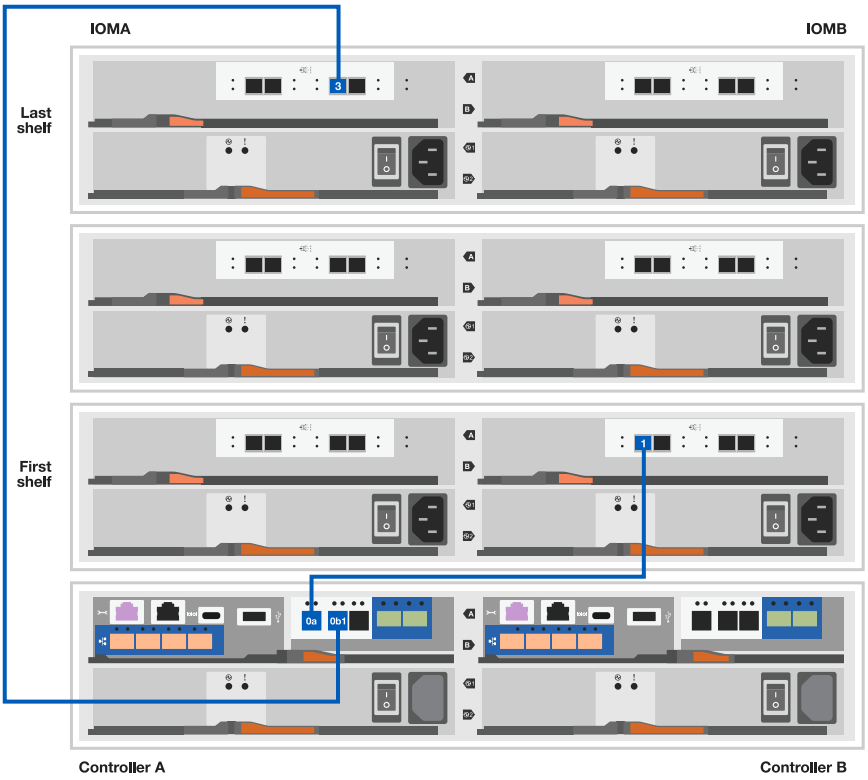
Animation - Drive shelf cabling

Steps

1. Cable the shelf-to-shelf ports.
   a. Port 1 on IOM A to port 3 on the IOM A on the shelf directly below.
   b. Port 1 on IOM B to port 3 on the IOM B on the shelf directly below.



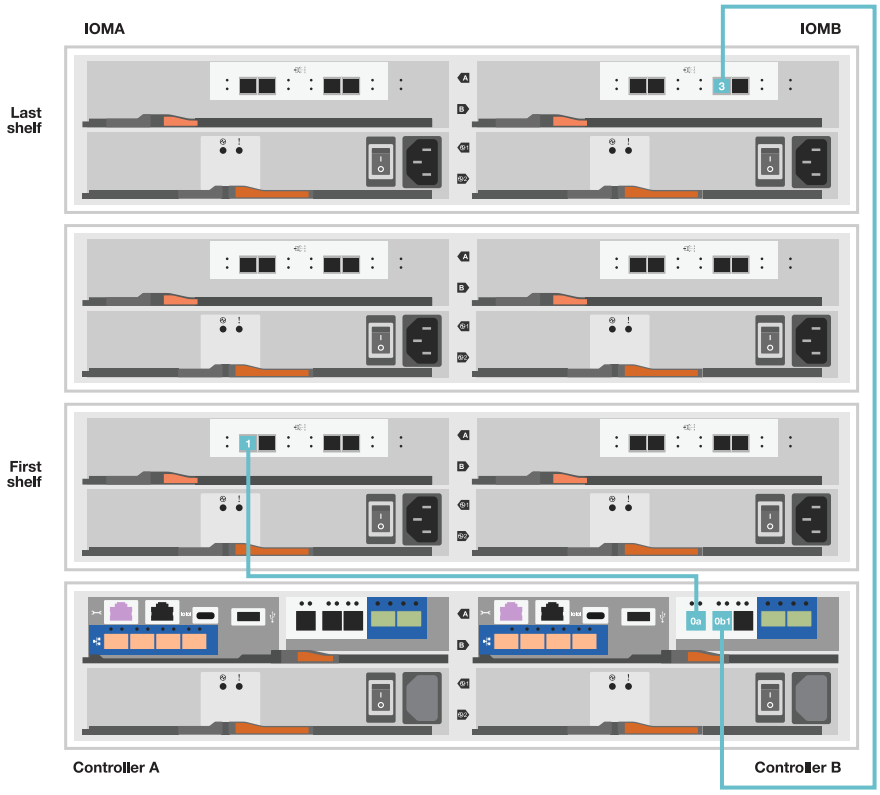
2. Cable controller A to the drive shelves.
   a. Controller A port 0a to IOM B port 1 on first drive shelf in the stack.
   b. Controller A port 0b1 to IOM A port 3 on the last drive shelf in the stack.



3. Connect controller B to the drive shelves.

a. Controller B port 0a to IOM A port 1 on first drive shelf in the stack.

b. Controller B port 0b1 to IOM B port 3 on the last drive shelf in the stack.



Step 5: Complete storage system setup and configuration

Complete your storage system setup and configuration using either Option 1: if network discovery enabled or Option 2: if network discovery is not enabled.

Use the following animation in either option where setting shelf ID is required:

Animation - Set drive shelf IDs

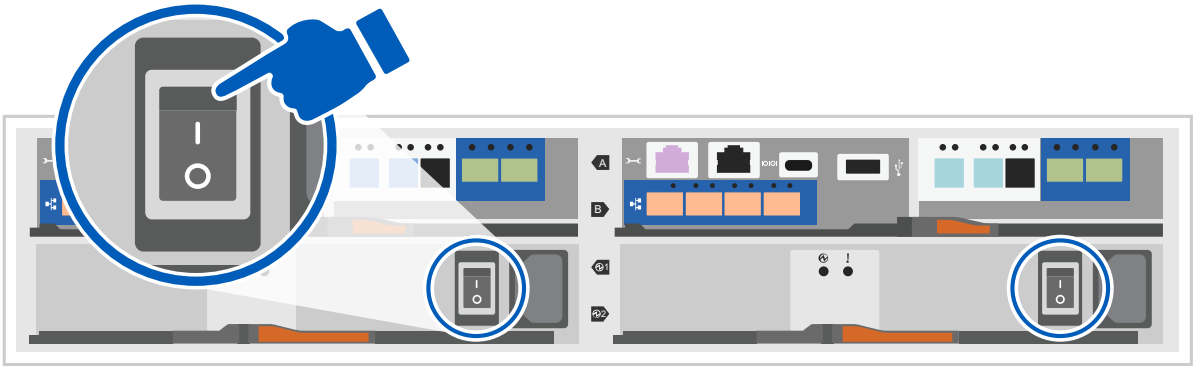
Option 1: If network discovery is enabled

If network discovery enabled on your laptop, complete storage system setup and configuration using automatic cluster discovery.

Steps

1. Turn on shelf power and set shelf IDsusing the animation at the beginning of this Step.
2. Power on the controllers
   a. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
   b. Turn on the power switches to both nodes.

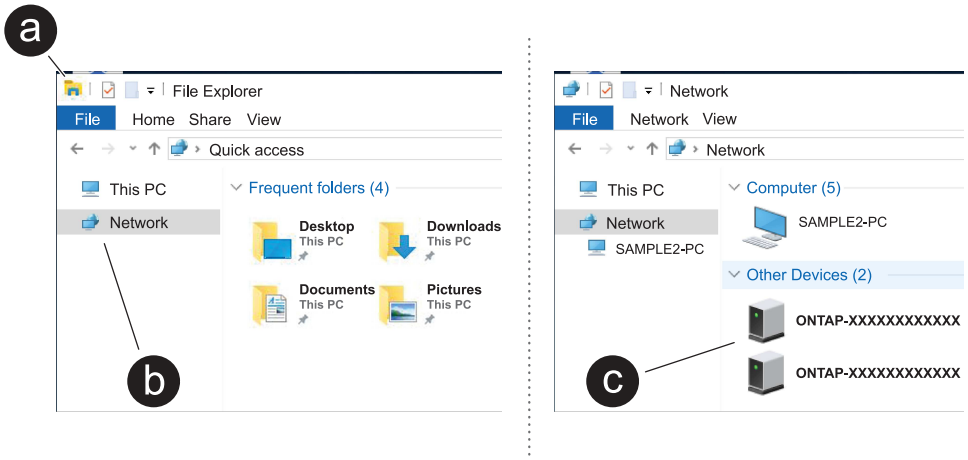
      Initial booting may take up to eight minutes.



3. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

4. Connect your laptop to the Management switch.
5. Use the graphic or steps to discover the storage system node to configure::



   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 storage system serial number for the target node.

System Manager opens.

6. Use System Manager guided setup to configure your storage system using the data you collected in Step 1: Prepare for installation.
7. Create an account or log into your account.
   a. Click mysupport.netapp.com
   b. Click *Create Account* if you need to create an account or log into your account.
8. Download and install Active IQ Config advisor
   a. Verify the health of your storage system by running Active IQ Config Advisor.
9. Register your system at https://mysupport.netapp.com/site/systems/register.
10. After you have completed the initial configuration, go to the NetApp ONTAP Resources page for information about configuring additional features in ONTAP.

Option 2: If network discovery is not enabled

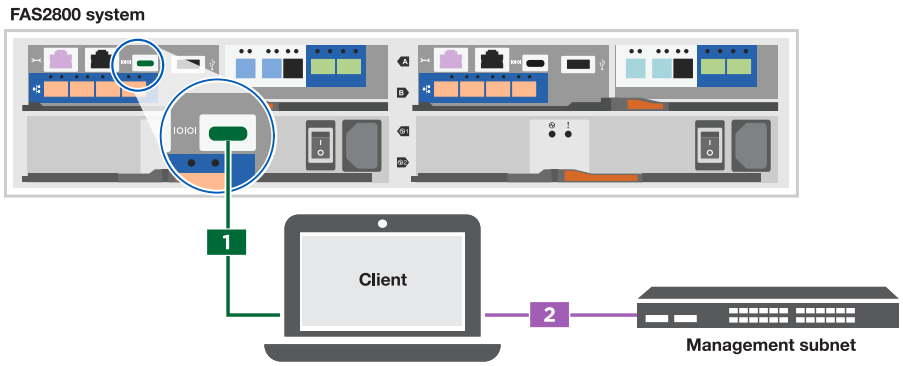
If network discovery is not enabled on your laptop, manually complete the configuration and setup.

Steps

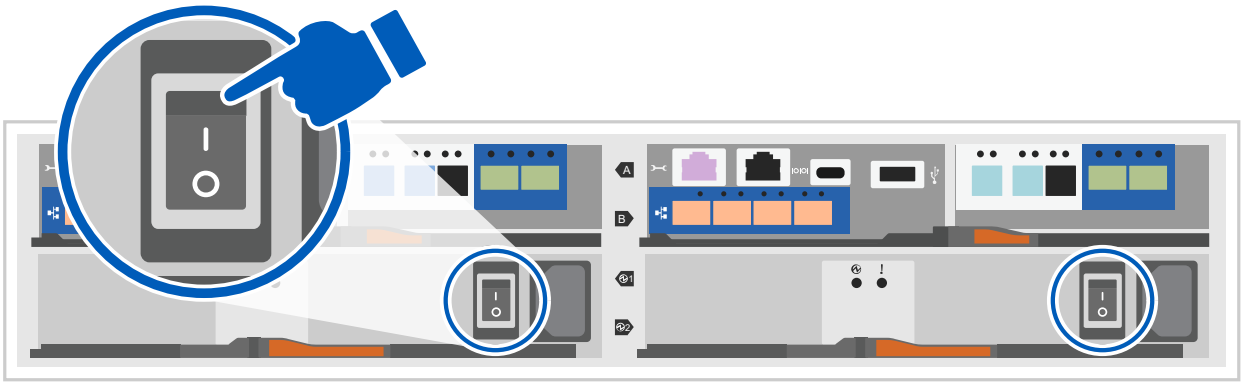
1. Cable and configure your laptop or console:
   a. Set the console port on the laptop or console to 115,200 baud with N-8-1.

      See your laptop or console's online help for 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 storage system, and then connect the laptop or console to the switch on the management subnet.



   c. Assign a TCP/IP address to the laptop or console, using one that is on the management subnet.
2. Turn on shelf power and set shelf IDs using the animation at the beginning of this Step.
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 storage system using the data you collected in Step 1: Prepare for installation..
7. Create an account or log into your account.
   a. Click mysupport.netapp.com
   b. Click *Create Account* if you need to create an account or log into your account.
8. Download and install Active IQ Config advisor
   a. Verify the health of your storage system by running Active IQ Config Advisor.
9. Register your system at https://mysupport.netapp.com/site/systems/register.
10. After you have completed the initial configuration, go to the NetApp ONTAP Resources page for information about configuring additional features in ONTAP.

Maintain

Maintain FAS2820 hardware

Maintain the hardware of your FAS2820 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the FAS2820 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the FAS2820 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Caching module	You must replace the controller's caching module when your system registers a single AutoSupport (ASUP) message that the module has gone offline.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
Drive	A drive is a device that provides the physical storage media for data.
NVMEM battery	A battery is included with the controller and preserves cached data if the AC power fails.
Mezzanine card	A Mezzanine card is an expansion card that is designed to be inserted into a specialized slot on the motherboard and holds the card I/O cards.

Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - FAS2800

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your FAS2800 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3. **Replace the boot media**

Remove the failed boot media from the System Management module and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5. **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - FAS2800

Before replacing the boot media in your FAS2800 system, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - FAS2800

Shut down the impaired controller in your FAS2800 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   `cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node impaired_node_name -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - FAS2800

The boot media in your FAS2800 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

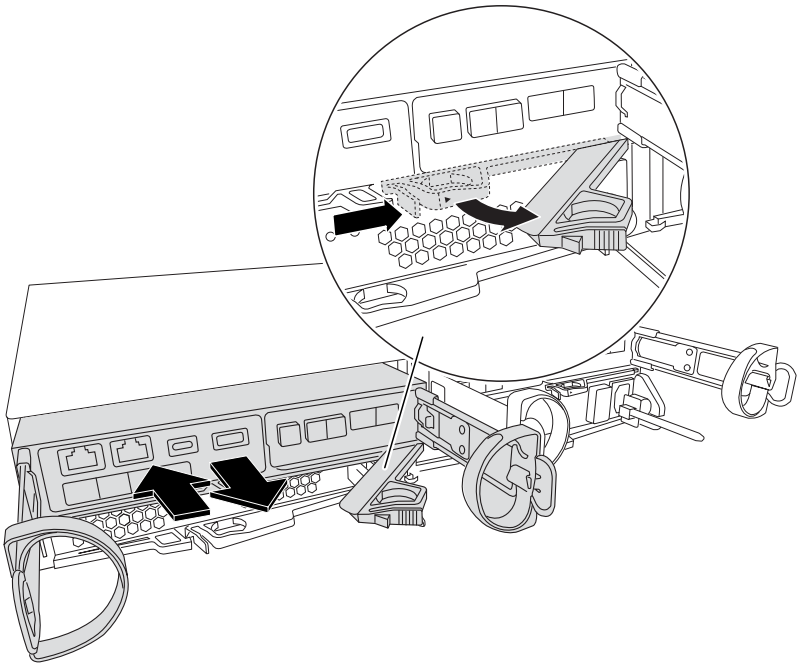
The boot media is located inside the controller module and is accessed by removing the controller module from the chassis and removing the controller module cover.module from the system.

Steps

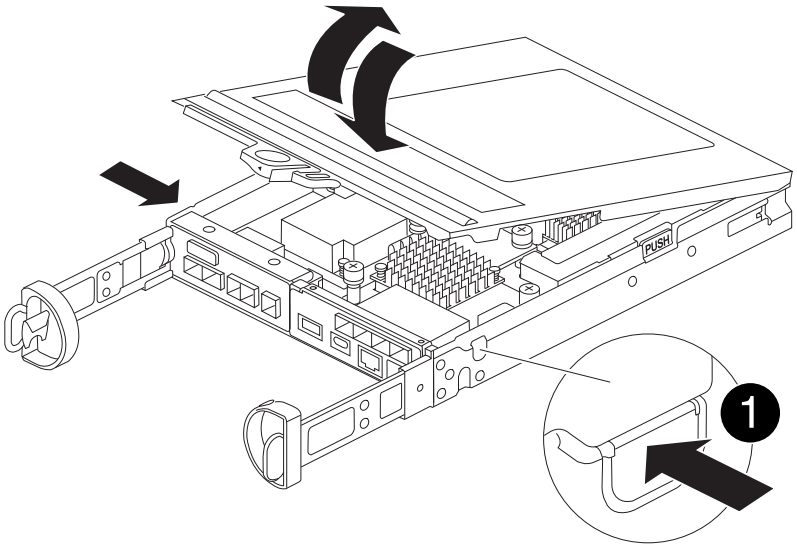
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the

system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

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.

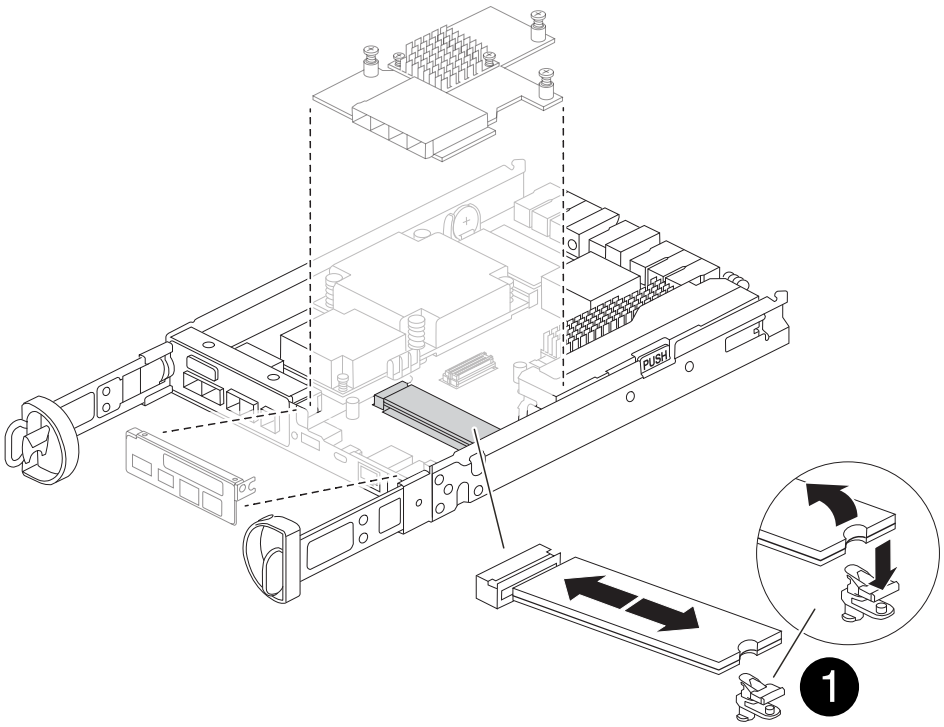


4. Turn the controller module over and place it on a flat, stable surface.
5. Open the cover by pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.



1	Controller module cover release button

6. Locate the boot media in the controller module, located under the mezzanine card and follow the directions to replace it.



1	Boot media locking tab

7. Remove the mezzanine card using the following illustration or the FRU map on the controller module:
   a. Remove the IO Plate by sliding it straight out from thr controller module.
   b. Loosen the thumbscrews on the mezzanine card.

      You can loosen the thumbscrews with your fingers or a screwdriver. If you use your fingers, you might need to rotate the NV battery up for better finger purchase on the thumbscrew next to it.

   c. Lift the mezzanine card straight up.
8. Replace the boot media:
   a. Press the blue button on the boot media housing to release the boot media from its housing, rotate the boot media up, 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.

   b. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
      Check the boot media to make sure that it is seated squarely and completely in the socket, and if necessary, remove the boot media and reseat it into the socket.
   c. Push the blue locking button, rotate the boot media all the way down, and then release the locking button to lock the boot media in place.
9. Reinstall the mezzanine card:
   a. Align the socket on the motherboard with the socket on the mezzanine card, and then gently seat the card in the socket.

b. Tighten the three thumbscrews on the mezzanine card.

c. Reinstall the IO Plate.

10. Reinstall the controller module cover and lock it into place.

11. Install the controller module:

    a. Align the end of the controller module with the opening in the chassis, and then gently push the controller module half-way into the way into the system.

    b. Recable the controller, 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 and stops at the LOADER prompt.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - FAS2800

After installing the new boot media device in your FAS2800 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- Determine your key manager type:
  - Onboard Key Manager (OKM): Requires cluster-wide passphrase and backup data
  - External Key Manager (EKM): Requires the following files from the partner node:
    - `/cfcard/kmip/servers.cfg`
    - `/cfcard/kmip/certs/client.crt`
    - `/cfcard/kmip/certs/client.key`
    - `/cfcard/kmip/certs/CA.pem`

Steps

1. From the LOADER prompt, start the boot media recovery process:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and displays one of the following messages:

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. a. Wait for the login prompt to display. b. Log into the node and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* c. Go to re-enabling automatic giveback if it was disabled.
`key manager is configured.`	Encryption is installed. Go to restoring the key manager.

If the system cannot identify the key manager configuration, it displays an error message and prompts you to confirm whether key manager is configured and which type (onboard or external). Answer the prompts to proceed.

1. Restore the key manager using the appropriate procedure for your configuration:

Onboard Key Manager (OKM)

The system displays the following message and begins running BootMenu Option 10:

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the passphrase for onboard key management when prompted.

c. Enter the passphrase again when prompted to confirm.

d. Enter the backup data for onboard key manager when prompted.

Show example of passphrase and backup data prompts

```
Enter the passphrase for onboard key management:
-----BEGIN PASSPHRASE-----
<passphrase_value>
-----END PASSPHRASE-----
Enter the passphrase again to confirm:
-----BEGIN PASSPHRASE-----
<passphrase_value>
-----END PASSPHRASE-----
Enter the backup data:
-----BEGIN BACKUP-----
<passphrase_value>
-----END BACKUP-----
```

e. Monitor the recovery process as it restores the appropriate files from the partner node.

When the recovery process is complete, the node reboots. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

f. After the node reboots, verify that the system is back online and operational.

g. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

h. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster:

```
security key-manager onboard sync
```

Go to re-enabling automatic giveback if it was disabled.

External Key Manager (EKM)

The system displays the following message and begins running BootMenu Option 11:

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. Enter the EKM configuration settings when prompted:

i. Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file:

Show example of client certificate contents

```
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----
```

ii. Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file:

Show example of client key file contents

```
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----
```

iii. Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file:

Show example of KMIP server file contents

```
-----BEGIN CERTIFICATE-----
<KMIP_certificate_CA_value>
-----END CERTIFICATE-----
```

iv. Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file:

Show example of server configuration file contents

```
xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx
xxx.xxx.xxx.xxx:5696.port=5696
xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem
xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4
1xxx.xxx.xxx.xxx:5696.timeout=25
xxx.xxx.xxx.xxx:5696.nbio=1
xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.cr
t
xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key
xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:
!RC2:!RC4:!SEED:!eNULL:!aNULL"
xxx.xxx.xxx.xxx:5696.verify=true
xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>
```

v. If prompted, enter the ONTAP Cluster UUID from the partner node. You can check the cluster UUID from the partner node using the `cluster identify show` command.

Show example of ONTAP Cluster UUID prompt

```
Notice: bootarg.mgwd.cluster_uuid is not set or is empty.
Do you know the ONTAP Cluster UUID? {y/n} y
Enter the ONTAP Cluster UUID: <cluster_uuid_value>


System is ready to utilize external key manager(s).
```

vi. If prompted, enter the temporary network interface and settings for the node:

- The IP address for the port
- The netmask for the port
- The IP address of the default gateway

Show example of temporary network setting prompts

```
In order to recover key information, a temporary network
interface needs to be
configured.

Select the network port you want to use (for example,
'e0a')
e0M

Enter the IP address for port : xxx.xxx.xxx.xxx
Enter the netmask for port : xxx.xxx.xxx.xxx
Enter IP address of default gateway: xxx.xxx.xxx.xxx
Trying to recover keys from key servers....
[discover_versions]
[status=SUCCESS reason= message=]
```

b. Verify the key restoration status:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored. The process restores the appropriate files from the partner node and reboots the node. Proceed to the next step.
- If the key is not successfully restored, the system halts and displays error and warning messages. Rerun the recovery process from the LOADER prompt: `boot_recovery -partner`

Show example of key recovery error and warning messages

```
ERROR: kmip_init: halting this system with encrypted
mroot...
WARNING: kmip_init: authentication keys might not be
available.
*****************************************************
*                 A T T E N T I O N                 *
*                                                   *
*       System cannot connect to key managers.      *
*                                                   *
*****************************************************
ERROR: kmip_init: halting this system with encrypted
mroot...
.
Terminated

Uptime: 11m32s
System halting...

LOADER-B>
```

c. After the node reboots, verify that the system is back online and operational.

d. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

Go to re-enabling automatic giveback if it was disabled.

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media to NetApp - FAS2800

If a component in your FAS2800 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot Media - manual recovery

Boot media manual recovery workflow - FAS2800

Get started with replacing the boot media in your FAS2800 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.



Review the boot media requirements

Review the requirements for replacing the boot media.



Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.



Shut down the controller

Shut down the controller when when you need to replace the boot media.



Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.



Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.



Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONATP boot menu.



Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - AFF A800

Before replacing the boot media in your AFF A800 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - FAS2820

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery process.

Step 1: Check NVE support and download the correct ONTAP image

Determine whether your ONTAP version supports NetApp Volume Encryption (NVE) so you can download the correct ONTAP image for the boot media replacement.

Steps

1. Check if your ONTAP version supports encryption:

   ```
   version -v
   ```

   If the output includes `1Ono-DARE`, NVE is not supported on your cluster version.

2. Download the appropriate ONTAP image based on NVE support:
   - If NVE is supported: Download the ONTAP image with NetApp Volume Encryption
   - If NVE is not supported: Download the ONTAP image without NetApp Volume Encryption

Download the ONTAP image from the NetApp Support Site to your HTTP or FTP server or a local folder. You will need this image file during the boot media replacement procedure.

Step 2: Verify key manager status and back up configuration

Before shutting down the impaired controller, verify the key manager configuration and back up the necessary information.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manager is configured on your system, do one of the following:

   If no key manager is configured:

   You can safely shut down the impaired controller and proceed to the shutdown procedure.

   If a key manager is configured (EKM or OKM):

   a. Enter the following query command to display the status of the authentication keys in your key manager:

      ```
      security key-manager key query
      ```

   b. Review the output and check the value in the `Restored` column. This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Complete the appropriate procedure based on your key manager type:

External Key Manager (EKM)

Complete these steps based on the value in the `Restored` column.

If all keys show `true` in the Restored column:

You can safely shut down the impaired controller and proceed to the shutdown procedure.

If any keys show a value other than `true` in the Restored column:

a. Restore the external key management authentication keys to all nodes in the cluster:

   `security key-manager external restore`

   If the command fails, contact NetApp Support.

b. Verify that all authentication keys are restored:

   `security key-manager key query`

   Confirm that the `Restored` column displays `true` for all authentication keys.

c. If all keys are restored, you can safely shut down the impaired controller and proceed to the shutdown procedure.

Onboard Key Manager (OKM)

Complete these steps based on the value in the `Restored` column.

If all keys show `true` in the Restored column:

a. Back up the OKM information:
   i. Switch to advanced privilege mode:

      `set -priv advanced`

      Enter `y` when prompted to continue.

   ii. Display the key management backup information:

      `security key-manager onboard show-backup`

   iii. Copy the backup information to a separate file or your log file.

      You will need this backup information if you need to manually recover OKM during the replacement procedure.

   iv. Return to admin mode:

      `set -priv admin`

b. You can safely shut down the impaired controller and proceed to the shutdown procedure.

If any keys show a value other than `true` in the Restored column:

a. Synchronize the onboard key manager:

```
security key-manager onboard sync
```

Enter the 32-character alphanumeric onboard key management passphrase when prompted.

> This is the cluster-wide passphrase you created when you initially configured the Onboard Key Manager. If you do not have this passphrase, contact NetApp Support.

b. Verify all authentication keys are restored:

```
security key-manager key query
```

Confirm that the `Restored` column displays `true` for all authentication keys and the `Key Manager` type shows `onboard`.

c. Back up the OKM information:

i. Switch to advanced privilege mode:

```
set -priv advanced
```

Enter `y` when prompted to continue.

ii. Display the key management backup information:

```
security key-manager onboard show-backup
```

iii. Copy the backup information to a separate file or your log file.

You will need this backup information if you need to manually recover OKM during the replacement procedure.

iv. Return to admin mode:

```
set -priv admin
```

d. You can safely shut down the impaired controller and proceed to the shutdown procedure.

What's next?

After checking the encryption key support and status on the boot media, you need to shut down the controller.

Shut down the controller for manual boot media recovery - FAS2820

Shut down or take over the impaired controller.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

b. From the LOADER prompt, enter: `printenv` to capture all boot environmental variables. Save the output to your log file.

This command may not work if the boot device is corrupted or non-functional.

What's next?

After shutting down the controller, you need to replace the boot media.

Replace the boot media and prepare for manual boot recovery - FAS2820

You must remove and open the impaired controller module, locate and replace the boot media in the controller, transfer the boot image to a USB drive, insert the USB drive in the controller, and then boot the controller.

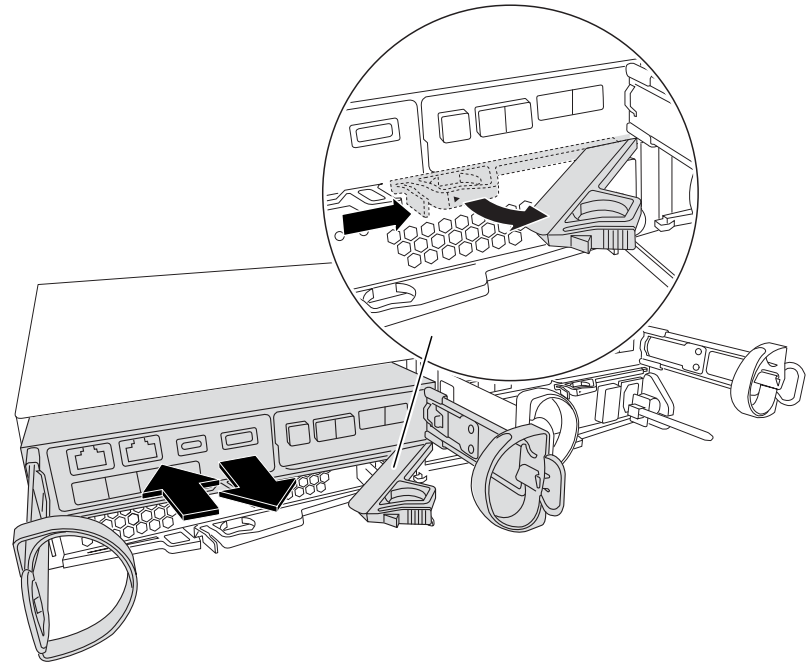
If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Remove the controller module

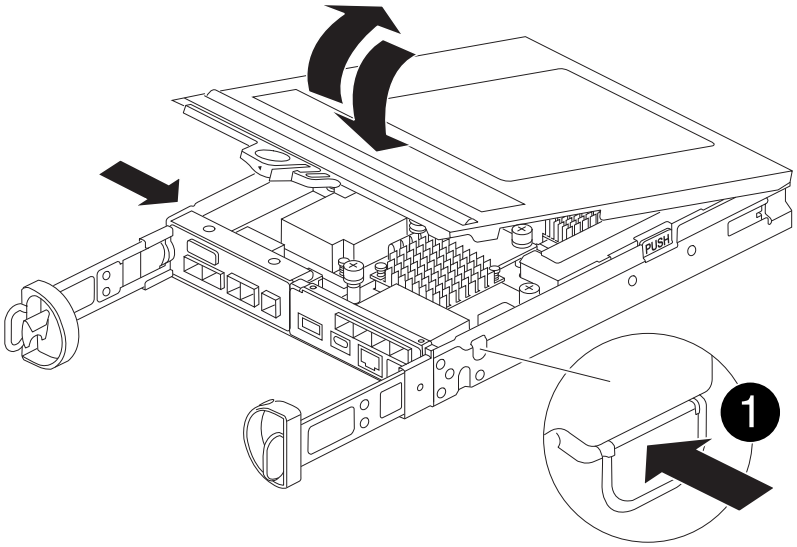
Steps

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.
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.



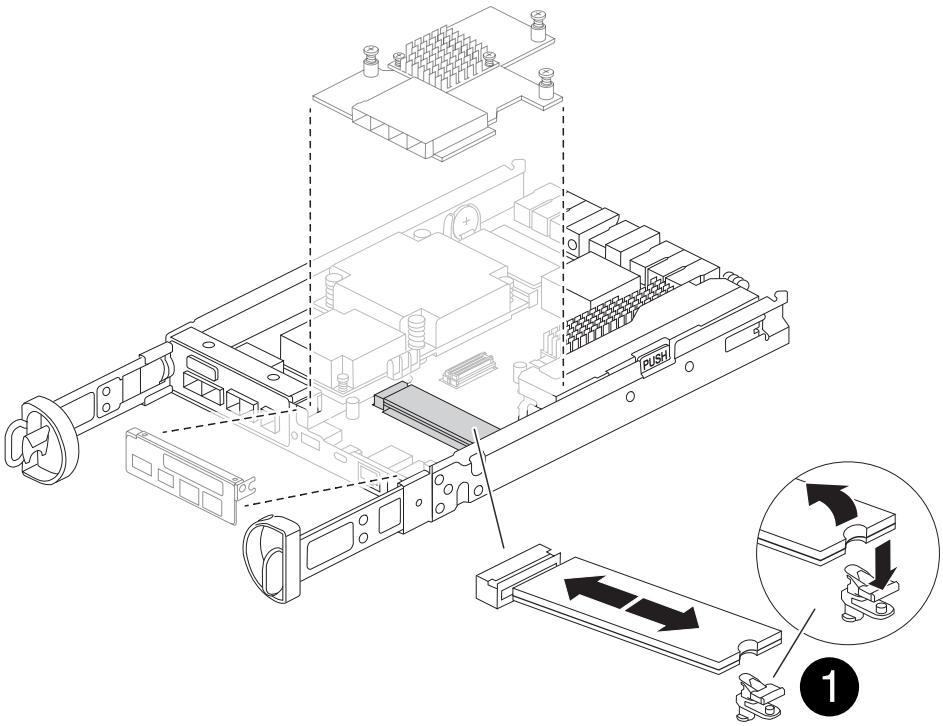
4. Turn the controller module over and place it on a flat, stable surface.
5. Open the cover by pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.



1	Controller module cover release button

Step 2: Replace the boot media

Locate the boot media in the controller module, located under the mezzanine card and follow the directions to replace it.



1	Boot media locking tab

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the mezzanine card using the following illustration or the FRU map on the controller module:
   a. Remove the IO Plate by sliding it straight out from thr controller module.
   b. Loosen the thumbscrews on the mezzanine card.

      You can loosen the thumbscrews with your fingers or a screwdriver. If you use your fingers, you might need to rotate the NV battery up for better finger purchase on the thumbscrew next to it.

   c. Lift the mezzanine card straight up.
3. Replace the boot media:
   a. Press the blue button on the boot media housing to release the boot media from its housing, rotate the boot media up, 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.

   b. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
      Check the boot media to make sure that it is seated squarely and completely in the socket, and if necessary, remove the boot media and reseat it into the socket.
   c. Push the blue locking button, rotate the boot media all the way down, and then release the locking button to lock the boot media in place.

4. Reinstall the mezzanine card:
    a. Align the socket on the motherboard with the socket on the mezzanine card, and then gently seat the card in the socket.
    b. Tighten the three thumbscrews on the mezzanine card.
    c. Reinstall the IO Plate.
5. Reinstall the controller module cover and lock it into place.

Step 3: Transfer the boot image to the boot media

Install the system image on the replacement boot media using a USB flash drive with the image installed on it. You must restore the var file system during this procedure.

Before you begin

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity.
- You must have a network connection.

Steps

1. Download the appropriate image version of ONTAP to the formatted USB flash drive:
    a. Use How to determine if the running ONTAP version supports NetApp Volume Encryption (NVE) to determine if volume encryption is currently supported.
        - If NVE is supported on the cluster, download the image with NetApp Volume Encryption.
        - If NVE is not supported on the cluster, download the image without NetApp Volume Encryption. See Which ONTAP image should I download? With or without Volume Encryption? for more details.
2. Remove the USB flash drive from your laptop.
3. Install the controller module:
    a. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
    b. Recable the controller module.

    When recabling, remember to reinstall the media converters (SFPs) if they were removed.
4. 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.
5. 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 and stops at the LOADER prompt.

What's next?

After replacing the boot media, you need to boot the recovery image.

Manual boot media recovery from a USB drive - FAS2820

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Before you begin

- Ensure your console is connected to the impaired controller.
- Verify you have a USB flash drive with the recovery image.
- Determine if your system uses encryption. You will need to select the appropriate option in step 3 based on whether encryption is enabled.

Steps

1. From the LOADER prompt on the impaired controller, boot the recovery image from the USB flash drive:

   ```
   boot_recovery
   ```

   The recovery image is downloaded from the USB flash drive.

2. When prompted, enter the name of the image or press **Enter** to accept the default image displayed in brackets.
3. Restore the var file system using the procedure for your ONTAP version:

ONTAP 9.16.0 or earlier

Complete the following steps on the impaired controller and partner controller:

a. **On the impaired controller:** Press `Y` when you see `Do you want to restore the backup configuration now?`

b. **On the impaired controller:** If prompted, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. **On the partner controller:** Set the impaired controller to advanced privilege level:

```
set -privilege advanced
```

d. **On the partner controller:** Run the restore backup command:

```
system node restore-backup -node local -target-address
impaired_node_IP_address
```

> If you see any message other than a successful restore, contact NetApp Support.

e. **On the partner controller:** Return to admin level:

```
set -privilege admin
```

f. **On the impaired controller:** Press `Y` when you see `Was the restore backup procedure successful?`

g. **On the impaired controller:** Press `Y` when you see `...would you like to use this restored copy now?`

h. **On the impaired controller:** Press `Y` when prompted to reboot, then press `Ctrl-C` when you see the Boot Menu.

i. **On the impaired controller:** Do one of the following:

- If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.
- If the system uses encryption, go to Restore encryption.

ONTAP 9.16.1 or later

Complete the following steps on the impaired controller:

a. Press `Y` when prompted to restore the backup configuration.

   After the restore procedure is successful, this message displays: `syncflash_partner: Restore from partner complete`

b. Press `Y` when prompted to confirm that the restore backup was successful.

c. Press `Y` when prompted to use the restored configuration.

d. Press `Y` when prompted to reboot the node.

e. Press `Y` when prompted to reboot again, then press `Ctrl-C` when you see the Boot Menu.

f. Do one of the following:

- If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.

- If the system uses encryption, go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -fromnode local
```

6. If you disabled automatic giveback, reenable it:

```
storage failover modify -node local -auto-giveback true
```

7. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next?

After booting the recovery image, you need to restore encryption on the boot media.

Restore encryption - FAS2820

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Complete the appropriate steps to restore encryption on your system based on your key manager type. If you are unsure which key manager your system uses, check the settings you captured at the beginning of the boot media replacement procedure.

Onboard Key Manager (OKM)

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

Ensure you have the following information available:

- Cluster-wide passphrase entered while enabling onboard key management
- Backup information for the Onboard Key Manager
- Verification that you have the correct passphrase and backup data using the How to verify onboard key management backup and cluster-wide passphrase procedure

Steps

On the impaired controller:

1. Connect the console cable to the impaired controller.
2. From the ONTAP boot menu, select the appropriate option:

ONTAP version	Select this option
ONTAP 9.8 or later	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
   ```

What's next?

After restoring encryption on the boot media, you need to return the failed part to NetApp.

Return the failed boot media to NetApp - FAS2820

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the caching module - FAS2820

You must replace the caching module in the controller module when your system registers a single AutoSupport (ASUP) message that the module has gone offline; failure to do so results in performance degradation.

- You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

If you have a 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.

ONTAP 9 System Administration Reference

You might want to erase the contents of your caching module before replacing it.

Steps

1. Although data on the caching module is encrypted, you might want to erase any data from the impaired caching module and verify that the caching module has no data:

   a. Erase the data on the caching module: `system controller flash-cache secure-erase run -node node name localhost -device-id device_number`

      Run the `system controller flash-cache show` command if you don't know the Flash Cache device ID.

   b. Verify that the data has been erased from the caching module: `system controller flash-cache secure-erase show`

      The output should display the caching module status as erased.

2. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

3. If the impaired controller is part of an HA pair, disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

4. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `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: • For an HA pair, take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`. • For a stand-alone system: `system node halt` *`impaired_node_name`*

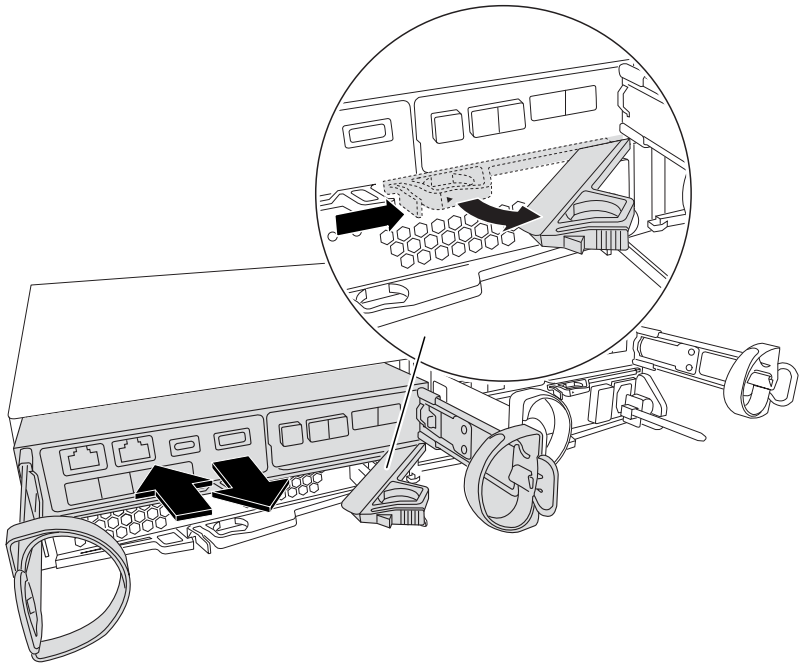
Step 2: Remove controller module

Remove the controller module from the system and then remove the cover on the controller module.

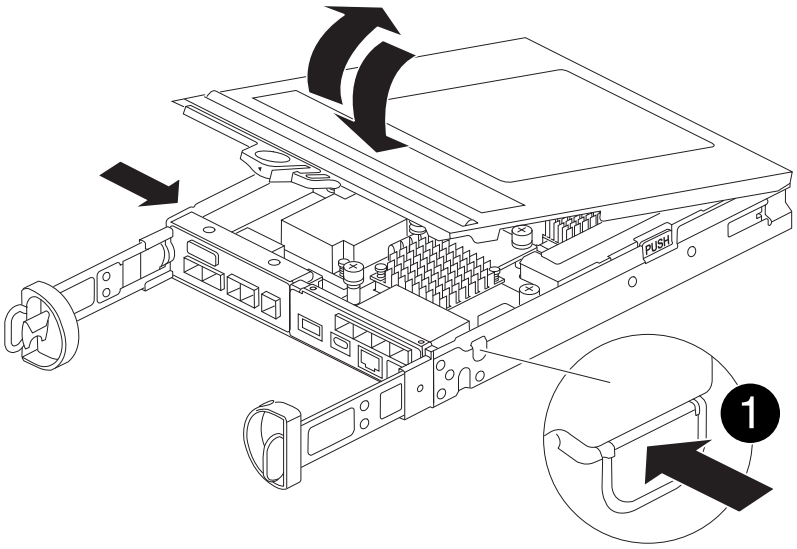
Steps

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 pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.



1	Controller module cover release button

Step 3: Replace a caching module

Locate the caching module inside the controller, remove the failed caching module and replace it.

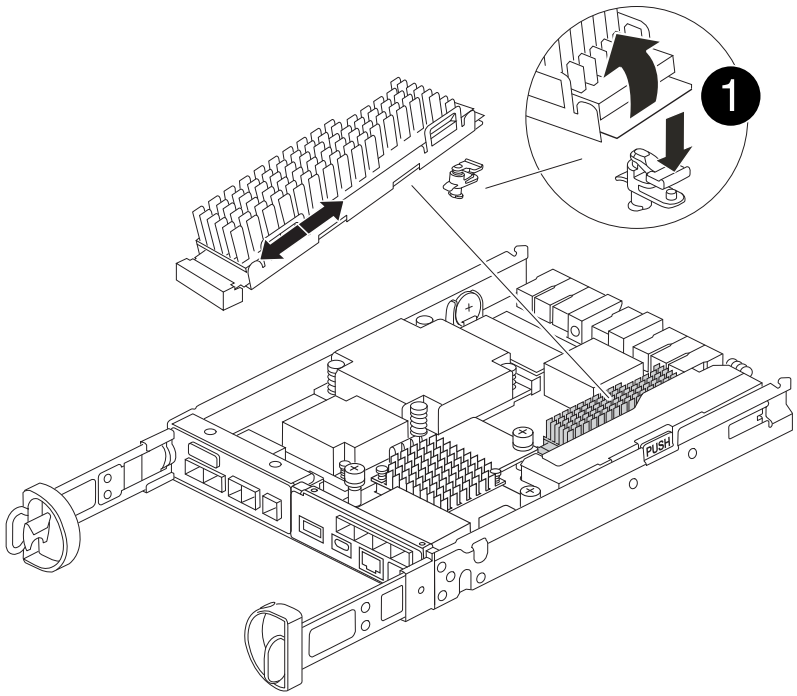
Animation - Replace the caching module

Your storage system must meet certain criteria depending on your situation:

- It must have the appropriate operating system for the caching module you are installing.
- It must support the caching capacity.
- All other components in the storage system must be functioning properly; if not, you must contact technical support.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the failed caching module near the rear of the controller module and remove it.
   a. Press the blue release button and rotate the caching module upward.
   b. Gently pull the caching module straight out of the housing.



1	Caching module release button

3. Align the edges of the replacement caching module with the socket in the housing, and then 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. Push the blue locking button, rotate the caching module all the way down, and then release the locking button to lock the caching module in place.
6. Reinstall the controller module cover and lock it into place.

Step 4: Reinstall the controller module

Reinstall the controller module 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. Turn the controller module over and align the end with the opening in the chassis.
4. Gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

5. 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.

6. 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.

      The controller begins to boot as soon as it is completely 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.

Step 5: Restore automatic giveback and AutoSupport

Restore automatic giveback and AutoSupport if they have been disabled.

1. Restore automatic giveback by using the `storage failover modify -node local -auto -giveback true` command.
2. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END`

Step 6: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - FAS2820

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 - FAS2820

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace
   chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   > For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Move and replace hardware - FAS2820

Move the power supplies, hard drives, and controller module or modules from the impaired chassis to the replacement chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the replacement 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 impaired 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.

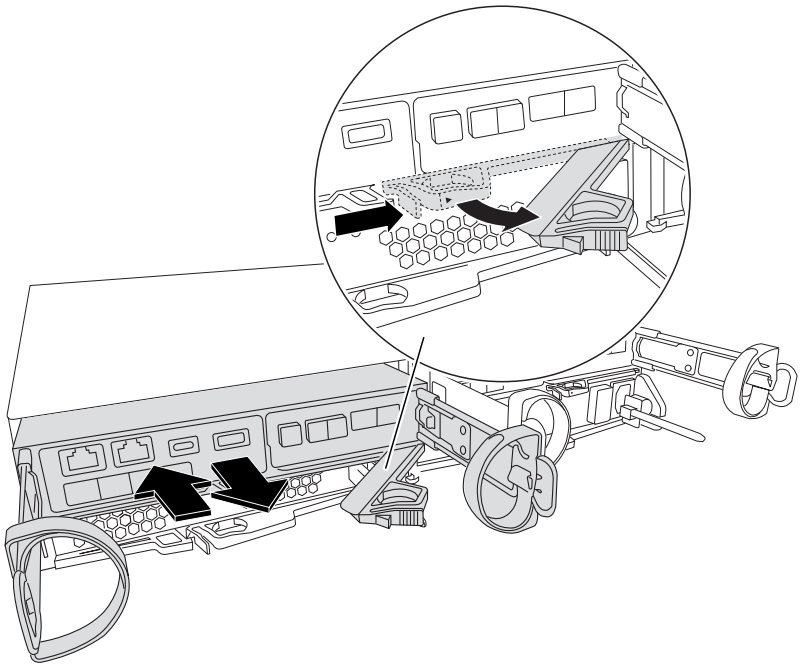
Step 2: Remove the controller module

Remove the controller module or modules from the impaired 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.
5. Repeat these steps for the second controller module in the chassis.

Step 3: Move drives to the replacement chassis

Move the drives from each drive bay opening in the impaired chassis to the same bay opening in the replacement chassis.

1. Gently remove the bezel from the front of the system.
2. Remove the drives:
   a. Press the release button on the opposite side of 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 impaired chassis with the same bay opening in the replacement chassis.
4. Gently push the drive into the chassis as far as it will go.

   The cam handle engages and begins to rotate to the closed position.

5. Firmly push the drive the rest of the way into the chassis, and then lock the cam handle by pushing it 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

Remove the existing chassis from the equipment rack or system cabinet and install the replacement chassis in the equipment rack or system cabinet.

1. Remove the screws from the chassis mount points.
2. With the help of two or three people, slide the impaired 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 impaired chassis.
7. If you have not already done so, install the bezel.

Step 5: Install the controller

Install the controller module and any other components into the replacement chassis, boot it tto Maintenance mode.

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 for the second controller in the replacement chassis.
4. Complete the installation 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. Repeat the preceding steps for the second controller module in the replacement chassis.

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 - FAS2820

Verify the HA state of the chassis bring up the system,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 based on the system's existing configuration: `ha-config modify chassis` *`ha-state`*

      The value for HA-state can be one of the following:

      - `ha`
      - `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. Exit Maintenance mode: `halt`. The LOADER prompt appears.
5. Boot the controller modules.

Step 2: Bring up the system

1. If you have not done so, plug the power cables back into the PSUs.
2. Turn on the PSUs by toggling the rocker switched to **ON**, and wait for the controllers to power up completely.
3. Check the front and the back of the chassis and controllers for any fault lights after power up.
4. Connect to the SP or BMC IP address of the nodes via SSH. This will be the same address used to shut

down the nodes.

5. Perform additional health checks as described in How_to_perform_a_cluster_health_check_with_a_script_in_ONTAP
6. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

As a best practice, you should do the following:

- Resolve any Active IQ Wellness Alerts and Risks (Active IQ will take time to process post-power up AutoSupports - expect a delay in results)
- Run Active IQ Config Advisor
- Check system health using How_to_perform_a_cluster_health_check_with_a_script_in_ONTAP

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 - FAS2820

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”).
- 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 - FAS2820

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

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_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, 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 Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
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 controller module hardware - FAS2820

Replace the impaired controller module hardware by removing the impaired controller, moving FRU components to the replacement controller module, installing the replacement controller module in the chassis, and then booting the replacement controller module.

Animation - Replace a controller module

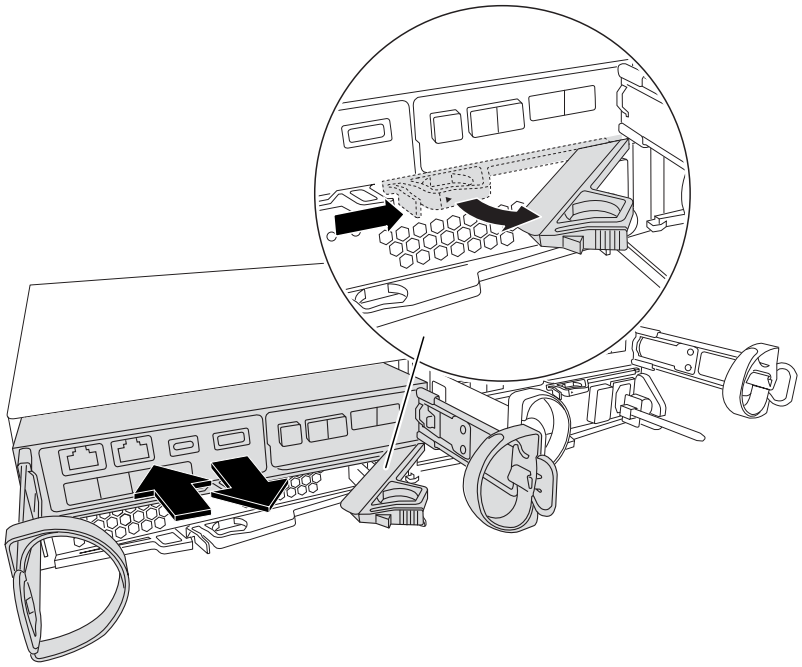
Step 1: Remove controller module

Remove the impaired controller module from the chassis.

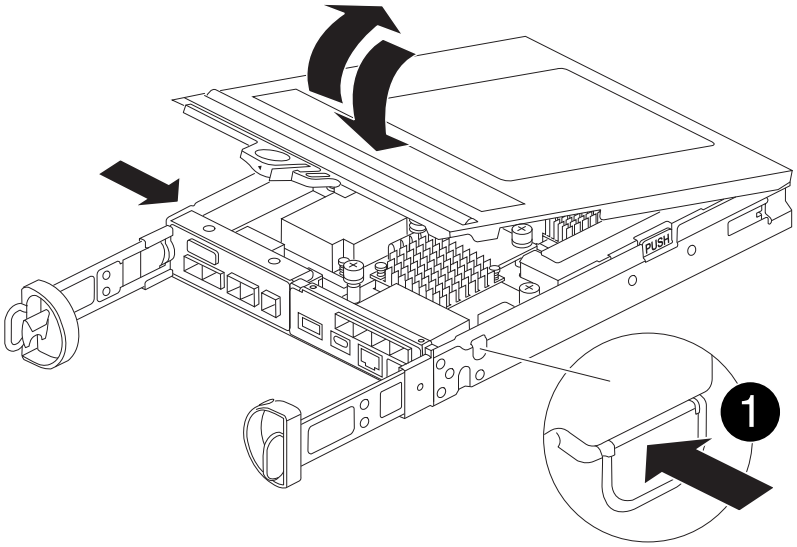
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   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 replacement 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 pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.

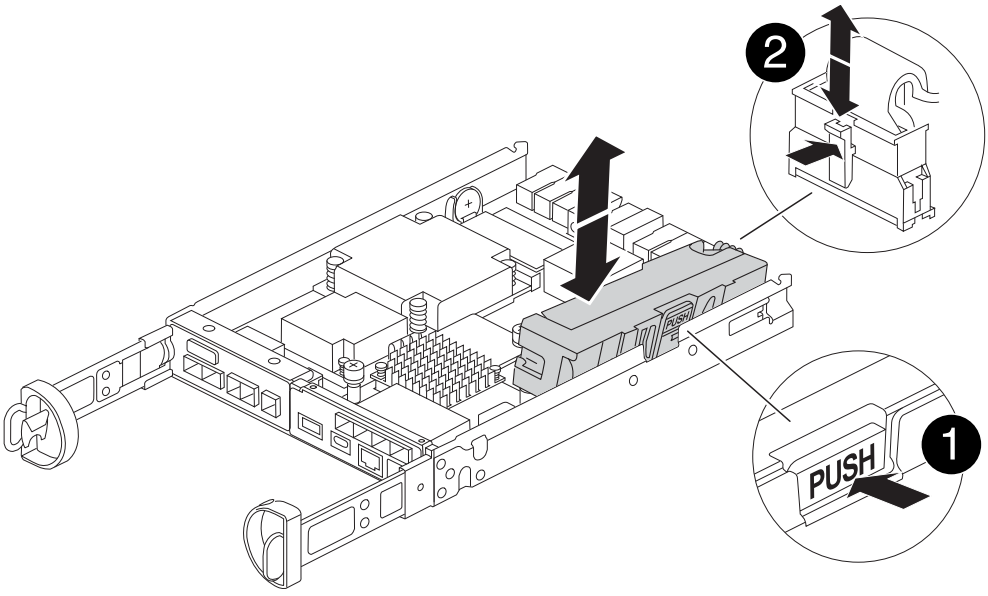


1	Controller module cover release button

Step 2: Move the NVMEM battery

Remove the NVMEM battery from the impaired controller module and install it into the replacement controller module.

Do not plug the NVMEM battery in until directed to do so.



1	NVMEM battery release button
2	NVMEM battery plug

1. Remove the battery from the controller module:

a. Press the blue button on the side of the controller module.

b. Slide the battery up until it clears the holding brackets, and then lift the battery out of the controller module.

c. Unplug the battery plug by squeezing the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.

2. Move the battery to the replacement controller module and install it:

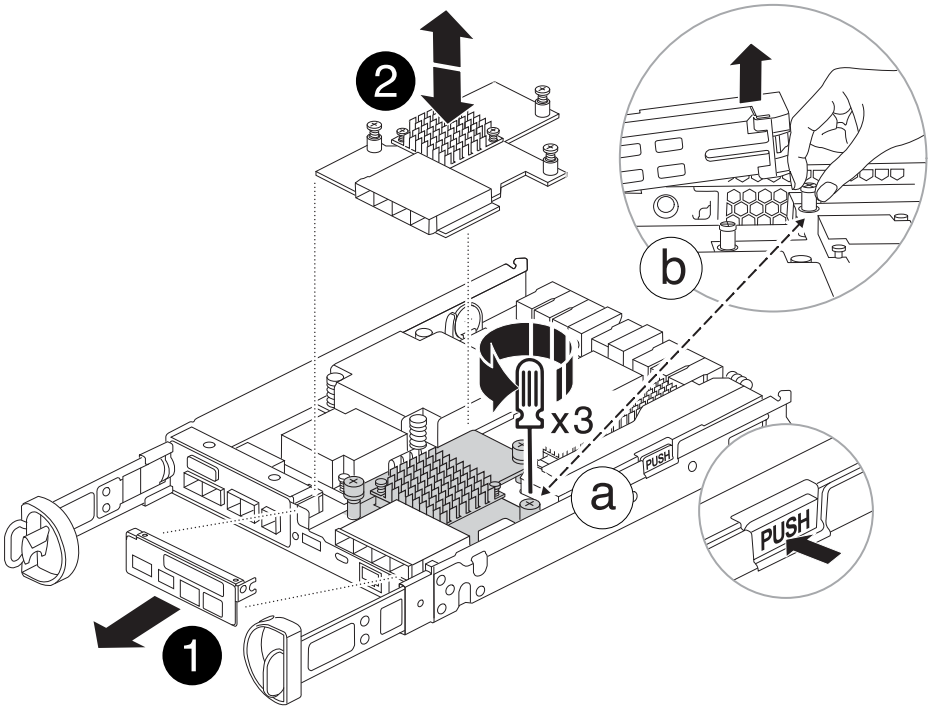
a. Aligning the battery with the holding brackets on the sheet metal side wall.

b. Slide the battery pack down until the battery latch engages and clicks into the opening on the side wall.

> Do not plug the battery in yet. You will plug it in once the rest of the components are moved to the replacement controller module.

Step 3: Remove the mezzanine card

Remove the IO Plate and PCIe mezzanine card from the impaired controller module.



1	IO Plate
2	PCIe mezzanine card

1. Remove the IO Plate by sliding it straight out from the controller module.
2. Loosen the thumbscrews on the mezzanine card.

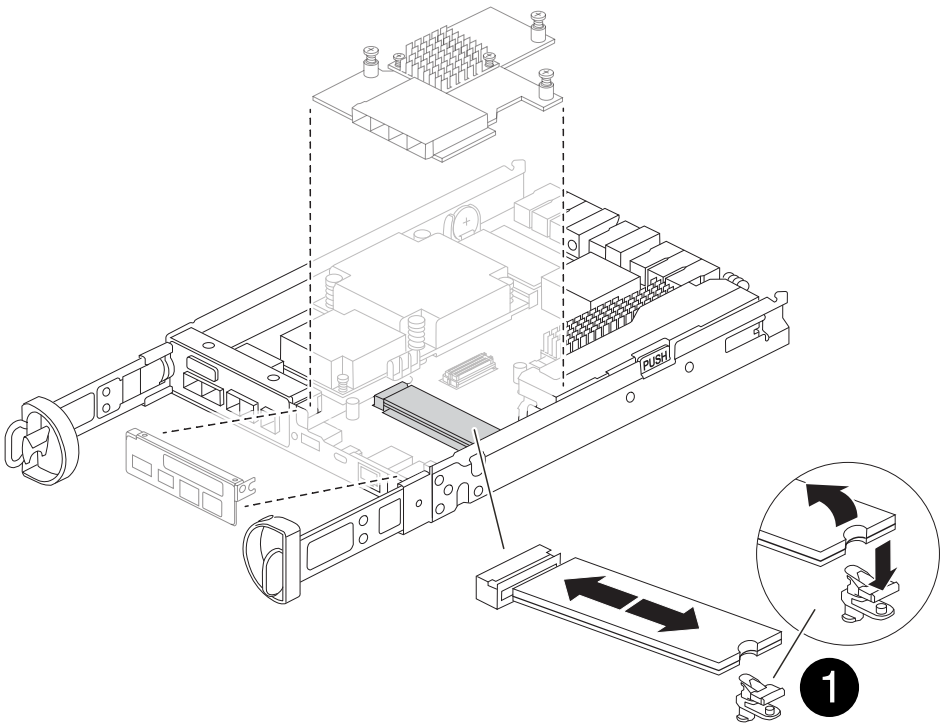
> You can loosen the thumbscrews with your fingers or a screwdriver.

3. Lift the mezzanine card straight up and set it aside on an anti-static surface.

Step 4: Move the boot media

Remove the boot media from the impaired controller module and install it in the replacement controller module.

1. After removing the mezzanine card, locate the boot media using the following illustration or the FRU map on the controller module:



1	Boot media release button

2. Remove the boot media:
   a. Press the blue button on the boot media housing to release the boot media from its housing.
   b. Rotate the boot media up, 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. Install the the boot media to the replacement controller module:
   a. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
   b. 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.

   c. Push the blue locking button on the boot media housing, rotate the boot media all the way down, and then release the locking button to lock the boot media in place.

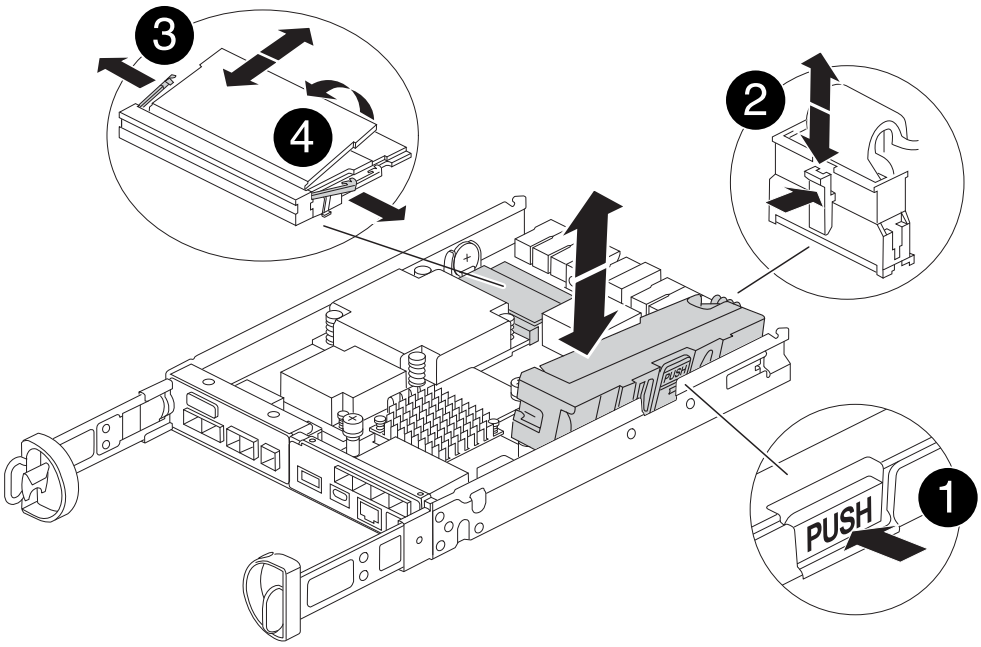
Step 5: Install the mezzanine card in the replacement controller

Install the mezzanine card in the replacement controller module.

1. Reinstall the mezzanine card:
   a. Align mezzanine card with the socket on the motherboard.
   b. Gently push down on the card to seat the card in the socket.
   c. Tighten the three thumbscrews on the mezzanine card.
2. Reinstall the IO Plate.

Step 6: Move the DIMMs

Remove the DIMMs from the impaired controller module and install them into the replacement controller module.



1	DIMM locking latches
2	DIMM

1. Locate the DIMMs on your controller module

   Note the location of the DIMM in the sockets so that you can insert the DIMM in the same location in the replacement controller module and in the proper orientation.

2. Remove the DIMMs from the impaired controller module:
   a. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM.

      The DIMM will rotate up a little.

   b. Rotate the DIMM as far as it will go, and then slide the DIMM out of the socket.

Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

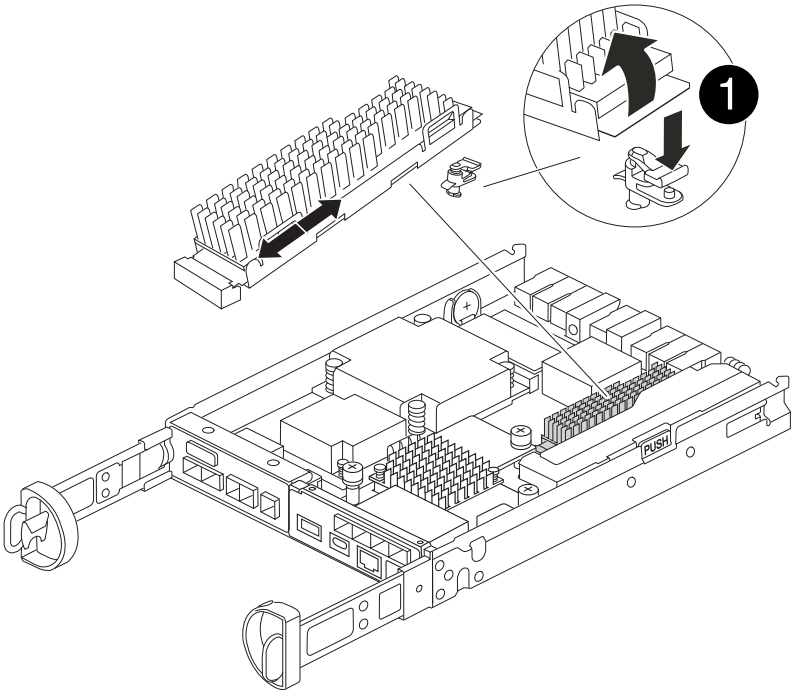
3. Verify that the NVMEM battery is not plugged into the replacement controller module.
4. Install the DIMMs in the replacement controller in the same place they were in the impaired controller:
   a. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.

      The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

      Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.
5. Repeat these steps for the other DIMM.

Step 7: Move a caching module

Remove the caching module from the impaired controller module install it into replacement controller module.



1	Caching module locking button

1. Locate the caching module near the rear of the controller module and remove it:
   a. Press the blue locking button and rotate the caching module upward.
   b. Gently pull the caching module straight out of the housing.
2. Install the caching module in the replacement controller module:
   a. Align the edges of the caching module with the socket in the housing, and then gently push it into the socket.

b. 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.

c. Push the blue locking button, rotate the caching module all the way down, and then release the locking button to lock the caching module in place.

3. Plug in the NVMEM battery.

   Make sure that the plug locks down into the battery power socket on the motherboard.

   If plugging in the battery is difficult, remove the battery from the controller module, plug it in, and then reinstall the battery into the controller module.

4. Reinstall the controller module cover.

Step 8: Install the NV battery

Install the NV battery into the replacement controller module.

1. Plug the battery plug back into the socket on the controller module.

   Make sure that the plug locks down into the battery socket on the motherboard.

2. Aligning the battery with the holding brackets on the sheet metal side wall.
3. Slide the battery pack down until the battery latch engages and clicks into the opening on the side wall.
4. Reinstall the controller module cover and lock it into place.

Step 9: Install the controller

Install the replacement controller module into the system chassis and boot ONTAP.

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. Turn the controller module.
4. 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.

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.

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.

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. If the firmware update is aborted, the boot process exits to the LOADER prompt. You must run the `update_flash` command, and then enter `bye -g` to reboot the system.

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. Respond `y` to this prompt.
- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the healthy controller remains down. Respond `y` to this prompt.

Restore and verify the system configuration - FAS2820

After completing the hardware replacement and booting the replacement controller, 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 controller does not match your system configuration, set the HA state for the replacement controller module: `ha-config modify controller HA-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mcc-2n`
   - `mccip`

     a. Confirm that the setting has changed: `ha-config show`

3. Reboot the controller module.

   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.

Recable the system and reassign disks - FAS2820

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

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

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. 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, resolve the veto issue. If the veto is not critical to resolve, you can override the veto.

      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.

6. 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
.
.
.
```

Complete system restoration - FAS2820

Restore your system to full operation by restoring the NetApp Storage Encryption or Volume Encryption configurations (if necessary), and installing licenses for the replacement controller, and returning 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: 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 - FAS2820

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.

Animation - Replace a DIMM

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

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=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, 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 Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
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`.

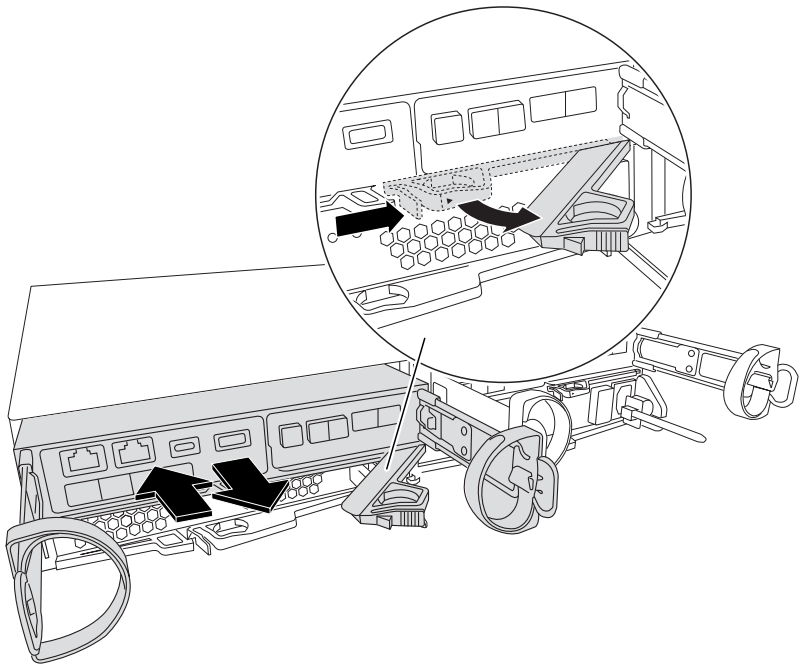
Step 2: Remove controller module

Remove the controller module from the system and then remove the controller module cover.

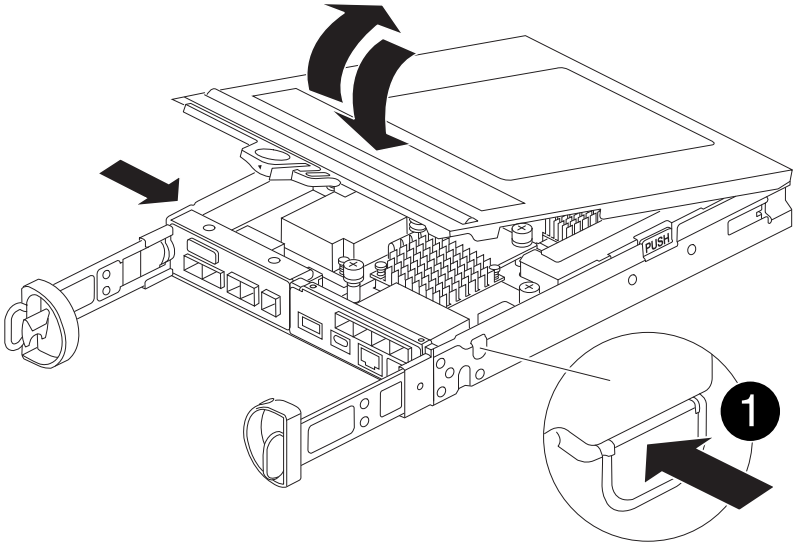
Steps

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 pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.



1	Controller module cover release button

Step 3: Replace the DIMMs

Locate the DIMM inside the controller, remove it, and replace it.

Before replacing a DIMM, you need to unplug the NVMEM battery from the controller module.

Steps

1. If you are not already grounded, properly ground yourself.

   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:

2. 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.
3. If the NVMEM LED is flashing, there is data in the NVMEM and you must disconnect the battery to clear the memory:
   a. Remove the battery from the controller module by pressing the blue button on the side of the controller module.
   b. Slide the battery up until it clears the holding brackets, and then lift the battery out of the controller module.
   c. Locate the battery cable, press the clip on the battery plug to release the lock clip from the plug socket, and then unplug the battery cable from the socket.

d. Confirm that the NVMEM LED is no longer lit.

e. Reconnect the battery connector and recheck the LED on the back of the controller.

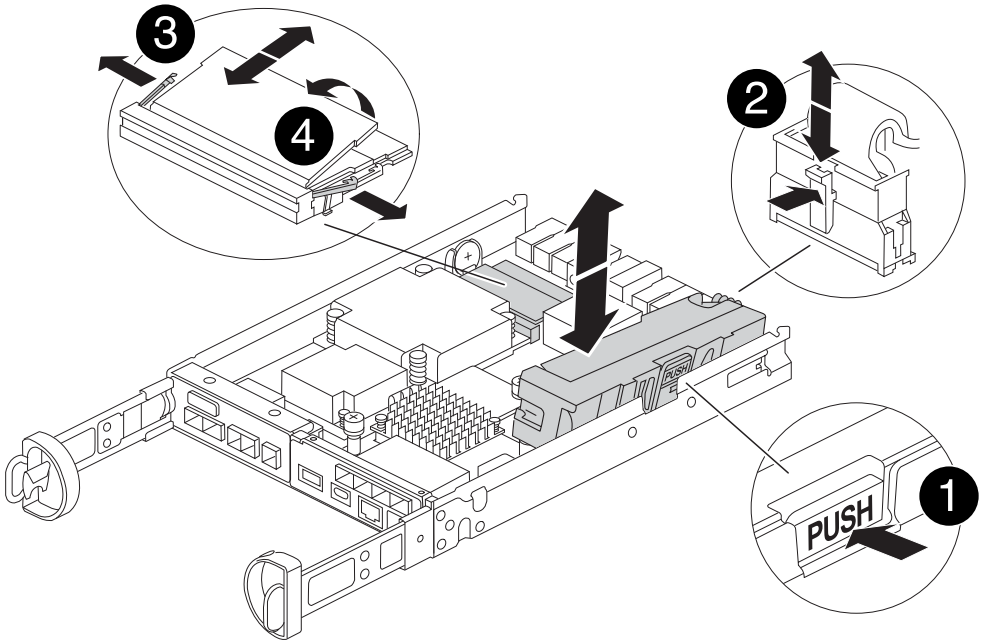
f. Unplug the battery cable.

4. Locate the DIMMs on your controller module.
5. Note the orientation and location of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
6. 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.

   The DIMM will rotate up a little.

7. Rotate the DIMM as far as it will go, and then slide the DIMM out of the socket.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.



1	NVRAM battery release button
2	NVRAM battery plug
3	DIMM ejector tabs
4	DIMMs

8. 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.

9. 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.

10. 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.
11. Reconnect the NVMRM battery:
    a. Plug in the NVRAM battery.

       Make sure that the plug locks down into the battery power socket on the motherboard.

    b. Align the battery with the holding brackets on the sheet metal side wall.
    c. Slide the battery pack down until the battery latch engages and clicks into the opening on the side wall.
12. Reinstall the controller module cover.

Step 4: Reinstall the controller module

Reinstall the controller module 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. Turn the controller module over and align the end with the opening in the chassis.
4. Gently push the controller module halfway into the system.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.

5. 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.

6. 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.

      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.
7. Reboot the controller module.

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.

Step 5: Restore automatic giveback and AutSupport

Restore automatic giveback and AutoSupport if they have been disabled.

1. Restore automatic giveback by using the `storage failover modify -node local -auto -giveback true` command.
2. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

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 - FAS2820

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 - FAS2820

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 NetApp Support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

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=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, 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 Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
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`.

Step 2: Remove and open the controller module

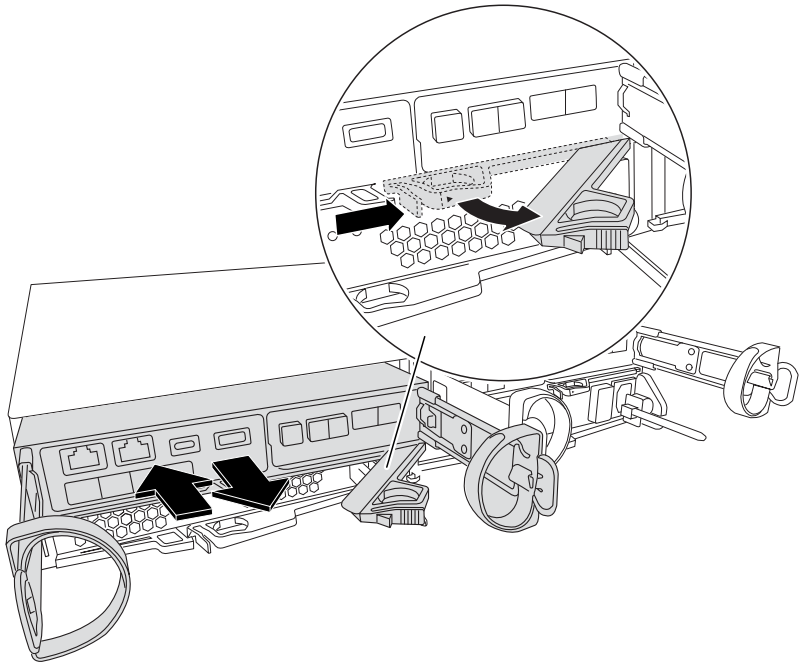
Remove and open the controller module.

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 half-way out of the chassis.



5. Check the NVMEM LED located on the back of the controller module. Look for the NV icon:





The green NV LED on the faceplate will start flashing when power is removed from the controller if the system was in the "waiting for giveback" state, or the system was not taken over or halted properly (uncommitted data). If the impaired controller module was not successfully taken over by the partner controller module, contact NetApp Support

   - If the green NV status LED begins flashing when the controller module is removed from the chassis:
     - Confirm that the controller had a clean takeover by the partner controller module or the impaired controller shows *waiting for giveback*, the flashing LED can be ignored and you can complete removing the impaired controller from the chassis.

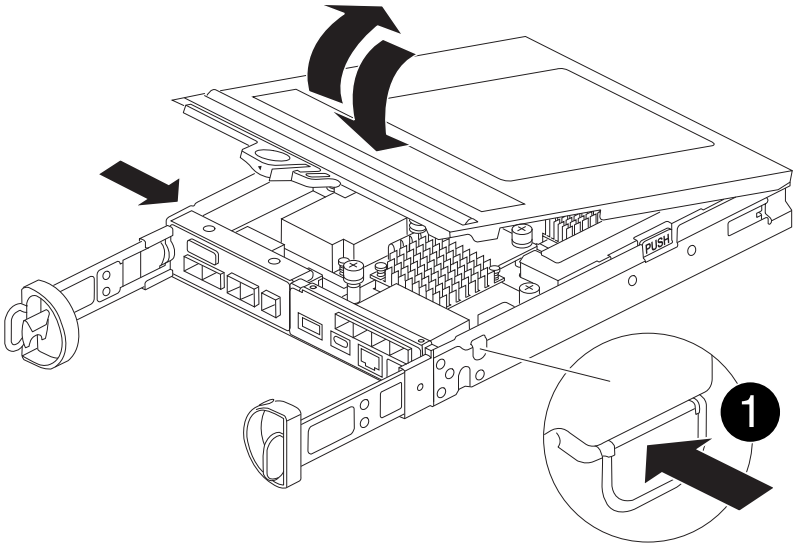
◦ If the green NV LED if off, you can complete removing the impaired controller from the chassis.

Step 3: Replace the NVMEM battery

Remove the failed NVMEM battery from the system and replace it with a new NVMEM battery.

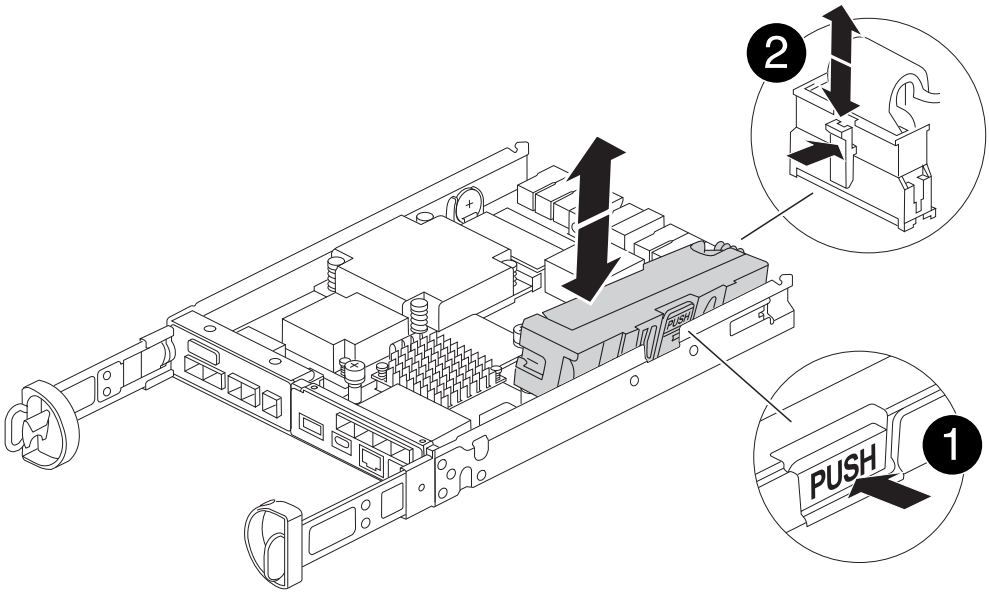
Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the controlle module from the chassis.
3. Turn the controller module over and place it on a flat, stable surface.
4. Open the cover by pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.



5. Locate the NVMEM battery in the controller module.

   Animation - Replace the NV battery



1	Battery release tab
2	Battery power connector

6. Remove the failed battery from the controller module:
   a. Press the blue button on the side of the controller module.
   b. Slide the battery up until it clears the holding brackets, and then lift the battery out of the controller module.
   c. Unplug the battery from the controller module
7. Remove the replacement battery from its package.
   Install the replacement battery:
   a. Plug the battery plug back into the socket on the controller module.

      Make sure that the plug locks down into the battery socket on the motherboard.

   b. Aligning the battery with the holding brackets on the sheet metal side wall.
   c. Slide the battery pack down until the battery latch engages and clicks into the opening on the side wall.
8. Reinstall the controller module cover and lock it into place.

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. Turn the controller module over and align the end with the opening in the chassis.
4. 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.

5. 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.

6. 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.

      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.

7. Reboot the controller module.

   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.

Step 5: Restore automatic giveback and AutSupport

Restore automatic giveback and AutoSupport if they have been disabled.

1. Restore automatic giveback by using the `storage failover modify -node local -auto -giveback true` command.
2. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

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 a mezzanine card - FAS2820

Replace the mezzanine card by disconnecting the cables and any SFP and QSFP modules from the card, replace the failed mezzanine card, and then recable the cards.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Animation - Replace the mezzanine card

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

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=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. If the impaired controller is part of an HA pair, 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 Remove controller module.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
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`.

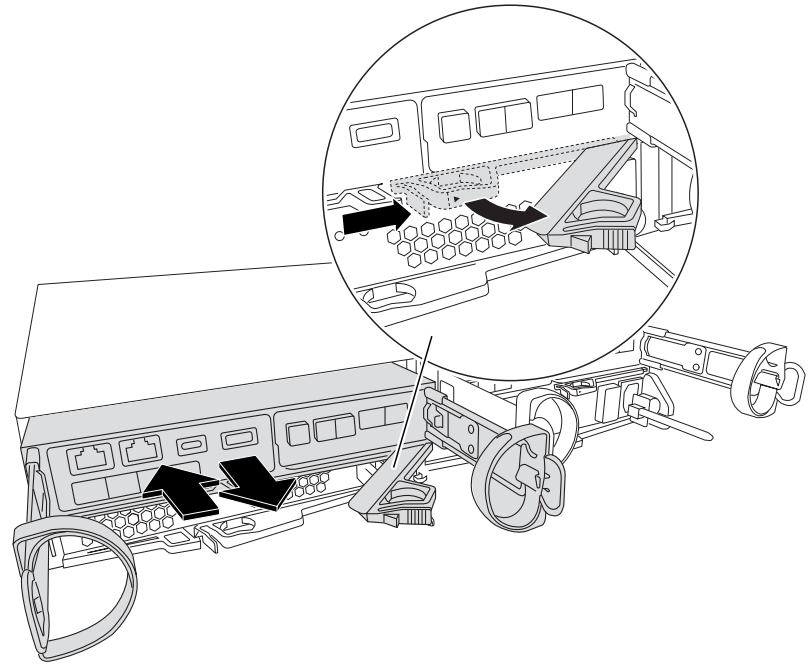
Step 2: Remove the controller module

Remove the controller module from the system and then remove the cover on the controller module.

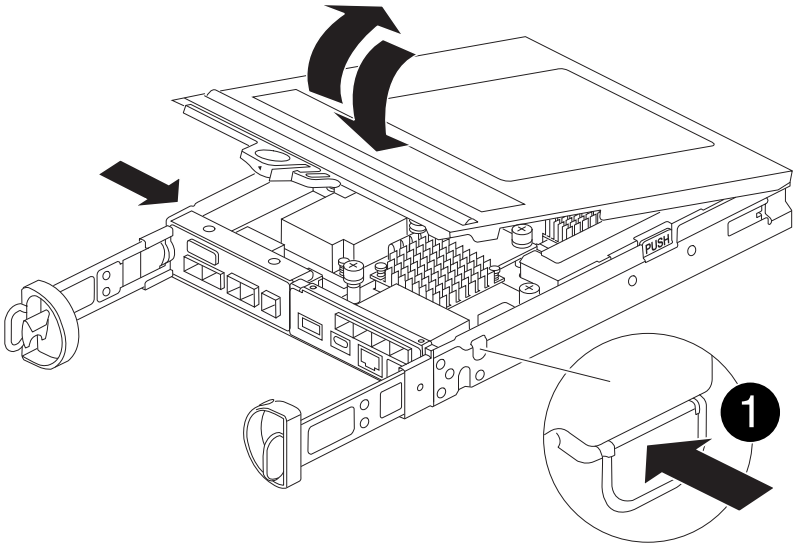
Steps

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 pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.

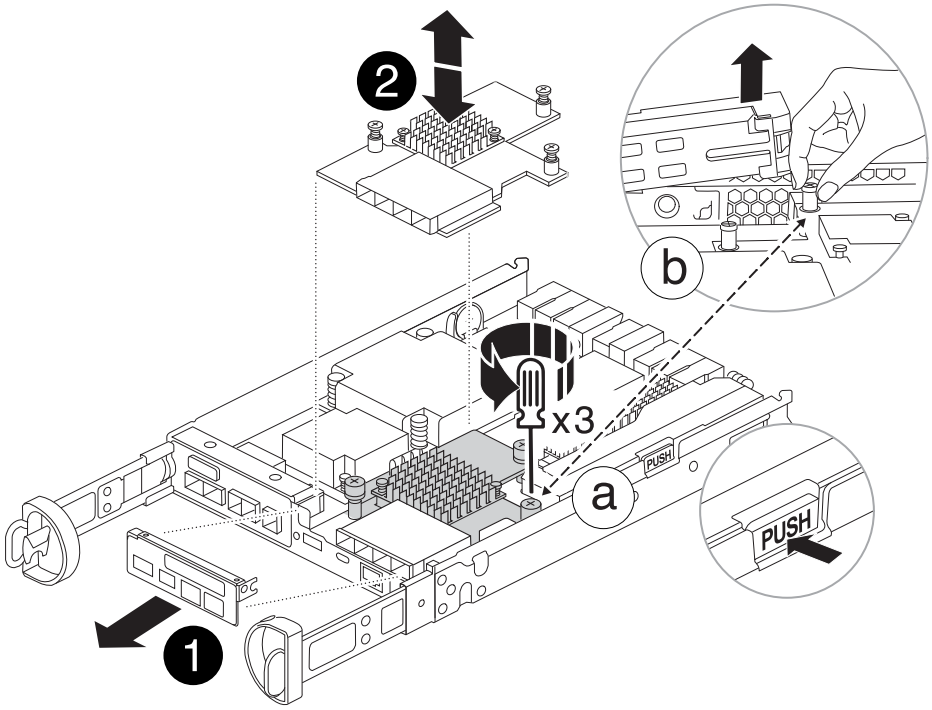


1	Controller module cover release button

Step 3: Replace the mezzanine card

Replace the mezzanine card.

1. If you are not already grounded, properly ground yourself.
2. Remove the mezzanine card using the following illustration or the FRU map on the controller module:



1	IO Plate
2	PCIe mezzanine card

a. Remove the IO Plate by sliding it straight out from thr controller module.

b. Loosen the thumbscrews on the mezzanine card and lift the mezzanine card straight up.

> You can loosen the thumbscrews with your fingers or a screwdriver. If you use your fingers, you might need to rotate the NV battery up for better finger purchase on the thumbscrew next to it.

3. Reinstall the mezzanine card:

   a. Align the socket on the replacement mezzanine card plug with the socket on the motherboard, and then gently seat the card squarely into the socket.

   b. Tighten the three thumbscrews on the mezzanine card.

   c. Reinstall the IO Plate.

4. Reinstall the controller module cover and lock it into place.

Step 4: Install the controller module

Reinstall the controller module.

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. Turn the controller module over and align the end with the opening in the chassis.

4. Gently push the controller module halfway into the system.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.

5. 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.

6. 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.

      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.
7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. Restore automatic giveback by using the `storage failover modify -node local -auto -giveback true` command.
9. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Swap out a power supply - FAS2820

Swapping out a power supply involves turning off, disconnecting, and removing the impaired 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.
- This procedure is written for replacing one power supply at a time.

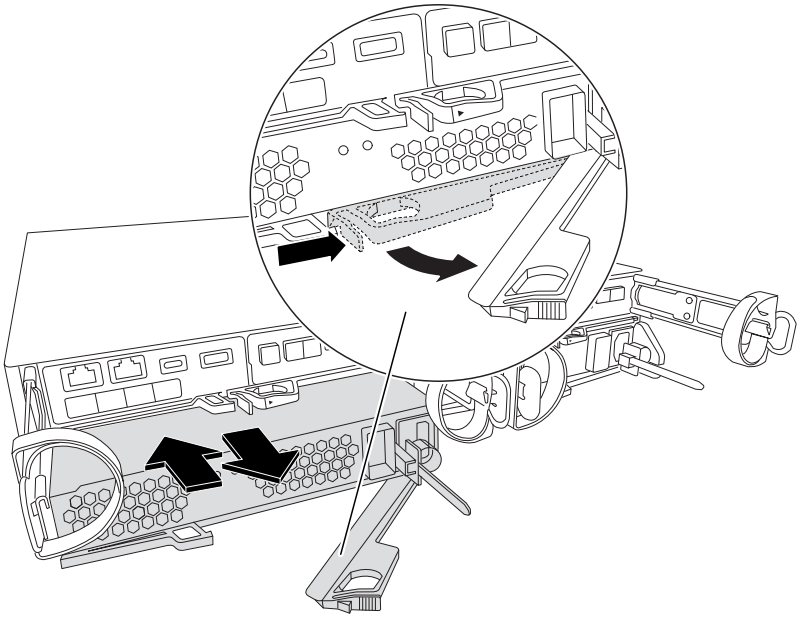
  It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

- Power supplies are auto-ranging.

Animation - Replace power supply

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 - FAS2820

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

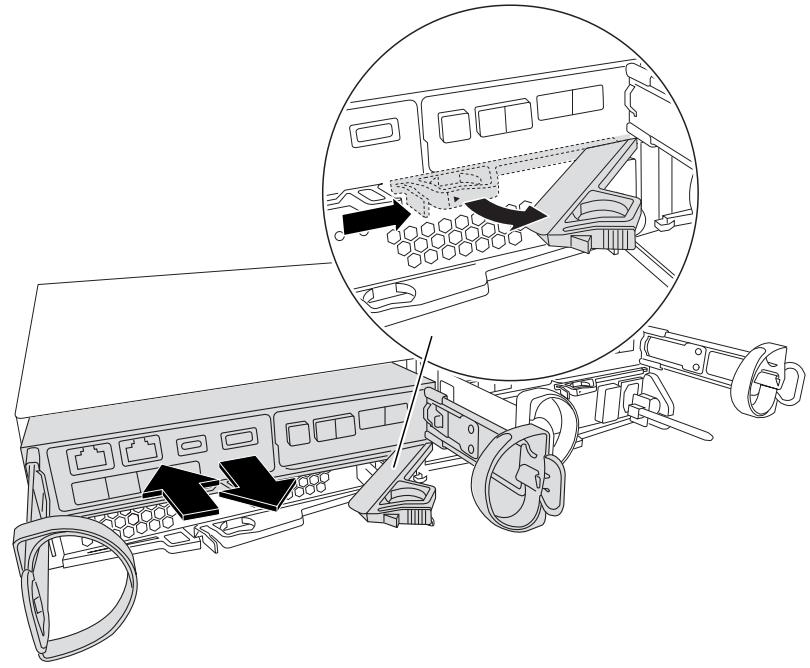
Step 2: Remove controller module

Remove the controller module from the system and then remove the cover on the controller module.

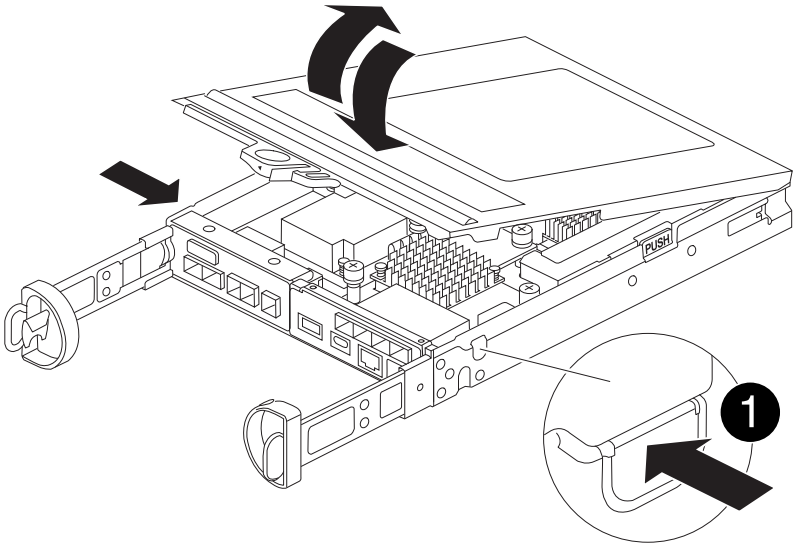
1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   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.



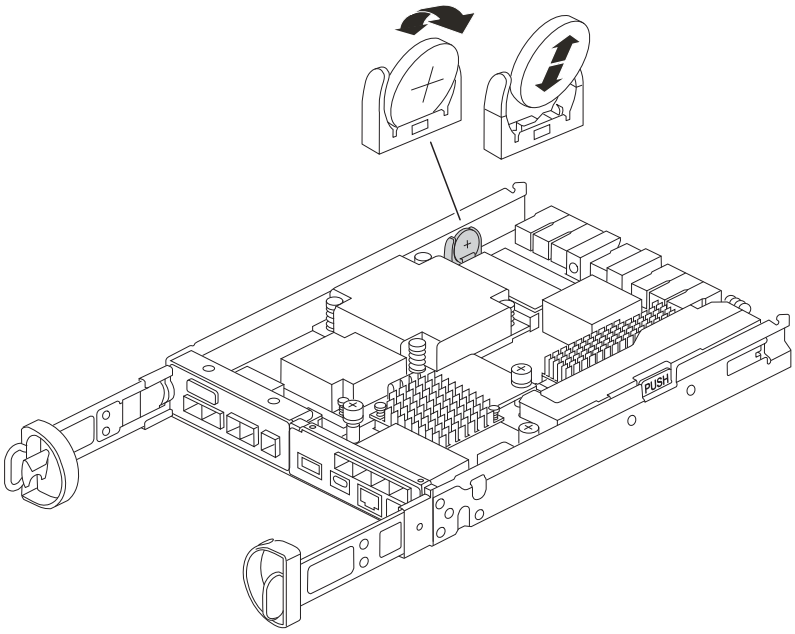
5. Turn the controller module over and place it on a flat, stable surface.
6. Open the cover by pressing the blue buttons on the sides of the controller module to release the cover, and then rotate the cover up and off of the controller module.



Step 3: Replace the RTC battery

Replace the RTC battery by locating it inside the controller and follow the specific sequence of steps.

Animation - Replace the RTC battery



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.
8. Reinstall the controller cover.

Step 4: Reinstall the controller module

Reinstall the controller module and boot it to the LOADER prompt..

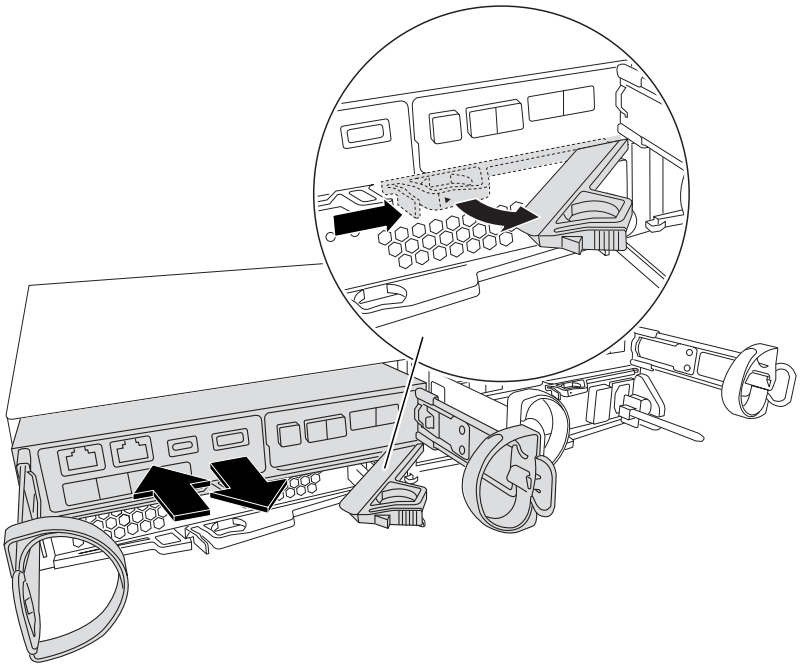
1. Turn the controller module over and align the end with the opening in the chassis.
2. Gently push the controller module halfway into the system.Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   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.

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. 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.

Step 5: Set time/date after RTC battery replacement

1. 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.
2. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
3. Return the controller to normal operation by giving back its storage: `storage failover giveback`

`-ofnode impaired_node_name`

4. Restore automatic giveback by using the `storage failover modify -node local -auto -giveback true` command.
5. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

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 FAS2820

The following are select specifications for the FAS2820 storage system in a single high availability pair. Visit NetApp Hardware Universe for the complete specifications for this storage system.

Key specifications for FAS2820

- Platform Configuration: FAS2820 Single Chassis HA Pair
- Max Raw Capacity: 3.1680 PB
- Memory: 128.0000 GB
- Form Factor: 2U chassis with 2 HA controllers and 12 drive slots
- ONTAP Version: ONTAP: 9.16.1P2
- PCIe Expansion Slots: 4
- Minimum ONTAP Version: ONTAP 9.13.1RC1

Scaleout maximums

- Type: NAS; HA Pairs: 12; Raw Capacity: 38.0 PB / 33.8 PiB; Max Memory: 1536 GB
- Type: SAN; HA Pairs: 6; Raw Capacity: 19.0 PB / 16.9 PiB; Max Memory: 768 GB
- Type: HA Pair; Raw Capacity: 3.2 PB / 2.8 PiB; Max Memory: 128.0000

I/O

Onboard I/O

- Protocol: Ethernet 25 Gbps; Ports: 4
- Protocol: SAS 12 Gbps; Ports: 4

Total I/O

- Protocol: Ethernet 25 Gbps; Ports: 12
- Protocol: Ethernet 10 Gbps; Ports: 8
- Protocol: FC 32 Gbps; Ports: 8
- Protocol: NVMe/FC 32 Gbps; Ports: 8

- Ports: 0
- Protocol: SAS 12 Gbps; Ports: 4

Management ports

- Protocol: Ethernet 1 Gbps; Ports: 2
- Protocol: RS-232 115 Kbps; Ports: 4
- Protocol: USB 600 Mbps; Ports: 2

Storage networking supported

- CIFS
- FC
- FCoE
- iSCSI
- NFS v3
- NFS v4.0
- NFS v4.1
- 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: 1815 BTU/hr
- Worst-case Power: 2339 BTU/hr
- Weight: 57.2 lb
  25.9 kg
- Height: 2U
- Width: 19” IEC rack-compliant (17.6” 44.7 cm)
- Depth: 20.0”
  (25.1” with cable management bracket)
- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°F) at up to 3048m (10000 ft) elevation;8% to 80% 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): 7.8
  Sound pressure (LpAm) (bystander positions): 68.4 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 for external shelves

FAS8300 and FAS8700 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.

For MetroCluster configurations, see either:

- Install MetroCluster IP configuration
- Install MetroCluster Fabric-Attached configuration

Quick guide - FAS8300 and FAS8700

This guide gives 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.

Access the *Installation and Setup Instructions* PDF poster:

FAS8300 and FAS8700 Installation and Setup Instructions

Video steps - FAS8300 and FAS8700

The following video shows how to install and cable your new system.

Animation - FAS8300 and FAS8700 Install and setup instructions

Detailed guide - FAS8300 and FAS8700

This guide gives detailed step-by-step instructions for installing a typical NetApp system. Use this guide if you want more detailed installation instructions.

Step 1: Prepare for installation

To install your system, you need to create an account, register the system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

You need to have access to the Hardware Universe for information about site requirements as well as additional information on your configured system. You might also want to have access to the Release Notes for your version of ONTAP for more information about this system.

NetApp Hardware Universe

Find the Release Notes for your version of ONTAP 9

You need to provide the following 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

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.



3. 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 the NetApp Hardware Universe to locate the cable and identify its use.

Type of cable…	Part number and length	Connector type	For…
100 GbE cable (QSF(28)	X66211A-05 (112-00595), 0.5m X66211A-1 (112-00573), 1m X66211A-2 (112-00574), 2m X66211A-5 (112-00574), 5m		Storage, cluster interconnect/HA, and Ethernet data (order-dependent)
25 GbE cable (SFP28s)	X66240-2 (112-00598), 2m X66240-5 (112-00639), 5m		GbE network connection (order-dependent)

Type of cable…	Part number and length	Connector type	For…
32 Gb FC (SFP+ Op)	X66250-2 (112-00342), 2m X66250-5 (112-00344), 5m X66250-15 (112-00346), 15m		FC network connection
Storage Cables	X66030A (112-00435), .5m X66031A (112-00436), 1m X66032A (112-00437), 2m X66033A (112-00438), 3m		mini-SAS HD to mini-SAS HD cables (order-dependent)
Optical cables	X66250-2-N-C (112-00342)		16 Gb FC or 25GbE cables for mezzanine cards (order-dependent)
RJ-45 (order dependent)	X6585-R6 (112-00291), 3m X6562-R6 (112-00196), 5m		Management network
Micro-USB console cable	Not applicable		Console connection used during software setup if laptop or console does not support network discovery.
Power cables	Not applicable		Powering up the system

4. Review the *NetApp ONTAP Confiuration Guide* and collect the required information listed in that guide.

   ONTAP Configuration Guide

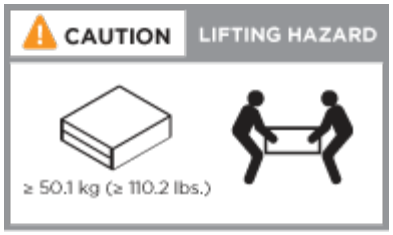
Step 2: Install the hardware

You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

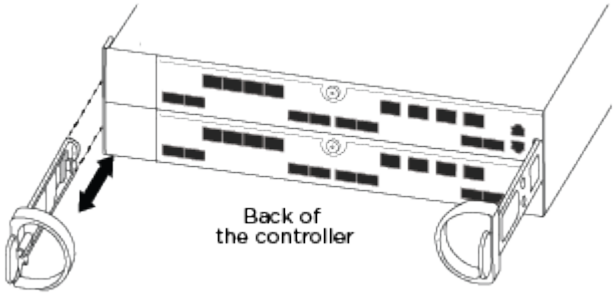
Steps

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 your network

You can cable the controllers to your network by using the two-node switchless cluster method or by using the cluster interconnect network.

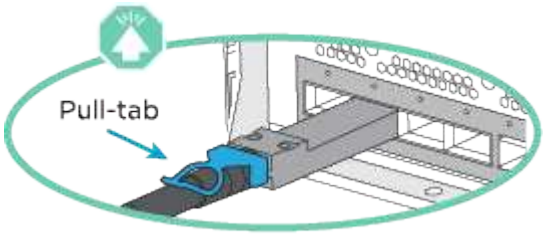
If the port labels on the card are not visible, check the card installation orientation (the PCIe connector socket is on the left side of the card slot in the A400 and FAS8300/8700), and then look for the card,then look for the card, by part number, in the NetApp Hardware Universe for a graphic of the bezel which will show the port labels. The card part number can be found using the `sysconfig -a` command or on the system packing list.

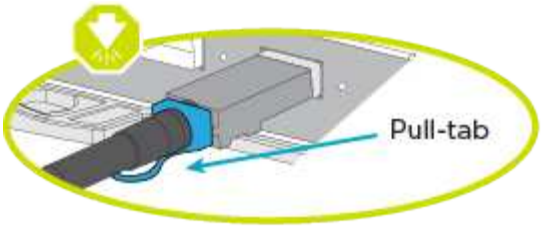
Option 1: Cable a two-node switchless cluster

The optional data ports, optional NIC cards, and management ports on the controller modules are connected to switches. The cluster interconnect and HA ports are cabled on both controller modules.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all onboard ports and down for expansion (NIC) cards.



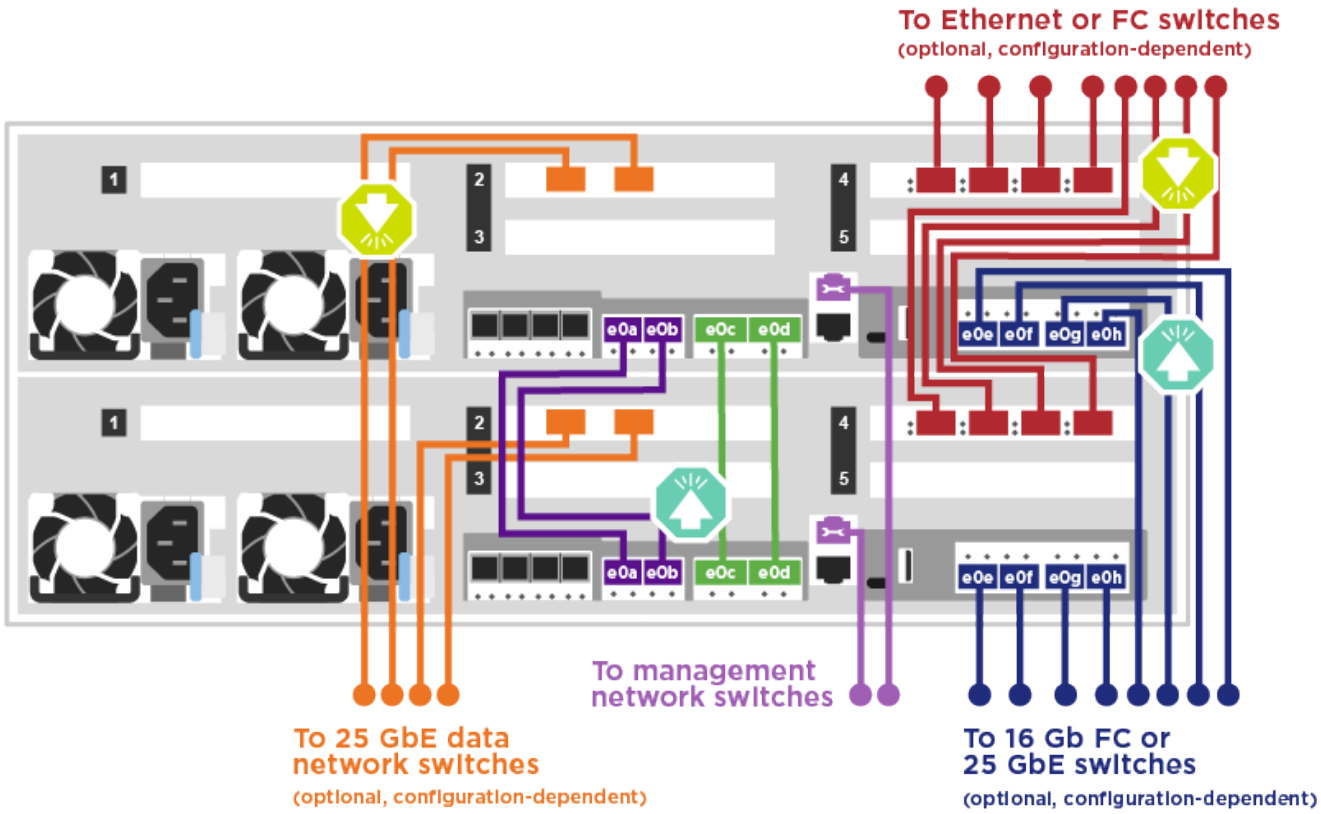


As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation - Two-node switchless cluster cabling



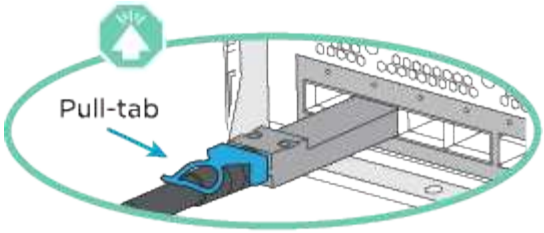
2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

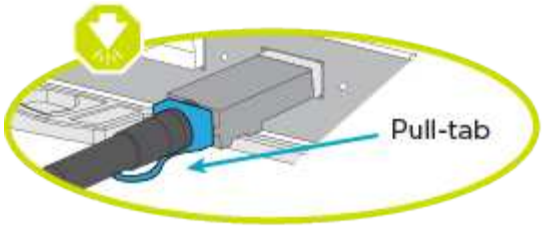
Option 2: Cable a switched cluster

The optional data ports, optional NIC cards, mezzanine cards, and management ports on the controller modules are connected to switches. The cluster interconnect and HA ports are cabled on to the cluster/HA switch.

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all onboard ports and down for expansion (NIC) cards.



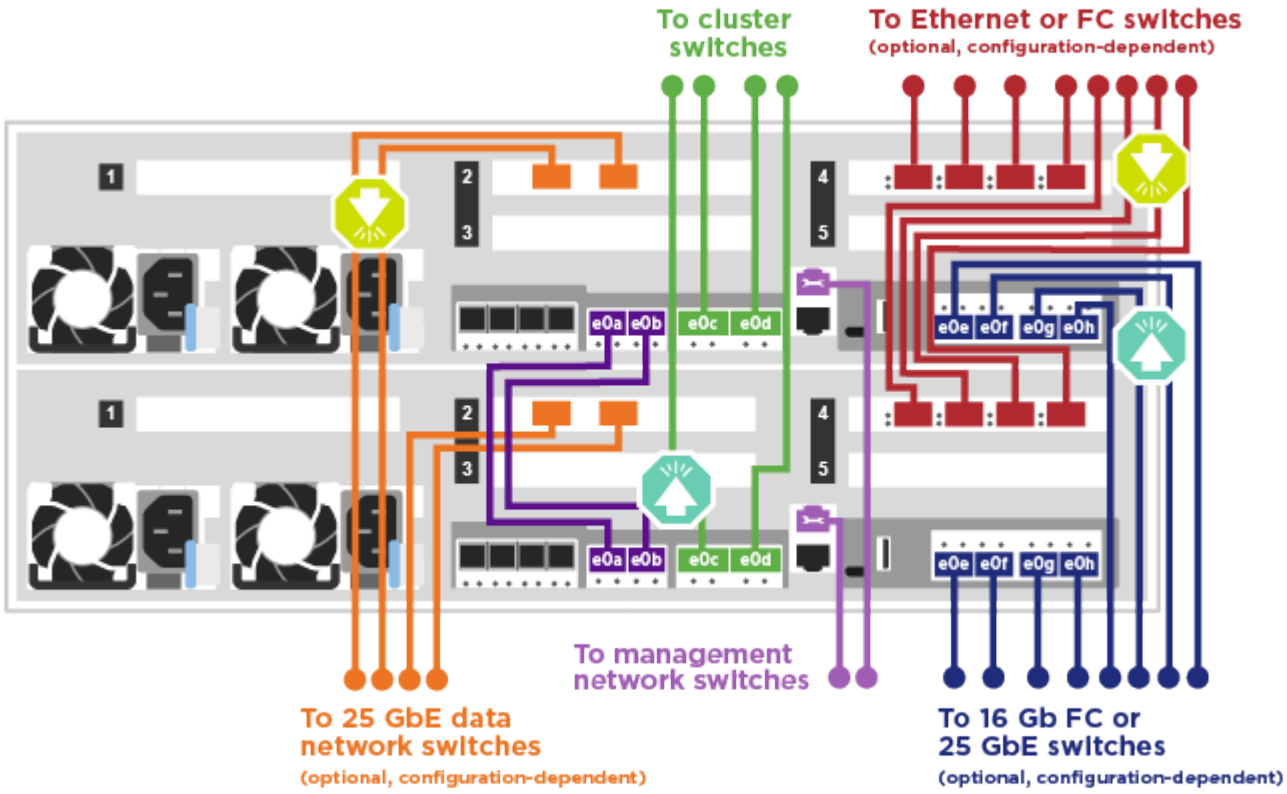


As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation - Switched cluster cabling



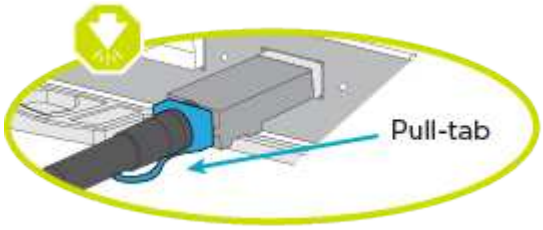
2. Go to Step 4: Cable controllers to drive shelves for drive shelf cabling instructions.

Step 4: Cable controllers to drive shelves

Option 1: Cable the controllers to SAS drive shelves

You must cable each controller to the IOM modules on both SAS drive shelves.

Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the DS224-C are down.

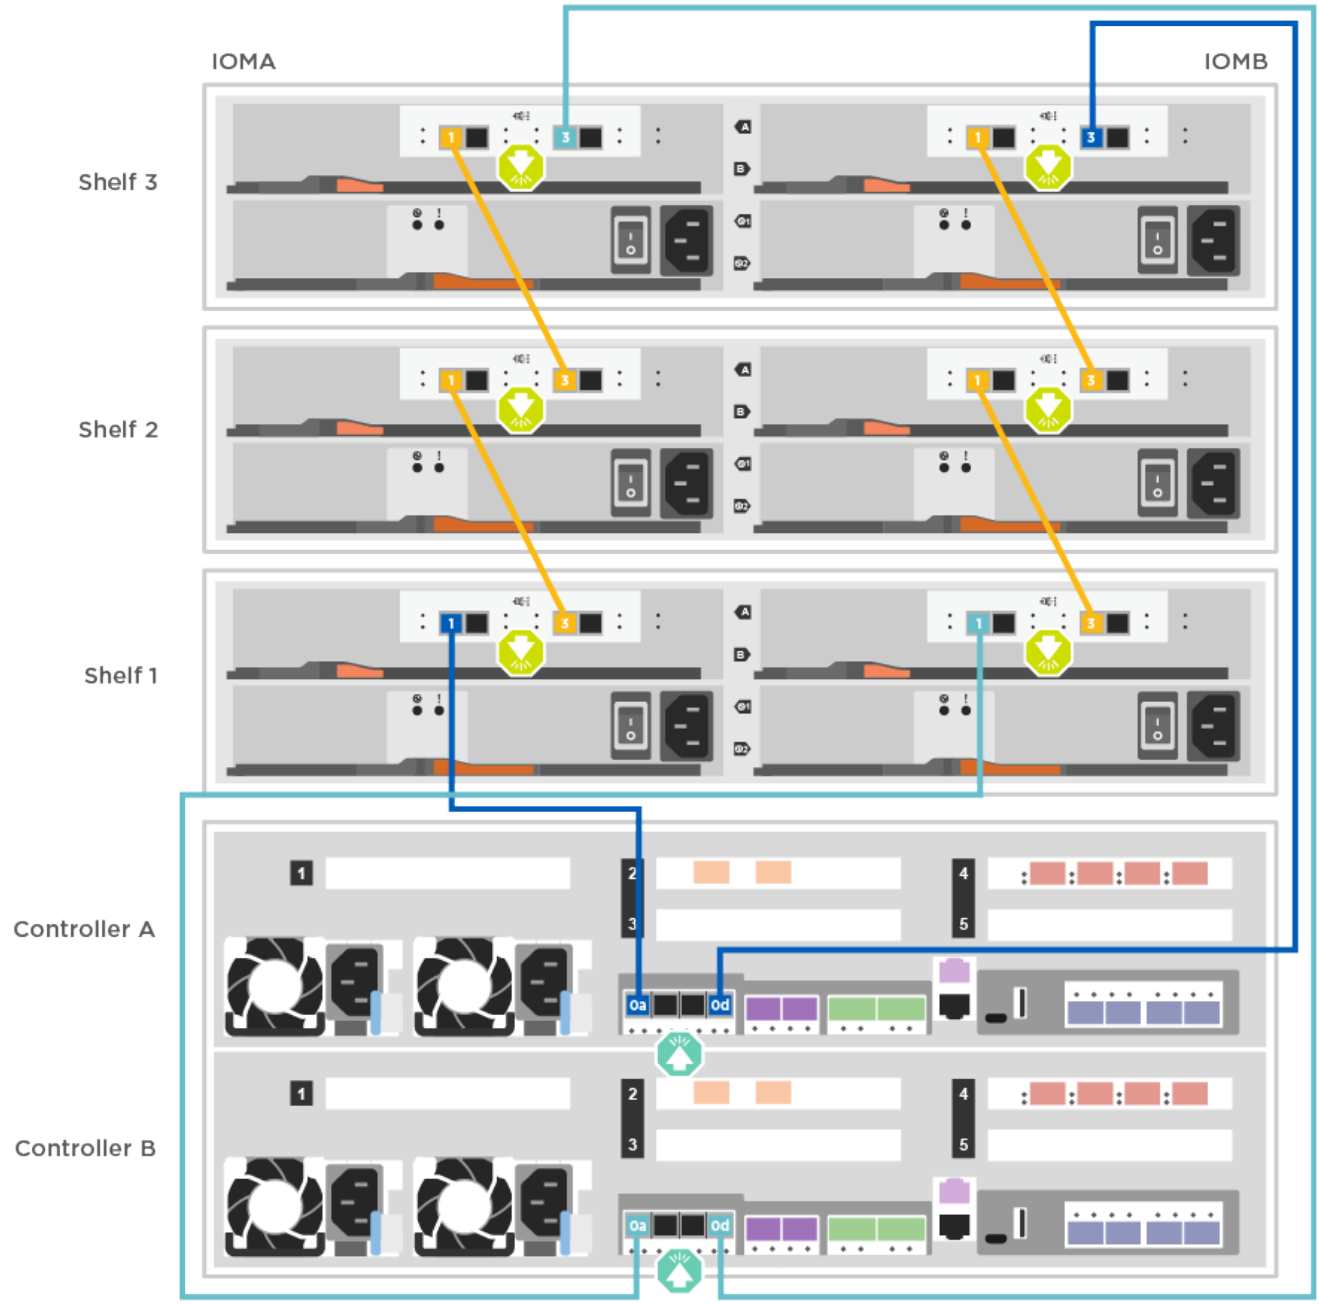


As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it around and try again.

Steps

1. Use the following animation or illustration to cable your controllers to two drive shelves.

   Animation - Cable the controllers to SAS drive shelves



2. Go to Step 5: Complete system setup and configuration to complete system setup and configuration.

Step 5: Complete system setup and configuration

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: Completing system setup and configuration 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

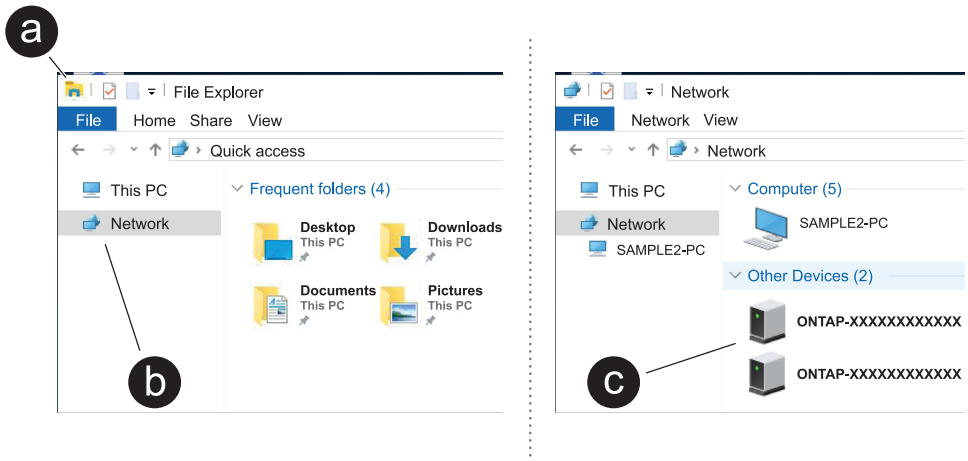
2. Plug the power cords into the controller power supplies, and then connect them to power sources on different circuits.
3. Make sure that your laptop has network discovery enabled.

   See your laptop's online help for more information.

4. Use the following animation to connect your laptop to the Management switch.

   Animation - Connect your laptop to the Management switch

5. 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.

6. Use System Manager guided setup to configure your system using the data you collected in the *NetApp ONTAP Configuration Guide*.

   ONTAP Configuration Guide

7. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: Completing system setup and configuration 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 how to configure the console port.

   b. Connect the console cable to the laptop or console using the console cable that came with your system, and then connect the laptop to the management switch on the management subnet .
   c. 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.

   FAS8300 and FAS8700 shown.

   Animation - Power on the controllers

   Initial booting may take up to eight minutes.

4. 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.

5. 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 *NetApp ONTAP Configuration guide*.

ONTAP Configuration Guide

6. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.

      NetApp Support Registration

   b. Register your system.

      NetApp Product Registration

   c. Download Active IQ Config Advisor.

      NetApp Downloads: Config Advisor

7. Verify the health of your system by running Config Advisor.
8. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain FAS8300 and FAS8700 hardware

Maintain the hardware of your FAS8300 and FAS8700 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the FAS8300 and FAS8700 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the FAS8300 and FAS8700 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the storage system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Caching module	You must replace the controller's caching module when your system registers a single AutoSupport (ASUP) message that the module has gone offline.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
Fan	The fan cools the controller.
NVDIMM	The NVDIMM (non-volatile dual in-line memory module) manages the data transfer from the volatile memory to the non-volatile storage, and maintains data integrity in the event of a power loss or system shutdown.
NVDIMM battery	A NVDIMM battery is responsible for maintaining power to the NVDIMM module.
PCIe card and risers	A PCIe (peripheral component interconnect express) card is an expansion card that plugs into the PCIe slot on the motherboard or into risers plugged into the motherboard.
Power supply	A power supply provides a redundant power source in a controller shelf.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - FAS8300 and FAS8700

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on the partner node to reinstall ONTAP on the replacement boot media in your FAS8300 or FAS8700 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1 Review the boot media requirements

Review the requirements for boot media replacement.

2 Shut down the controller

Shut down the controller in your storage system when when you need to replace the boot media.

3 Replace the boot media

Remove the failed boot media from the controller module and install the replacement boot media.

4 Restore the image on the boot media

Restore the ONTAP image from the partner controller.

5 Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - FAS8300 and FAS8700

Before replacing the boot media in your FAS8300 or FAS8700, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.

- /cfcard/kmip/certs/CA.pem file.

- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - FAS8300 and FAS8700

Shut down the impaired controller in your FAS8300 or FAS8700 storage system to prevent data loss and ensure system stability when replacing the boot media.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - FAS8300 and FAS8700

The boot media in your FAS8300 or FAS8700 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

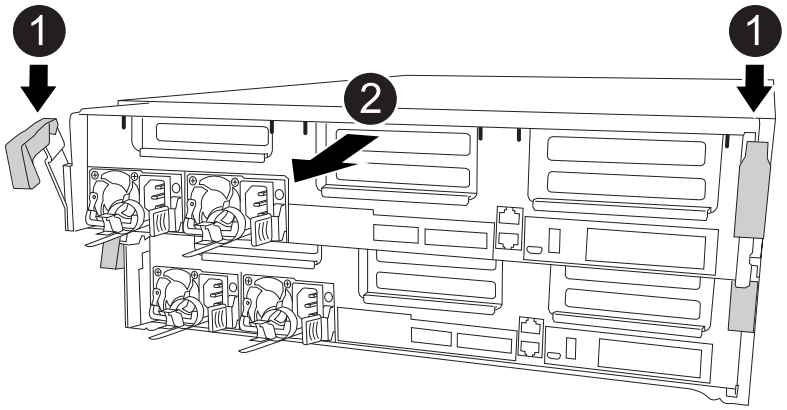
The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

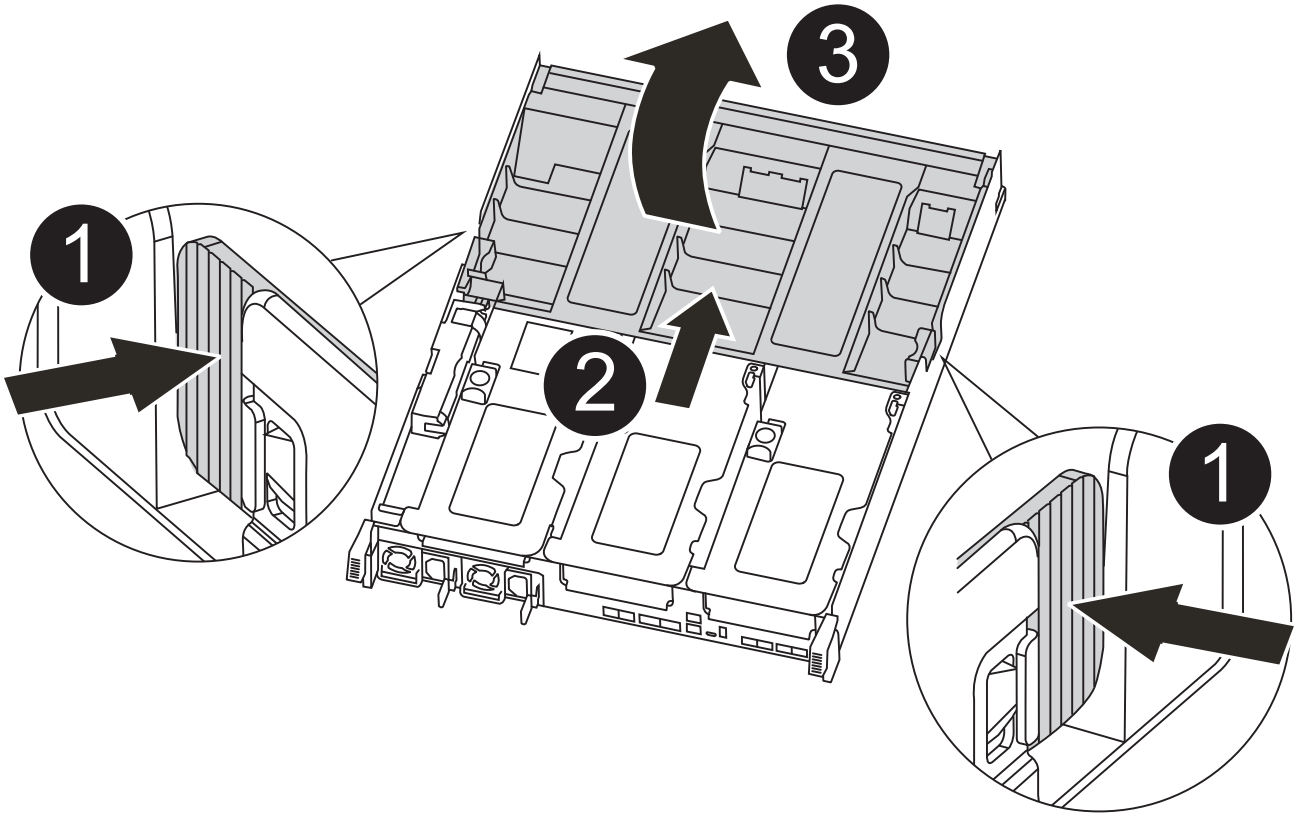


1	Locking latches
2	Controller moves slightly out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.
8. Open the air duct:

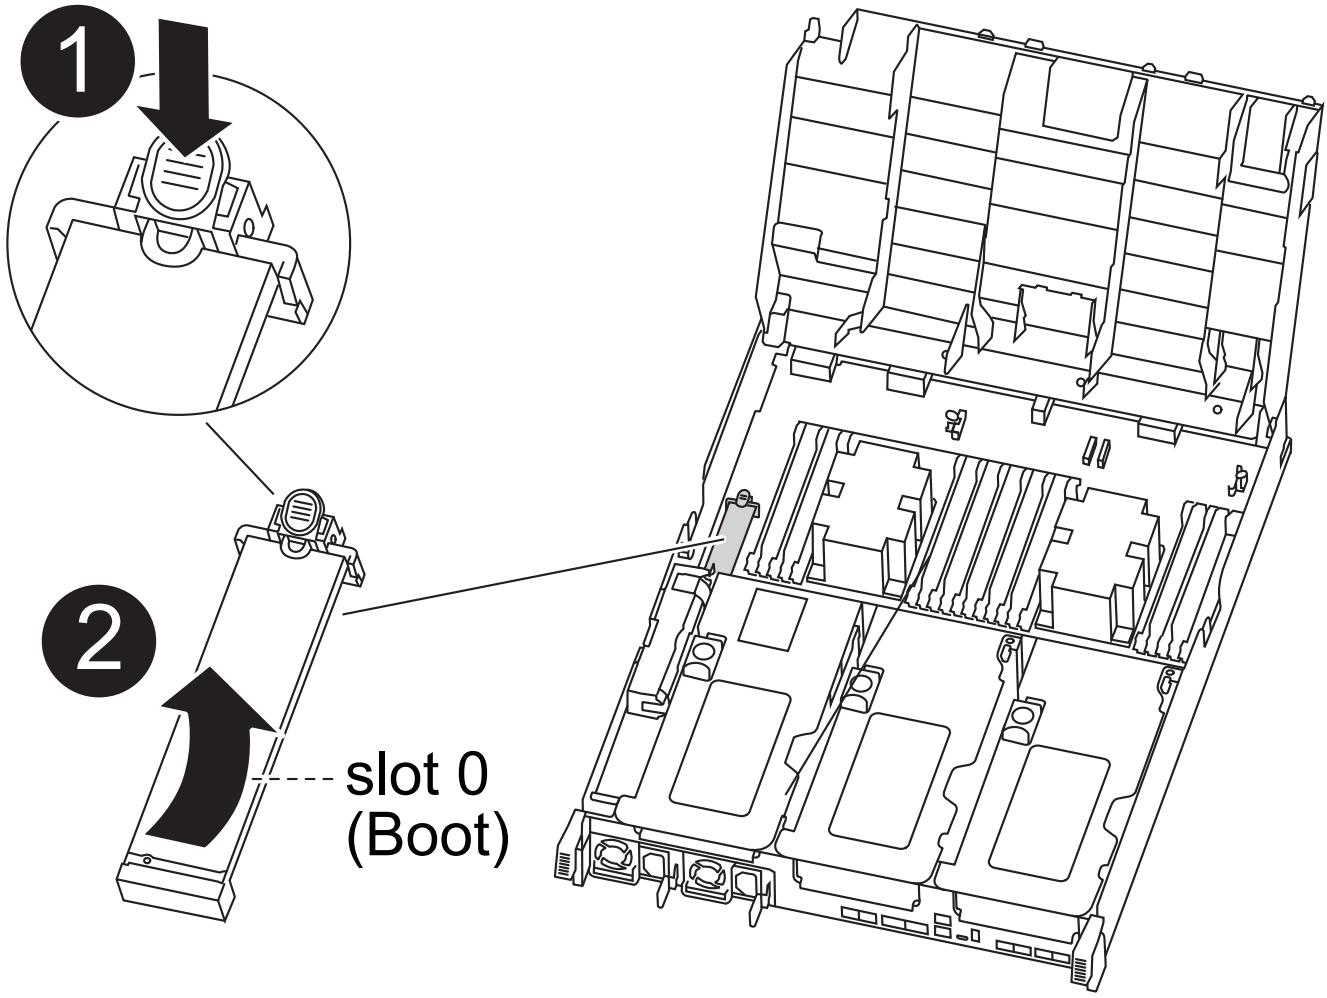


1	Locking tabs
2	Slide air duct toward back of controller
3	Rotate air duct up

a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.

b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

9. Locate and remove the boot media from the controller module:



1	Press blue button
2	Rotate boot media up and remove from socket

a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.

b. Rotate the boot media up and gently pull the boot media out of the socket.

10. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
11. Check the boot media to make sure that it is seated squarely and completely in the socket.

    If necessary, remove the boot media and reseat it into the socket.

12. Lock the boot media in place:
    a. Rotate the boot media down toward the motherboard.
    b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
    c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
13. Close the air duct.

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - FAS8300 and FAS8700

After installing the new boot media device in your FAS8300 or FAS8700 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- Determine your key manager type:
  - Onboard Key Manager (OKM): Requires cluster-wide passphrase and backup data
  - External Key Manager (EKM): Requires the following files from the partner node:
    - `/cfcard/kmip/servers.cfg`
    - `/cfcard/kmip/certs/client.crt`
    - `/cfcard/kmip/certs/client.key`
    - `/cfcard/kmip/certs/CA.pem`

Steps

1. From the LOADER prompt, start the boot media recovery process:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and displays one of the following messages:

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. a. Wait for the login prompt to display. b. Log into the node and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* c. Go to re-enabling automatic giveback if it was disabled.
`key manager is configured.`	Encryption is installed. Go to restoring the key manager.

> If the system cannot identify the key manager configuration, it displays an error message and prompts you to confirm whether key manager is configured and which type (onboard or external). Answer the prompts to proceed.

1. Restore the key manager using the appropriate procedure for your configuration:

Onboard Key Manager (OKM)

The system displays the following message and begins running BootMenu Option 10:

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the passphrase for onboard key management when prompted.

c. Enter the passphrase again when prompted to confirm.

d. Enter the backup data for onboard key manager when prompted.

Show example of passphrase and backup data prompts

```
Enter the passphrase for onboard key management:
-----BEGIN PASSPHRASE-----
<passphrase_value>
-----END PASSPHRASE-----
Enter the passphrase again to confirm:
-----BEGIN PASSPHRASE-----
<passphrase_value>
-----END PASSPHRASE-----
Enter the backup data:
-----BEGIN BACKUP-----
<passphrase_value>
-----END BACKUP-----
```

e. Monitor the recovery process as it restores the appropriate files from the partner node.

When the recovery process is complete, the node reboots. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

f. After the node reboots, verify that the system is back online and operational.

g. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

h. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster:

```
security key-manager onboard sync
```

Go to re-enabling automatic giveback if it was disabled.

External Key Manager (EKM)

The system displays the following message and begins running BootMenu Option 11:

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. Enter the EKM configuration settings when prompted:

i. Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file:

Show example of client certificate contents

```
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----
```

ii. Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file:

Show example of client key file contents

```
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----
```

iii. Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file:

Show example of KMIP server file contents

```
-----BEGIN CERTIFICATE-----
<KMIP_certificate_CA_value>
-----END CERTIFICATE-----
```

iv. Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file:

Show example of server configuration file contents

```
xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx
xxx.xxx.xxx.xxx:5696.port=5696
xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem
xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4
1xxx.xxx.xxx.xxx:5696.timeout=25
xxx.xxx.xxx.xxx:5696.nbio=1
xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.cr
t
xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key
xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:
!RC2:!RC4:!SEED:!eNULL:!aNULL"
xxx.xxx.xxx.xxx:5696.verify=true
xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>
```

v. If prompted, enter the ONTAP Cluster UUID from the partner node. You can check the cluster UUID from the partner node using the `cluster identify show` command.

Show example of ONTAP Cluster UUID prompt

```
Notice: bootarg.mgwd.cluster_uuid is not set or is empty.
Do you know the ONTAP Cluster UUID? {y/n} y
Enter the ONTAP Cluster UUID: <cluster_uuid_value>


System is ready to utilize external key manager(s).
```

vi. If prompted, enter the temporary network interface and settings for the node:

- The IP address for the port
- The netmask for the port
- The IP address of the default gateway

Show example of temporary network setting prompts

```
In order to recover key information, a temporary network
interface needs to be
configured.

Select the network port you want to use (for example,
'e0a')
e0M

Enter the IP address for port : xxx.xxx.xxx.xxx
Enter the netmask for port : xxx.xxx.xxx.xxx
Enter IP address of default gateway: xxx.xxx.xxx.xxx
Trying to recover keys from key servers....
[discover_versions]
[status=SUCCESS reason= message=]
```

b. Verify the key restoration status:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored. The process restores the appropriate files from the partner node and reboots the node. Proceed to the next step.
- If the key is not successfully restored, the system halts and displays error and warning messages. Rerun the recovery process from the LOADER prompt: `boot_recovery -partner`

Show example of key recovery error and warning messages

```
ERROR: kmip_init: halting this system with encrypted
mroot...
WARNING: kmip_init: authentication keys might not be
available.
*****************************************************
*                 A T T E N T I O N                 *
*                                                   *
*       System cannot connect to key managers.      *
*                                                   *
*****************************************************
ERROR: kmip_init: halting this system with encrypted
mroot...
.
Terminated

Uptime: 11m32s
System halting...

LOADER-B>
```

c. After the node reboots, verify that the system is back online and operational.

d. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

Go to re-enabling automatic giveback if it was disabled.

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media to NetApp - FAS8300 and FAS8700

If a component in your FAS8300 or FAS8700 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - FAS8300 and FAS8700

Get started with replacing the boot media in your FAS8300 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.



Review the boot media requirements

Review the requirements for replacing the boot media.



Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.



Shut down the controller

Shut down the controller when when you need to replace the boot media.



Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.



Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.



Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONATP boot menu.



Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - FAS8300 and FAS8700

Before replacing the boot media in your FAS8300 or FAS8700 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - FAS8300 and FAS8700

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Check NVE support and download the correct ONTAP image

Determine whether your ONTAP version supports NetApp Volume Encryption (NVE) so you can download the correct ONTAP image for the boot media replacement.

Steps

1. Check if your ONTAP version supports encryption:

   ```
   version -v
   ```

   If the output includes `10no-DARE`, NVE is not supported on your cluster version.

2. Download the appropriate ONTAP image based on NVE support:
   - If NVE is supported: Download the ONTAP image with NetApp Volume Encryption
   - If NVE is not supported: Download the ONTAP image without NetApp Volume Encryption

Download the ONTAP image from the NetApp Support Site to your HTTP or FTP server or a local folder. You will need this image file during the boot media replacement procedure.

Step 2: Verify key manager status and back up configuration

Before shutting down the impaired controller, verify the key manager configuration and back up the necessary information.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manager is configured on your system, do one of the following:

   If no key manager is configured:

   You can safely shut down the impaired controller and proceed to the shutdown procedure.

   If a key manager is configured (EKM or OKM):

   a. Enter the following query command to display the status of the authentication keys in your key manager:

      ```
      security key-manager key query
      ```

   b. Review the output and check the value in the `Restored` column. This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Complete the appropriate procedure based on your key manager type:

External Key Manager (EKM)

Complete these steps based on the value in the `Restored` column.

If all keys show `true` in the Restored column:

You can safely shut down the impaired controller and proceed to the shutdown procedure.

If any keys show a value other than `true` in the Restored column:

a. Restore the external key management authentication keys to all nodes in the cluster:

   ```
   security key-manager external restore
   ```

   If the command fails, contact NetApp Support.

b. Verify that all authentication keys are restored:

   ```
   security key-manager key query
   ```

   Confirm that the `Restored` column displays `true` for all authentication keys.

c. If all keys are restored, you can safely shut down the impaired controller and proceed to the shutdown procedure.

Onboard Key Manager (OKM)

Complete these steps based on the value in the `Restored` column.

If all keys show `true` in the Restored column:

a. Back up the OKM information:

   i. Switch to advanced privilege mode:

      ```
      set -priv advanced
      ```

      Enter `y` when prompted to continue.

   ii. Display the key management backup information:

      ```
      security key-manager onboard show-backup
      ```

   iii. Copy the backup information to a separate file or your log file.

      You will need this backup information if you need to manually recover OKM during the replacement procedure.

   iv. Return to admin mode:

      ```
      set -priv admin
      ```

b. You can safely shut down the impaired controller and proceed to the shutdown procedure.

If any keys show a value other than `true` in the Restored column:

a. Synchronize the onboard key manager:

```
security key-manager onboard sync
```

Enter the 32-character alphanumeric onboard key management passphrase when prompted.

> This is the cluster-wide passphrase you created when you initially configured the Onboard Key Manager. If you do not have this passphrase, contact NetApp Support.

b. Verify all authentication keys are restored:

```
security key-manager key query
```

Confirm that the `Restored` column displays `true` for all authentication keys and the `Key Manager` type shows `onboard`.

c. Back up the OKM information:

i. Switch to advanced privilege mode:

```
set -priv advanced
```

Enter `y` when prompted to continue.

ii. Display the key management backup information:

```
security key-manager onboard show-backup
```

iii. Copy the backup information to a separate file or your log file.

You will need this backup information if you need to manually recover OKM during the replacement procedure.

iv. Return to admin mode:

```
set -priv admin
```

d. You can safely shut down the impaired controller and proceed to the shutdown procedure.

Shut down the controller for manual boot media recovery - FAS8300 and FAS8700

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller. Shut down or take over the impaired controller using the appropriate procedure for your configuration.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Option 1: Most systems

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

Steps

a. Take the impaired controller to the LOADER prompt:

If the impaired controller displays…	Then…
The LOADER prompt	Go to Remove controller module.
`Waiting for giveback…`	Press Ctrl-C, and then respond `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.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 3: Controller is in a two-node Metrocluster

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired controller.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Replace the boot media and prepare for manual boot recovery - FAS8300 and FAS8700

To replace the boot media, you must remove the impaired controller module, install the replacement boot media, and transfer the boot image to a USB flash drive.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module

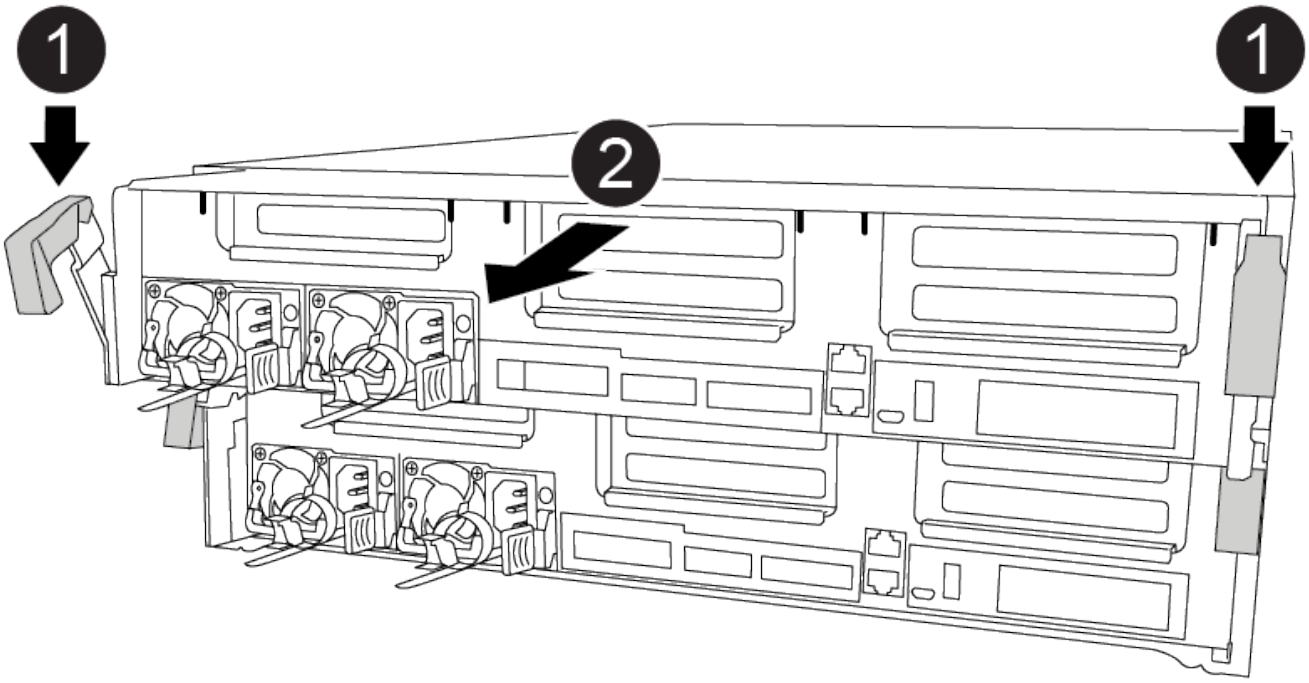
Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.



1	Locking latches
2	Slide controller out of chassis

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 2: Replace the boot media

You must locate the boot media in the controller module (see the FRU map on the controller module), and then follow the directions to replace it.

Before you begin

Although the contents of the boot media is encrypted, it is a best practice to erase the contents of the boot media before replacing it. For more information, see the Statement of Volatility for your system on the NetApp Support Site.

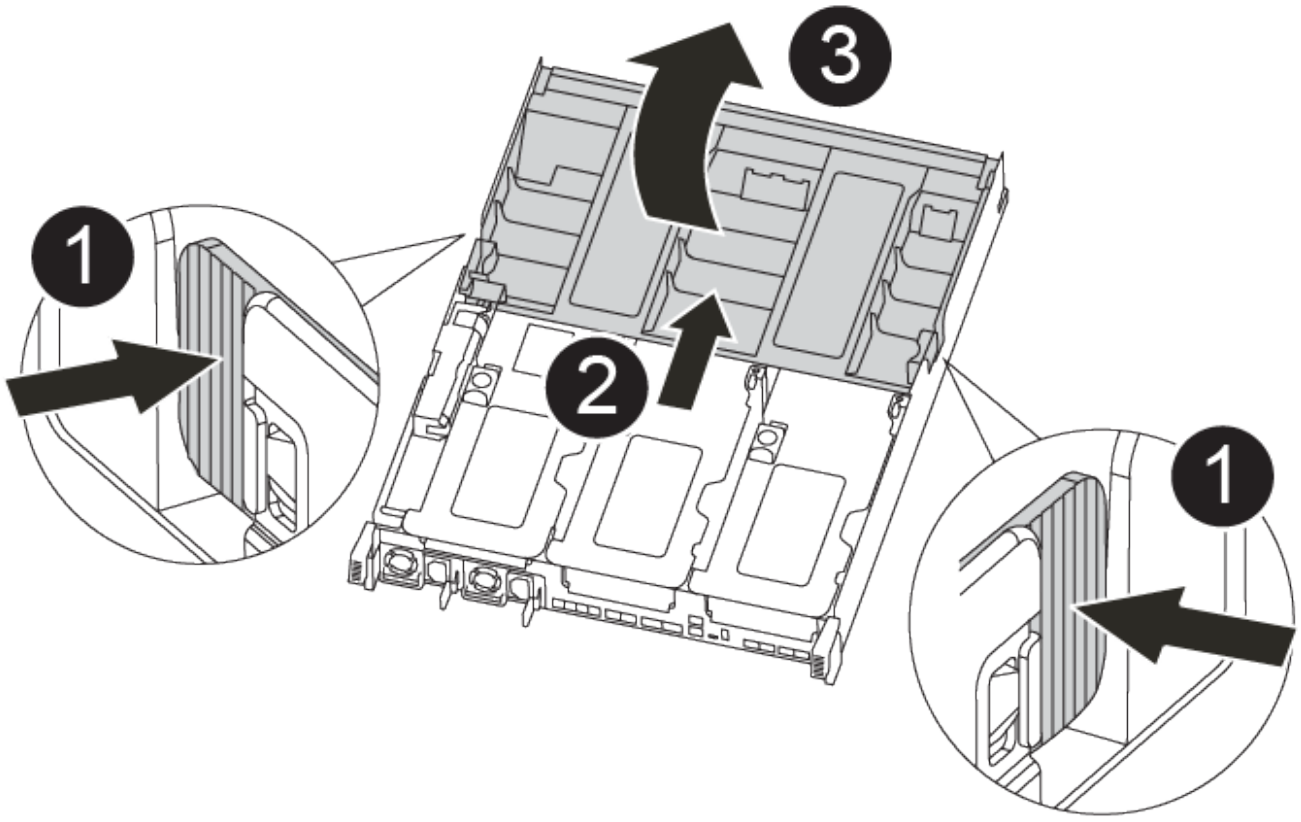
> You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

You can use the following animation, illustrations, or the written steps to replace the boot media.

Animation - Replace the boot media

Steps

1. Open the air duct:

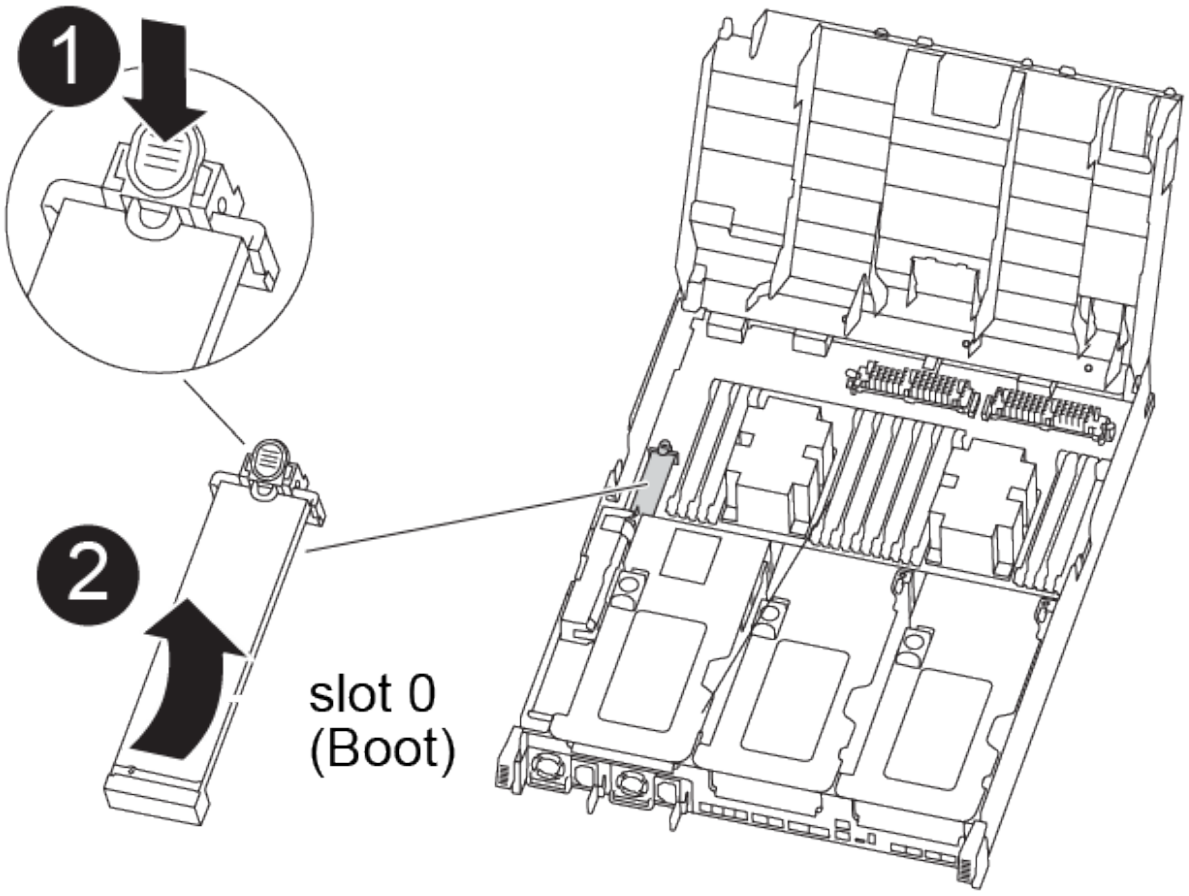


1	Locking tabs
2	Slide air duct toward back of controller
3	Rotate air duct up

a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.

b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.

2. Locate and remove the boot media from the controller module:



1	Press blue button
2	Rotate boot media up and remove from socket

a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
b. Rotate the boot media up and gently pull the boot media out of the socket.

3. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Lock the boot media in place:
   a. Rotate the boot media down toward the motherboard.
   b. Placing a finger at the end of the boot media by the blue button, push down on the boot media end to engage the blue locking button.
   c. While pushing down on the boot media, lift the blue locking button to lock the boot media in place.
6. Close the air duct.

Step 3: Transfer the boot image to the boot media

The replacement boot media that you installed does not have a boot image, so you need to transfer a boot image using a USB flash drive.

Before you begin

- You must have a USB flash drive, formatted to MBR/FAT32, with at least 4GB capacity
- A copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site
  - If NVE is enabled, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not enabled, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is an HA pair, you must have a network connection.
- If your system is a stand-alone system you do not need a network connection, but you must perform an additional reboot when restoring the `var` file system.

1. Download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image to your work space on your laptop.
   b. Unzip the service image.

      If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

      There are two folders in the unzipped service image file:

      - `boot`
      - `efi`

   c. Copy the `efi` folder to the top directory on the USB flash drive.

      If the service image has no efi folder, see EFI folder missing from Service Image download file used for boot device recovery for FAS and AFF models^ .

      The USB flash drive should have the efi folder and the same Service Image (BIOS) version of what the impaired controller is running.

   d. Remove the USB flash drive from your laptop.
2. If you have not already done so, close the air duct.
3. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
4. Reinstall the cable management device and recable the system, as needed.

   When recabling, remember to reinstall the media converters (SFPs or QSFPs) if they were removed.

5. Insert the USB flash drive into the USB slot on the controller module.

   Make sure that you install the USB flash drive in the slot labeled for USB devices, and not in the USB

console port.

6. Complete the installation of the controller module:
   a. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
7. Interrupt the boot process by pressing Ctrl-C to stop at the LOADER prompt.

   If you miss this message, press Ctrl-C, select the option to boot to Maintenance mode, and then `halt` the controller to boot to LOADER.

8. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:
   a. Boot to Maintenance mode: `boot_ontap maint`
   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t initiator adapter_name`
   c. Halt to return to Maintenance mode: `halt`

The changes will be implemented when the system is booted.

Manual boot media recovery from a USB drive - FAS8300 and FAS8700

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Before you begin

- Ensure your console is connected to the impaired controller.
- Verify you have a USB flash drive with the recovery image.
- Determine if your system uses encryption. You will need to select the appropriate option in step 3 based on whether encryption is enabled.

Steps

1. From the LOADER prompt on the impaired controller, boot the recovery image from the USB flash drive:

```
boot_recovery
```

The recovery image is downloaded from the USB flash drive.

2. When prompted, enter the name of the image or press **Enter** to accept the default image displayed in brackets.
3. Restore the var file system using the procedure for your ONTAP version:

ONTAP 9.16.0 or earlier

Complete the following steps on the impaired controller and partner controller:

a. **On the impaired controller:** Press `Y` when you see `Do you want to restore the backup configuration now?`

b. **On the impaired controller:** If prompted, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. **On the partner controller:** Set the impaired controller to advanced privilege level:

```
set -privilege advanced
```

d. **On the partner controller:** Run the restore backup command:

```
system node restore-backup -node local -target-address
impaired_node_IP_address
```

> If you see any message other than a successful restore, contact NetApp Support.

e. **On the partner controller:** Return to admin level:

```
set -privilege admin
```

f. **On the impaired controller:** Press `Y` when you see `Was the restore backup procedure successful?`

g. **On the impaired controller:** Press `Y` when you see `...would you like to use this restored copy now?`

h. **On the impaired controller:** Press `Y` when prompted to reboot, then press `Ctrl-C` when you see the Boot Menu.

i. **On the impaired controller:** Do one of the following:

- If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.
- If the system uses encryption, go to Restore encryption.

ONTAP 9.16.1 or later

Complete the following steps on the impaired controller:

a. Press `Y` when prompted to restore the backup configuration.

   After the restore procedure is successful, this message displays: `syncflash_partner: Restore from partner complete`

b. Press `Y` when prompted to confirm that the restore backup was successful.

c. Press `Y` when prompted to use the restored configuration.

d. Press `Y` when prompted to reboot the node.

e. Press `Y` when prompted to reboot again, then press `Ctrl-C` when you see the Boot Menu.

f. Do one of the following:

- If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.

- If the system uses encryption, go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -fromnode local
```

6. If you disabled automatic giveback, reenable it:

```
storage failover modify -node local -auto-giveback true
```

7. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Restore encryption - FAS8300 and FAS8700

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Complete the appropriate steps to restore encryption on your system based on your key manager type. If you are unsure which key manager your system uses, check the settings you captured at the beginning of the boot media replacement procedure.

Onboard Key Manager (OKM)

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

Ensure you have the following information available:

- Cluster-wide passphrase entered while enabling onboard key management
- Backup information for the Onboard Key Manager
- Verification that you have the correct passphrase and backup data using the How to verify onboard key management backup and cluster-wide passphrase procedure

Steps

On the impaired controller:

1. Connect the console cable to the impaired controller.
2. From the ONTAP boot menu, select the appropriate option:

ONTAP version	Select this option
ONTAP 9.8 or later	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
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
```

```
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 boot media to NetApp - FAS8300 and FAS8700

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the caching module - FAS8300 and FAS8700

You must replace the caching module in the controller module when your system registers a single AutoSupport (ASUP) message that the module has gone offline; failure to do so results in performance degradation.

> The Ver2 controller module has only one caching module socket in the FAS8300. FAS8700 does not have a VER2 controller module. The caching module functionality is not impacted by the socket removal.

- You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this tasks

If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller.

Synchronize a node with the cluster

You might want to erase the contents of your caching module before replacing it.

Steps

1. Although data on the caching module is encrypted, you might want to erase any data from the impaired caching module and verify that the caching module has no data:
   a. Erase the data on the caching module: `system controller flash-cache secure-erase run -node node name localhost -device-id device_number`

      Run the `system controller flash-cache show` command if you don't know the Flash Cache device ID.

   b. Verify that the data has been erased from the caching module: `system controller flash-cache secure-erase show`
2. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=_number_of_hours_down_h`

   The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`
3. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
4. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y`.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

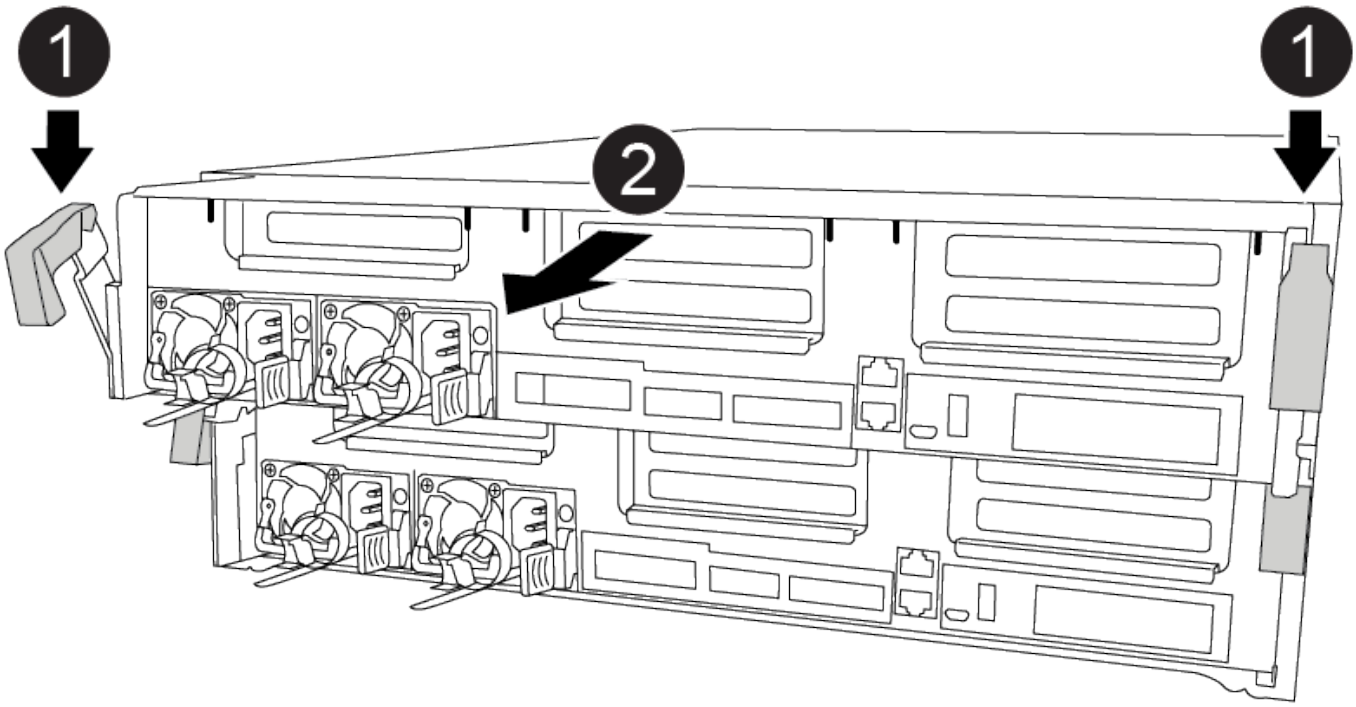
8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module



Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace a caching module

To replace a caching module, referred to as the Flash Cache on the label on your controller, locate the slot inside the controller and follow the specific sequence of steps. See the FRU map on the controller module for the location of the Flash Cache.

Slot 6 is only available in FAS8300 VER2 Controller.

Your storage system must meet certain criteria depending on your situation:

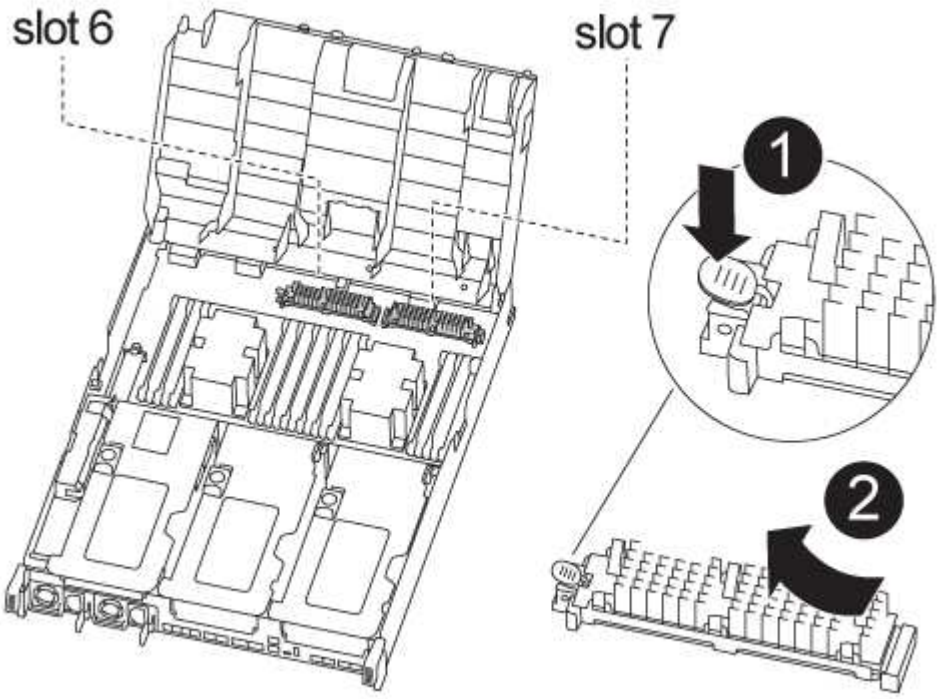
- It must have the appropriate operating system for the caching module you are installing.
- It must support the caching capacity.
- Although the contents of the caching module is encrypted, it is a best practice to erase the contents of the module before replacing it. For more information, see the Statement of Volatility for your system on the NetApp Support Site.

> You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

- All other components in the storage system must be functioning properly; if not, you must contact technical support.

You can use the following animation, illustration, or the written steps to replace a caching module.

Animation - Replace the caching module



Steps

1. If you are not already grounded, properly ground yourself.
2. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
3. Using the FRU map on the controller module, locate the failed caching module and remove it:

   Depending on your configuration, there may be zero, one, or two caching modules in the controller module. Use the FRU map inside the controller module to help locate the caching module.

   a. Press the blue release tab.

The caching module end rises clear of the release tab.

b. Rotate the caching module up and slide it out of the socket.

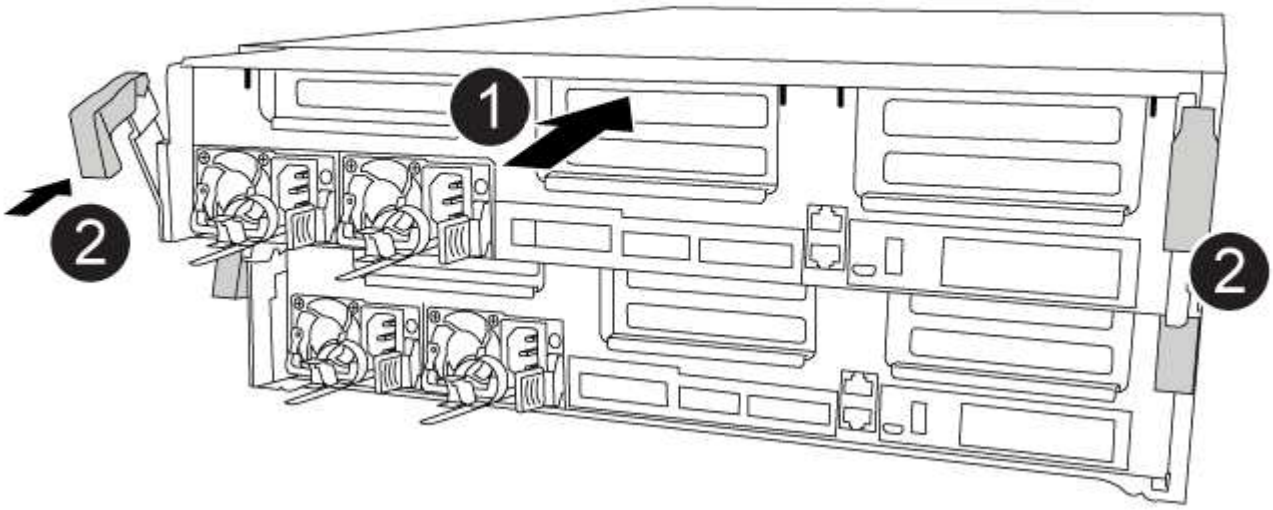
4. Install the replacement caching module:
   a. Align the edges of the replacement caching module with the socket and gently insert it into the socket.
   b. Rotate the caching module downward toward the motherboard.
   c. Placing your finger at the end of the caching module by the blue button, firmly push down on the caching module end, and then lift the locking button to lock the caching module in place.
5. Close the air duct:
   a. Rotate the air duct down to the controller module.
   b. Slide the air duct toward the risers to lock it in place.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module



Steps

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

   f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 7: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 8: Complete the replacement process

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Chassis

Overview of chassis replacement - FAS8300 and FAS8700

All other components in the system must be functioning properly; if not, you must contact technical support.

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is disruptive. For a two-controller cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - FAS8300 and FAS8700

Option 1: Most configurations

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

   ```
   system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
   true -ignore-quorum-warnings true -inhibit-takeover true
   ```

   For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

   ```
   Warning: Are you sure you want to halt node <node_name>? {y|n}:
   ```

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster configuration

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes            RAID
Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Move and replace hardware - FAS8300 and FAS8700

Move the fans, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Remove the controller modules

To replace the chassis, you must remove the controller modules from the old chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move the fans

To move the fan modules to the replacement chassis when replacing the chassis, you must perform a specific sequence of tasks.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and

then pulling it toward you until the bezel releases from the ball studs on the chassis frame.

3. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

   The fan module moves a little bit away from the chassis.

4. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

5. Set the fan module aside.
6. Repeat the preceding steps for any remaining fan modules.
7. Insert the fan module into the replacement chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

   The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.
10. Repeat these steps for the remaining fan modules.

Step 3: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

Steps

1. Remove the screws from the chassis mount points.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two people, install the replacement chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or equipment rack.
5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot it.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller

module halfway into the system.

Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

   f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.
   g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - FAS8300 and FAS8700

You must verify the HA state of the chassis and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set the HA state of the chassis

You must verify the HA state of the chassis, and, if necessary, update the state to match your system configuration.

Steps

1. In Maintenance mode, from either controller module, display the HA state of the local controller module and chassis: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state for the chassis does not match your system configuration:

a. Set the HA state for the chassis: `ha-config modify chassis HA-state`

   The value for HA-state can be one of the following:

   - `ha`
   - `mcc`
   - `mcc-2n`
   - `mccip`
   - `non-ha`

b. Confirm that the setting has changed: `ha-config show`

3. If you have not already done so, recable the rest of your system.

Step 2: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 3: Complete the replacement process

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Controller

Overview of controller module replacement - FAS8300 and FAS8700

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- If your system is in a MetroCluster configuration, you must review the section Choosing the correct recovery procedure to determine whether you should use this procedure.

  If this is the procedure you should use, note that the controller replacement procedure for a controller in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired controller to the *replacement* controller so that the *replacement* controller will boot up in the same version of ONTAP as the old controller module.

- It is important that you apply the commands in these steps on the correct systems:
  - The *impaired* controller is the controller that is being replaced.
  - The *replacement node* is the new controller that is replacing the impaired controller.
  - The *healthy* controller is the surviving controller.
- You must always capture the controller's console output to a text file.

  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - FAS8300 and FAS8700

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

To replace the controller module hardware, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

The Ver2 controller module has only one caching module socket (Slot 6) in the FAS8300. FAS8700 does not have a VER2 controller module. The caching module functionality is not impacted by the socket removal.

Step 1: Remove the controller module

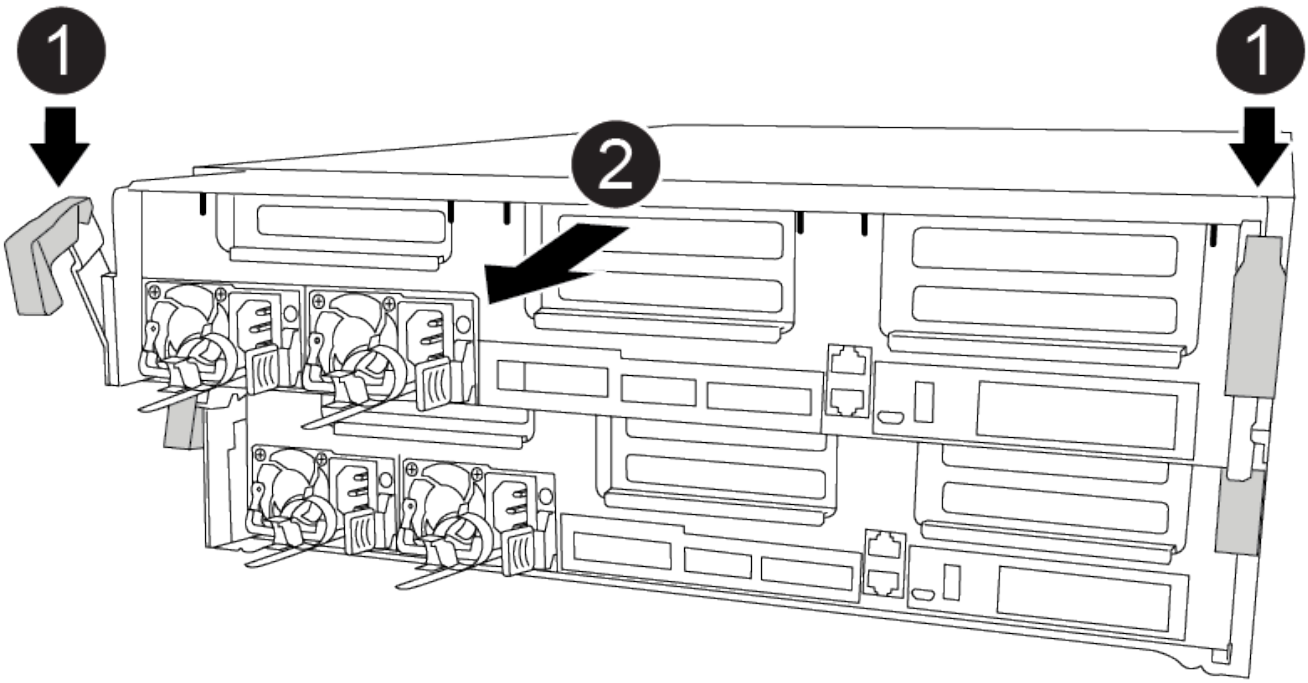
To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.



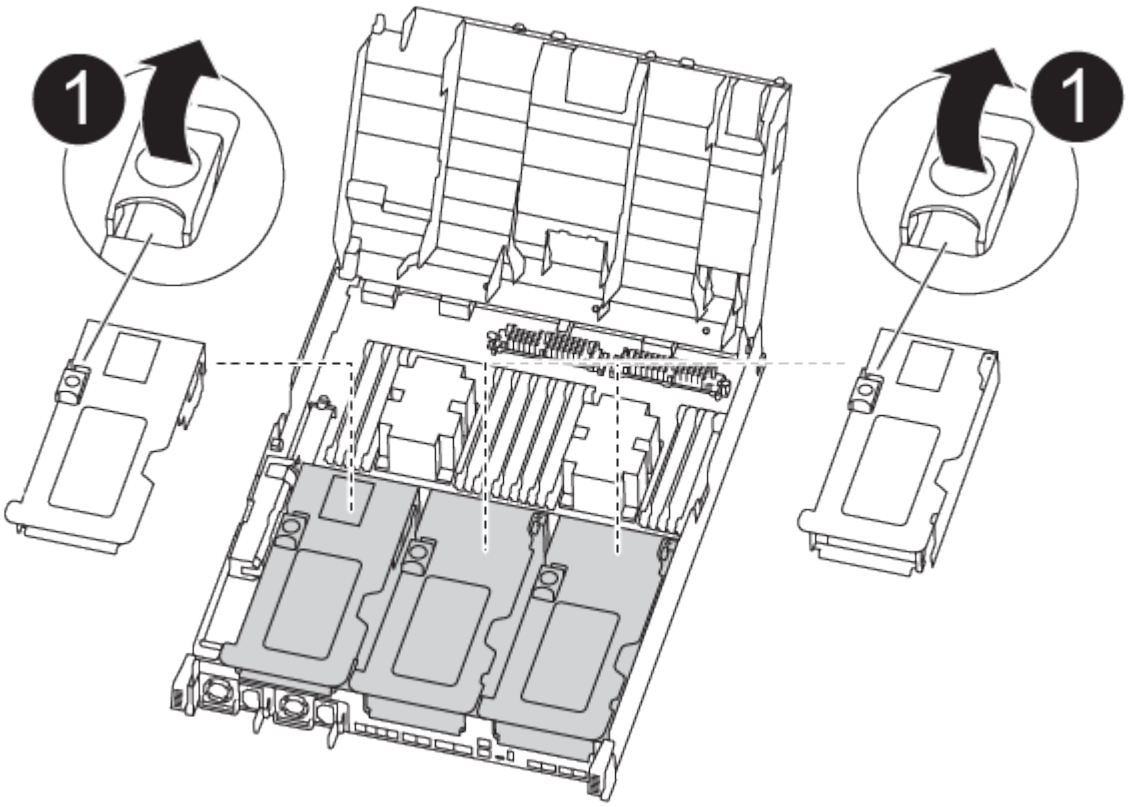
The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.
8. On the replacement controller module, open the air duct and remove the empty risers from the controller module using the animation, illustration, or the written steps:

   Removing the empty risers from the replacement controller module



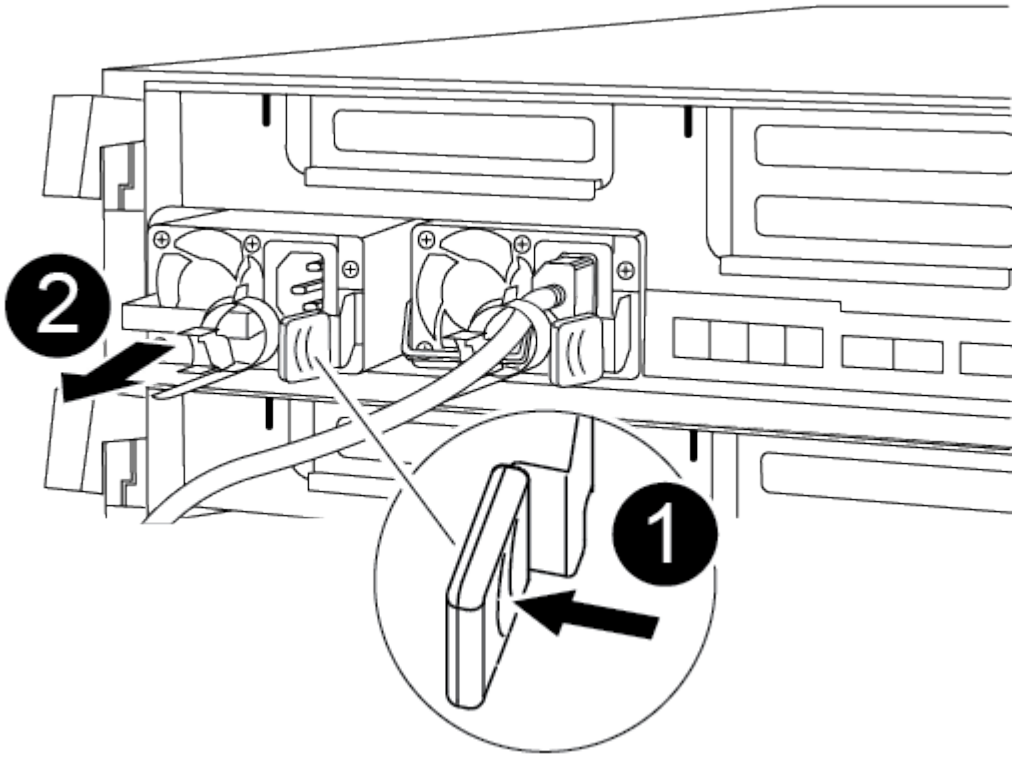
a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
c. Rotate the riser locking latch on the left side of riser 1 up and toward air duct, lift the riser up, and then set it aside.
d. Repeat the previous step for the remaining risers.

Step 2: Move the power supplies

You must move the power supply from the impaired controller module to the replacement controller module when you replace a controller module.

You can use the following animation, illustration, or the written steps to move the power supplies to the replacement controller module.

Animation - Move the power supplies



1. Remove the power supply:
   a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.
   b. Press the blue locking tab to release the power supply from the chassis.
   c. Using both hands, pull the power supply out of the chassis, and then set it aside.
2. Move the power supply to the new controller module, and then install it.
3. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   > To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

4. Repeat the preceding steps for any remaining power supplies.

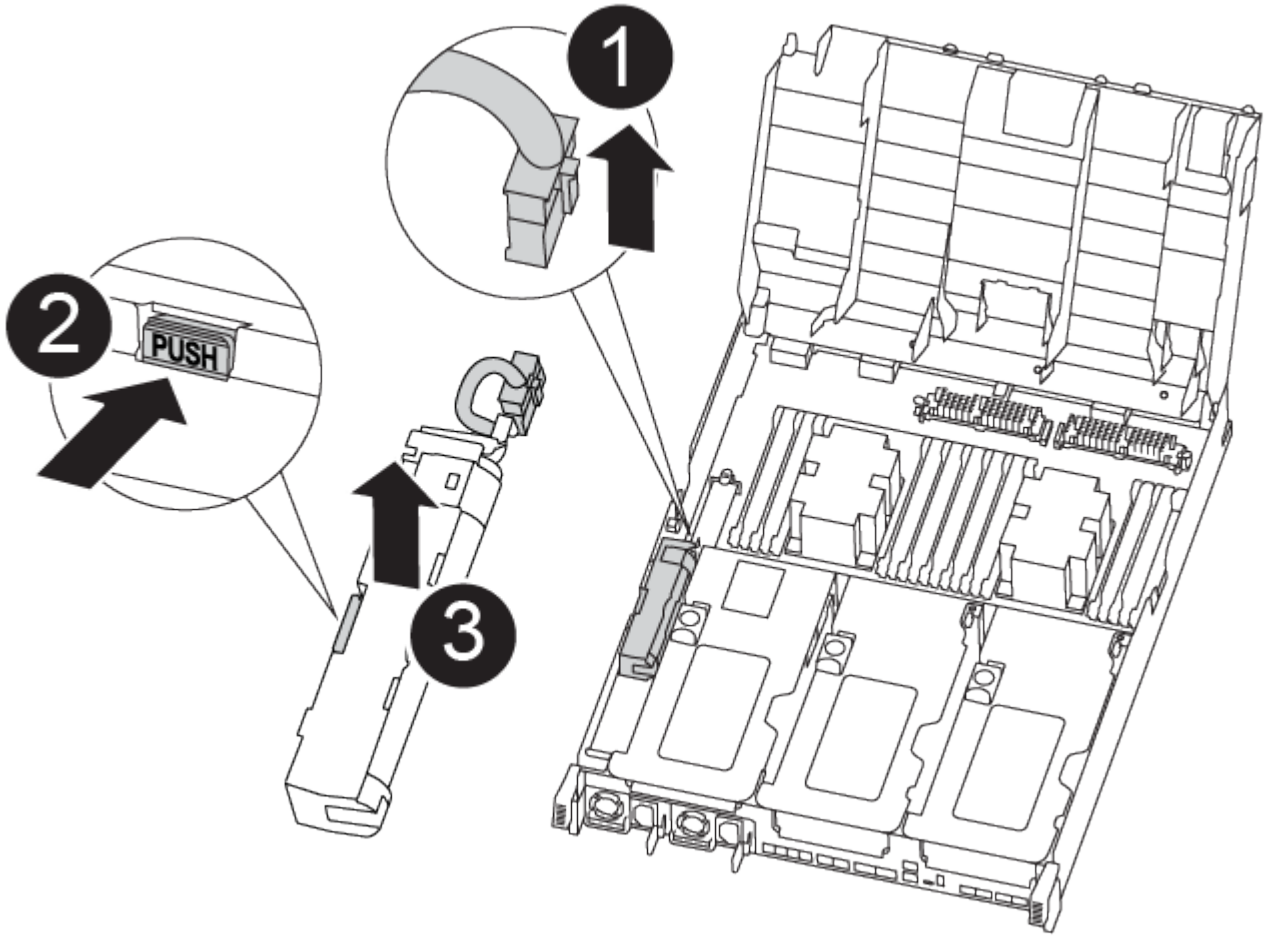
Step 3: Move the NVDIMM battery

To move the NVDIMM battery from the impaired controller module to the replacement controller module, you must perform a specific sequence of steps.

You can use the following animation, illustration, or the written steps to move the NVDIMM battery from the impaired controller module to the replacement controller module.

Animation - Move the NVDIMM battery

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the NVDIMM battery in the controller module.



1. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
2. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder and controller module.
3. Move the battery to the replacement controller module.
4. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.

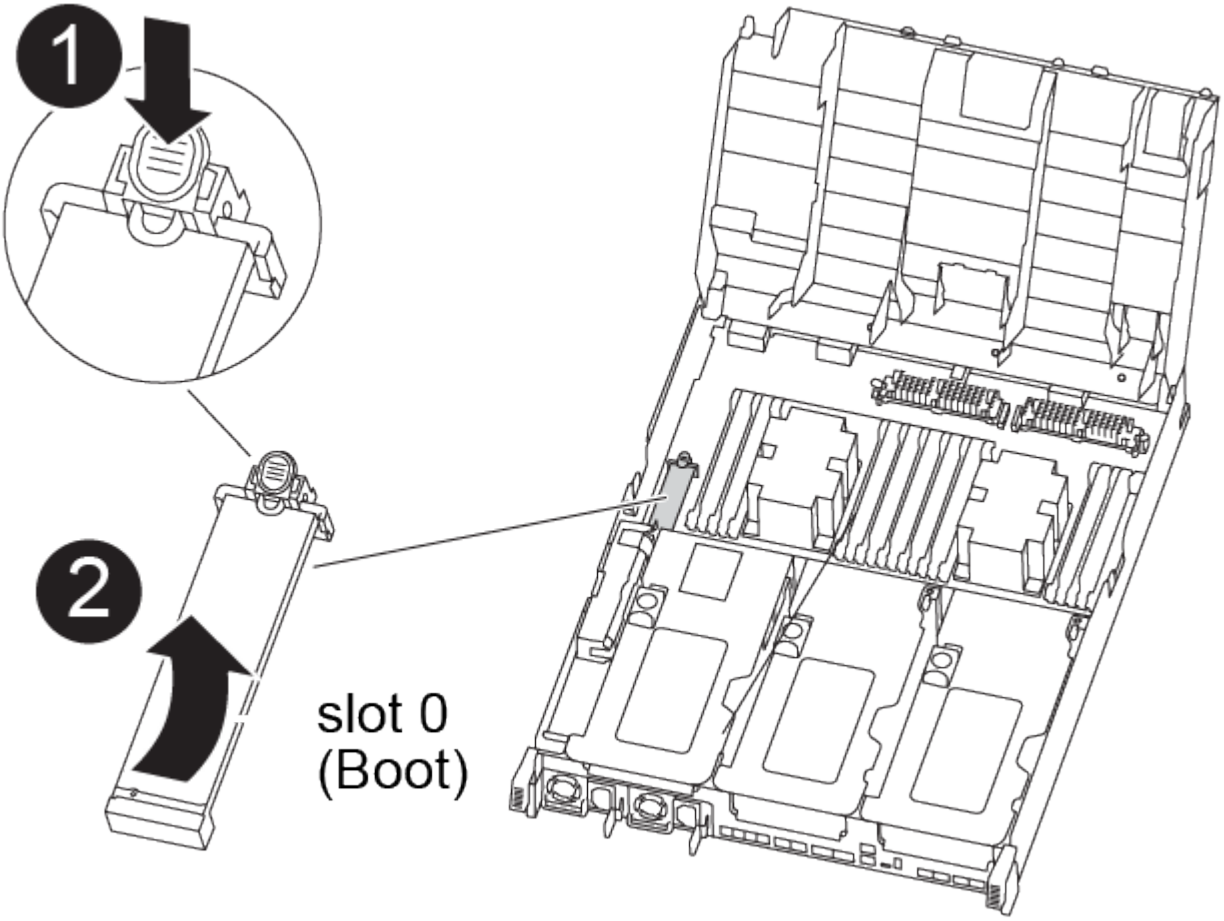
> Do not plug the battery cable back into the motherboard until instructed to do so.

Step 4: Move the boot media

You must locate the boot media, and then follow the directions to remove it from the impaired controller module and insert it into the replacement controller module.

You can use the following animation, illustration, or the written steps to move the boot media from the impaired controller module to the replacement controller module.

Animation - Move the boot media



1. Locate and remove the boot media from the controller module:
   a. Press the blue button at the end of the boot media until the lip on the boot media clears the blue button.
   b. Rotate the boot media up and gently pull the boot media out of the socket.
2. Move the boot media to the new controller module, align the edges of the boot media with the socket housing, and then gently push it into the socket.
3. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

4. Lock the boot media in place:
   a. Rotate the boot media down toward the motherboard.
   b. Press the blue locking button so that it is in the open position.
   c. Placing your fingers at the end of the boot media by the blue button, firmly push down on the boot media end to engage the blue locking button.

Step 5: Move the PCIe risers and mezzanine card

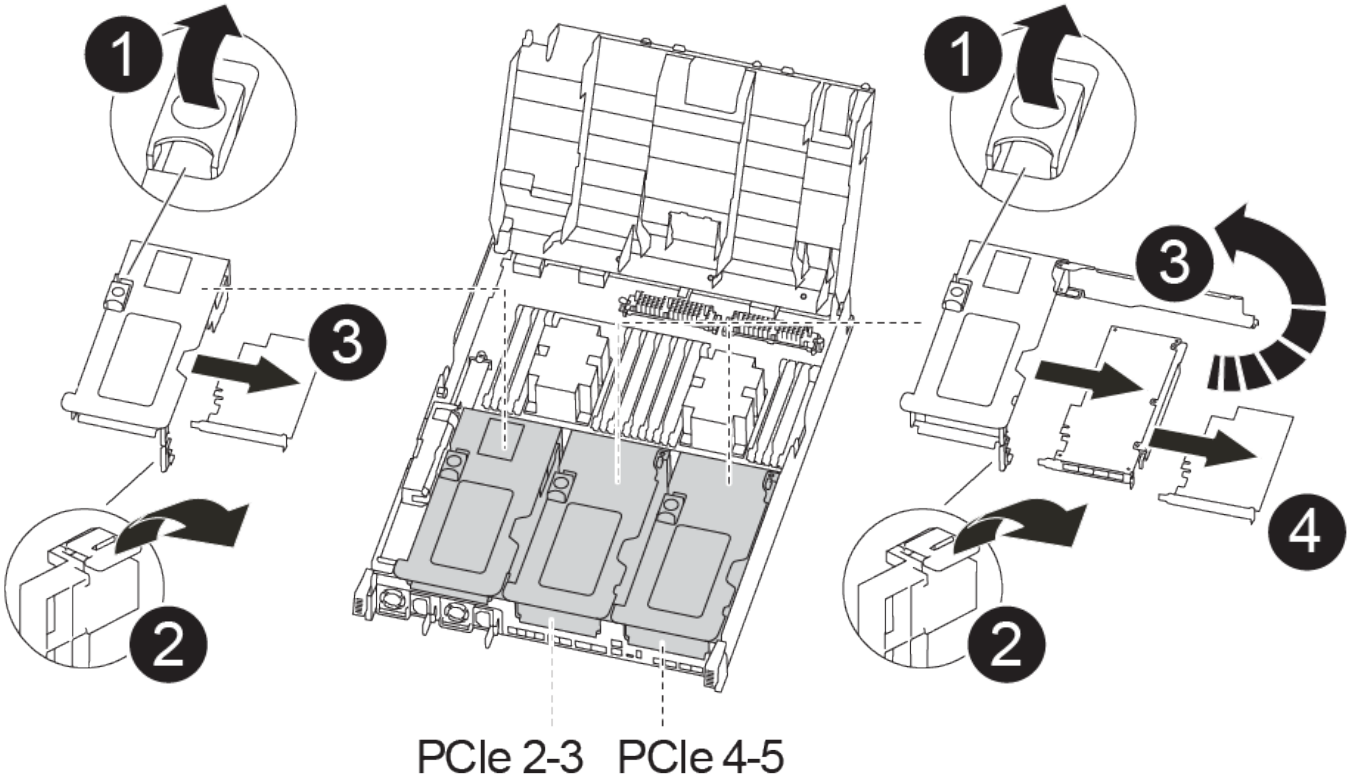
As part of the controller replacement process, you must move the PCIe risers and mezzanine card from the impaired controller module to the replacement controller module.

You can use the following animations, illustrations, the FUR map on the system, or the written steps to move the PCIe risers and mezzanine card from the impaired controller module to the replacement controller module.

> ⓘ You do not have to remove the PCIe cards from the risers. Transfer the risers, with the PCIe cards still installed, to the replacement controller module.

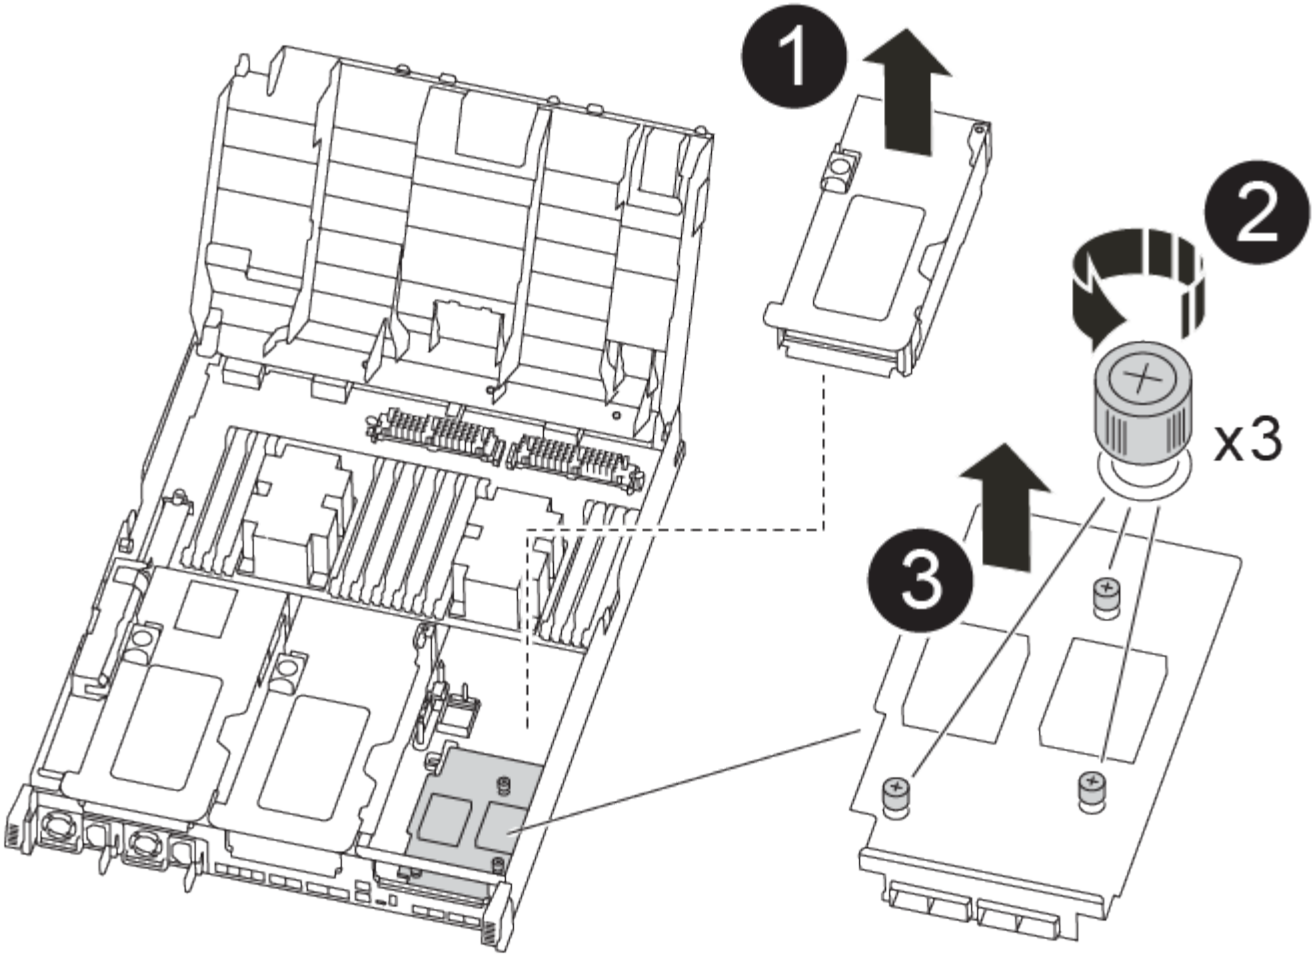
Moving PCIe riser 1 and 2 (left and middle risers):

Animation - Move PCI risers 1 and 2



Moving the mezzanine card and riser 3 (right riser):

Animation - Move the mezzanine card and riser 3



1. Move PCIe risers one and two from the impaired controller module to the replacement controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, and then move it to the replacement controller module.
   d. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins, push the riser squarely into the socket on the motherboard, and then rotate the latch down flush with the sheet metal on the riser.
   e. Repeat this step for riser number 2.
2. Remove riser number 3, remove the mezzanine card, and install both into the replacement controller module:
   a. Remove any SFP or QSFP modules that might be in the PCIe cards.
   b. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   c. Lift the riser up, and then set it aside on a stable, flat surface.
   d. Loosen the thumbscrews on the mezzanine card, and gently lift the card directly out of the socket, and then move it to the replacement controller module.

e. Install the mezzanine in the replacement controller and secure it with the thumbscrews.

f. Install the third riser in the replacement controller module.

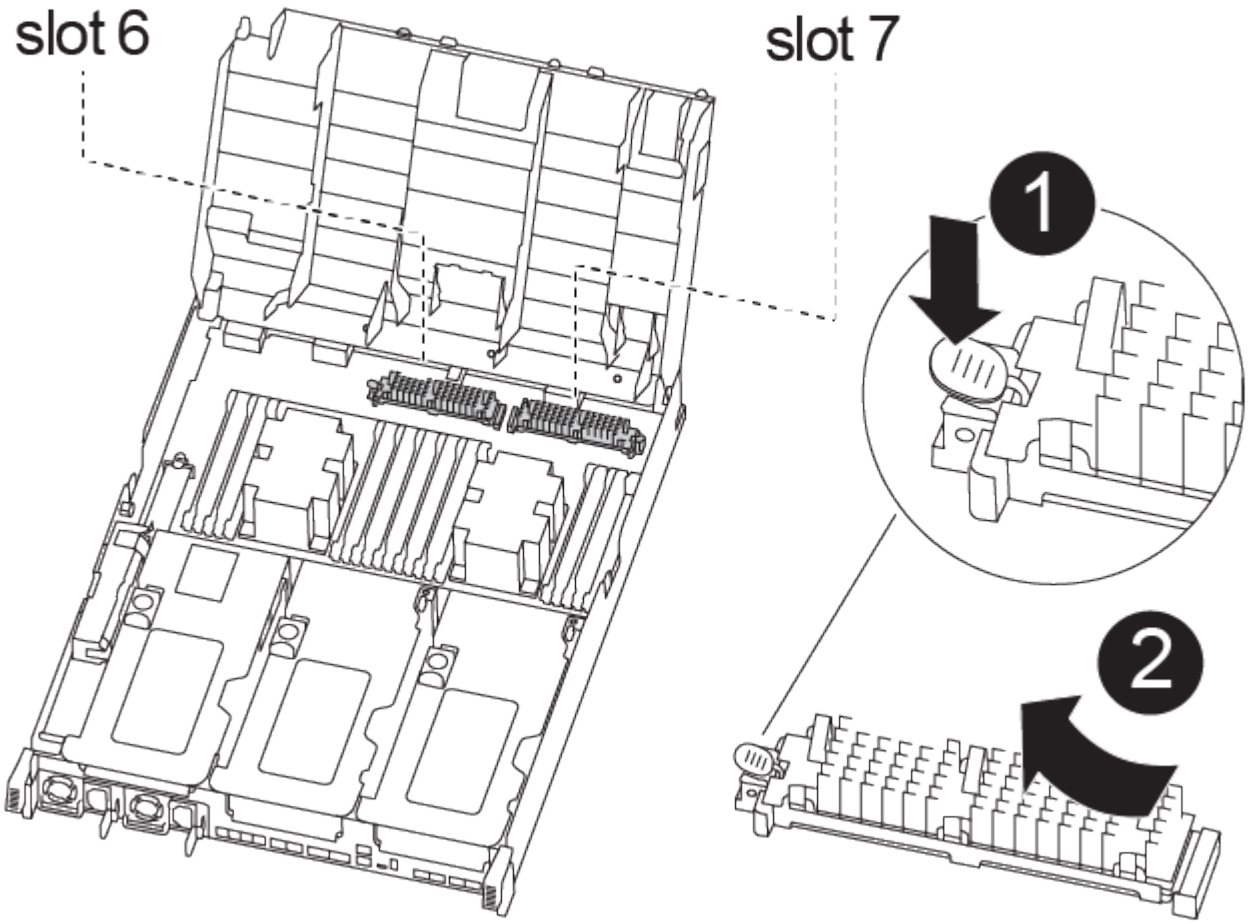
Step 6: Move caching modules

You must move the caching modules from the impaired controller modules to the replacement controller module when replacing a controller module.

> The Ver2 controller module has only one caching module socket in the FAS8300. FAS8700 does not have a VER2 controller module. The caching module functionality is not impacted by the socket removal.

You can use the following animation, illustration, or the written steps to move caching modules to the new controller module.

Animation - Move the caching modules



1. If you are not already grounded, properly ground yourself.
2. Move the caching modules from the impaired controller module to the replacement controller module:
   a. Press the blue release tab at the end of the caching module, rotate the module up, and then remove the module from the socket.
   b. Move the caching module to the same socket on the replacement controller module.

c. Align the edges of the caching module with the socket and gently insert the module as far into the socket as it will go.

d. Rotate the caching module downward toward the motherboard.

e. Placing your finger at the end of the caching module by the blue button, firmly push down on the caching module end, and then lift the locking button to lock the caching module in place.

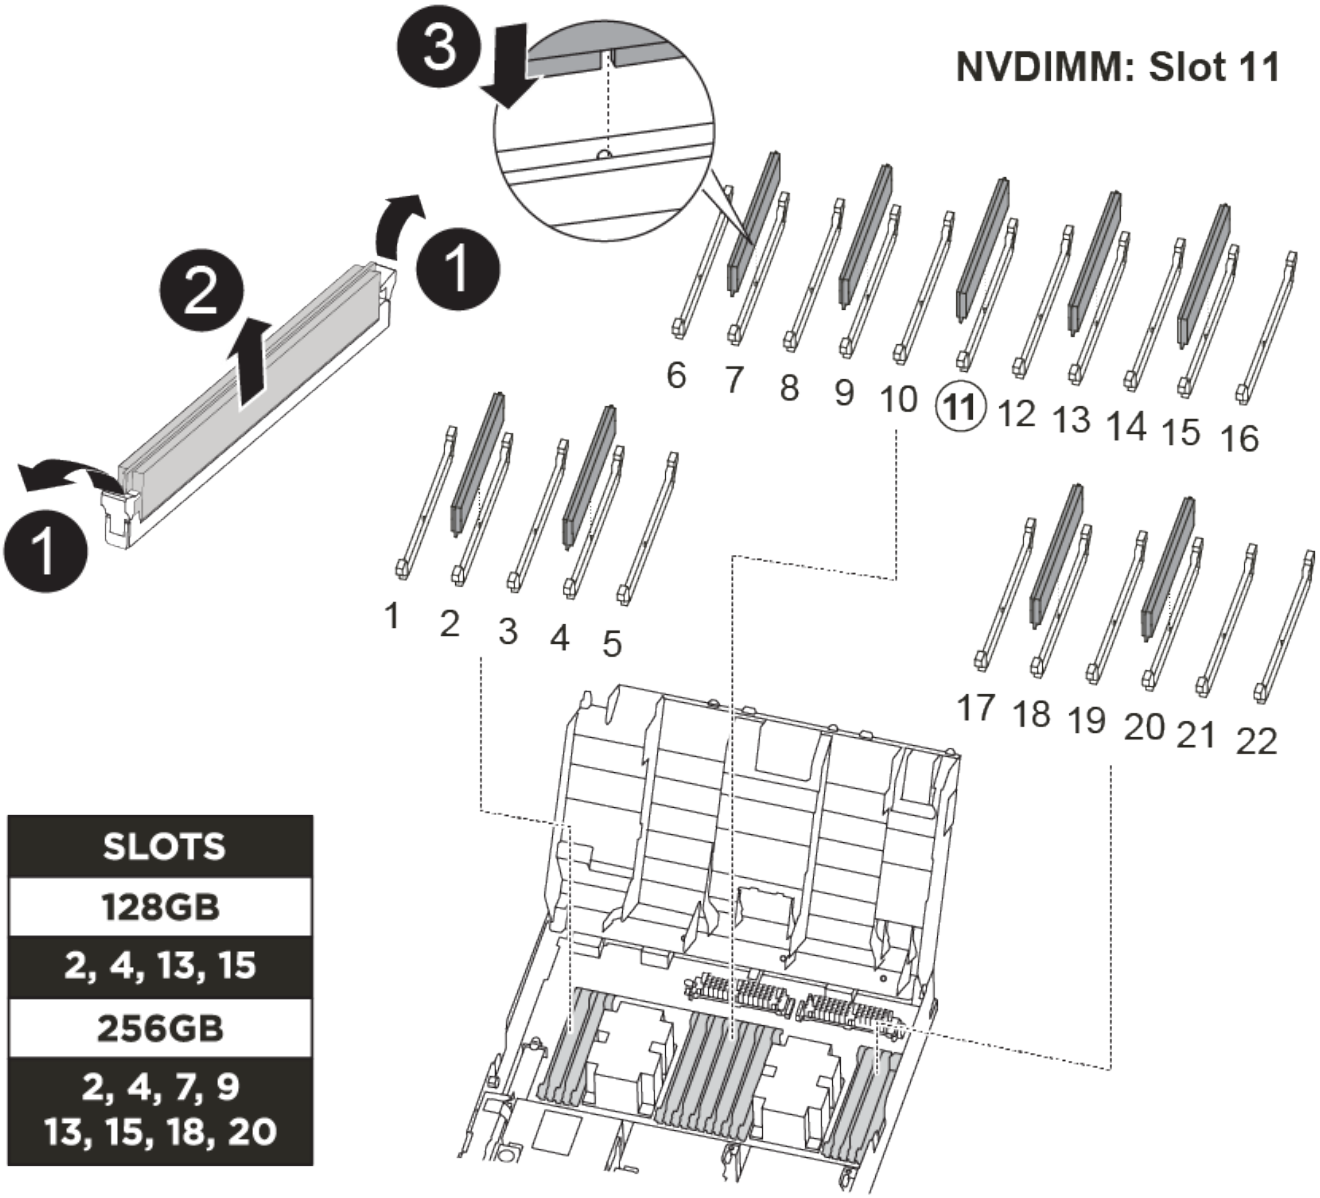
Step 7: Move the DIMMs

You need to locate the DIMMs, and then move them from the impaired controller module to the replacement controller module.

You must have the new controller module ready so that you can move the DIMMs directly from the impaired controller module to the corresponding slots in the replacement controller module.

You can use the following animation, illustration, or the written steps to move the DIMMs from the impaired controller module to the replacement controller module.

Animation - Move the DIMMs



1. Locate the DIMMs on your controller module.
2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Verify that the NVDIMM battery is not plugged into the new controller module.
4. Move the DIMMs from the impaired controller module to the replacement controller module:

> Make sure that you install the each DIMM into the same slot it occupied in the impaired controller module.

a. Eject the DIMM from its slot by slowly pushing apart the DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

> Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

b. Locate the corresponding DIMM slot on the replacement controller module.

c. Make sure that the DIMM ejector tabs on the DIMM socket are in the open position, and then insert the DIMM squarely into the socket.

   The DIMMs fit tightly in the socket, but should go in easily. If not, realign the DIMM with the socket and reinsert it.

d. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the socket.

e. Repeat these substeps for the remaining DIMMs.

5. Plug the NVDIMM battery into the motherboard.

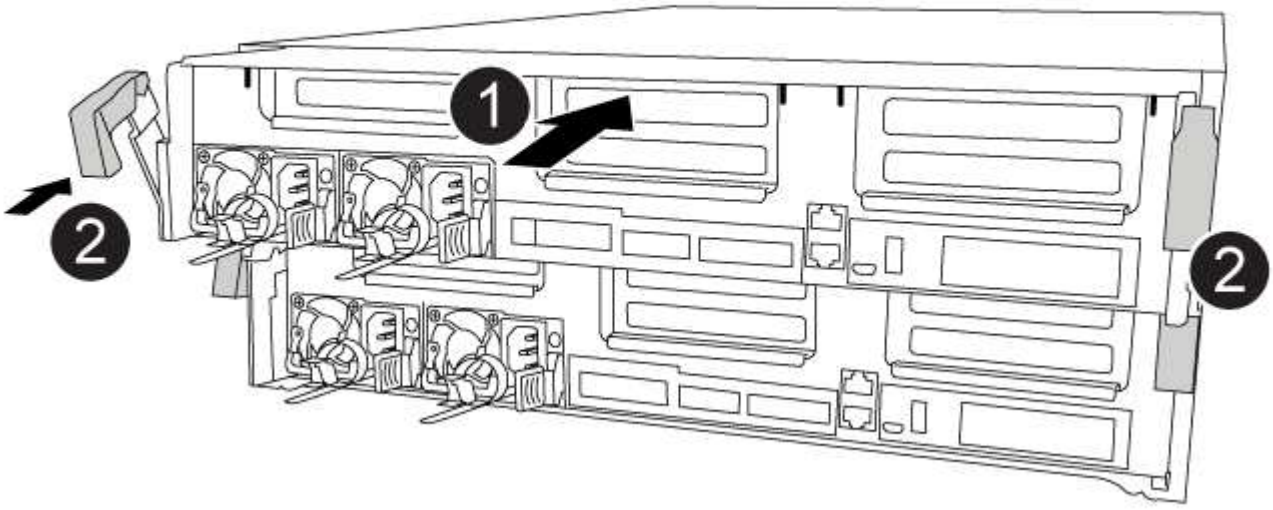
   Make sure that the plug locks down onto the controller module.

Step 8: Install the controller module

After all of the components have been moved from the impaired controller module to the replacement controller module, you must install the replacement controller module into the chassis, and then boot it to Maintenance mode.

You can use the following animation, illustration, or the written steps to install the replacement controller module in the chassis.

Animation - Install the controller module



1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

   f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.
   g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

Restore and verify the system configuration - FAS8300 and FAS8700

After completing the hardware replacement and booting to Maintenance mode, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time after replacing the controller

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.

2. On the *healthy* node, check the system time: `cluster date show`

   The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

   The date and time are given in GMT.

Step 2: Verify and set the HA state of the controller module

You must verify the `HA` state of the controller module and, if necessary, update the state to match your system configuration.

1. In Maintenance mode from the new controller module, verify that all components display the same `HA` state: `ha-config show`

   The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`

   The value for `HA-state` can be one of the following:

   - `ha`
   - `mcc`
   - `mcc-2n`
   - `mccip`
   - `non-ha`

3. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Recable the system and reassign disks - FAS8300 and FAS8700

You must complete a series of tasks before restoring your system to full operation.

Step 1: Recable the system

Verify the controller module's storage and network connections by using Active IQ Config Advisor.

Steps

1. Download and install Config Advisor.

2. Enter the information for the target system, and then click Collect Data.
3. Click the Cabling tab, and then examine the output. Make sure that all disk shelves are displayed and all disks appear in the output, correcting any cabling issues you find.
4. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. In a stand-alone system, you must manually reassign the ID to the disks.

You must use the correct procedure for your configuration:

Controller redundancy	Then use this procedure…
HA pair	Option 1: Verify the system ID change on an HA system
Two-node MetroCluster configuration	Option 2: Manually reassign the system ID on systems in a two-node MetroCluster configuration

Option 1: Verify the system ID change on an HA system

You must confirm the system ID change when you boot the *replacement* controller and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* controller is in Maintenance mode (showing the `*>` prompt, exit Maintenance mode and go to the LOADER prompt: `halt`
2. From the LOADER prompt on the *replacement* controller, boot the controller, entering `y` if you are prompted to override the system ID due to a system ID mismatch:
3. Wait until the `Waiting for giveback…` message is displayed on the *replacement* controller console and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, node2 has undergone replacement and has a new system ID of 151759706.

```
node1> `storage failover show`
                                    Takeover
Node                Partner         Possible     State Description
------------      ------------      --------
-------------------------------------
node1               node2             false        System ID changed on
partner (Old:
                                                    151759755, New:
151759706), In takeover
node2               node1             -            Waiting for giveback
(HA mailboxes)
```

4. From the healthy controller, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. Save any coredumps: `system node run -node local-node-name partner savecore`
   c. Wait for the `savecore`command to complete before issuing the giveback.

      You can enter the following command to monitor the progress of the savecore command: `system node run -node local-node-name partner savecore -s`

   d. Return to the admin privilege level: `set -privilege admin`
5. If your storage system has Storage or Volume Encryption configured, you must restore Storage or Volume Encryption functionality by using one of the following procedures, depending on whether you are using onboard or external key management:
   - Restore onboard key management encryption keys
   - Restore external key management encryption keys
6. Give back the controller:
   a. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* controller takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      Find the High-Availability Configuration Guide for your version of ONTAP 9

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed

on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> `storage disk show -ownership`

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

Option 2: Manually reassign the system ID on systems in a two-node MetroCluster configuration

In a two-node MetroCluster configuration running ONTAP, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies only to systems in a two-node MetroCluster configuration running ONTAP.

You must be sure to issue the commands in this procedure on the correct node:

- The *impaired* node is the node on which you are performing maintenance.
- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the DR partner of the impaired node.

Steps

1. If you have not already done so, reboot the *replacement* node, interrupt the boot process by entering `Ctrl-C`, and then select the option to boot to Maintenance mode from the displayed menu.

   You must enter `Y` when prompted to override the system ID due to a system ID mismatch.

2. View the old system IDs from the healthy node: `metrocluster node show -fields node-systemid`,dr-partner-systemid`

   In this example, the Node_B_1 is the old node, with the old system ID of 118073209:

```
dr-group-id cluster         node                 node-systemid dr-
partner-systemid
 ----------- -------------------- ------------------- ------------
------------------
 1           Cluster_A            Node_A_1             536872914
118073209
 1           Cluster_B            Node_B_1             118073209
536872914
 2 entries were displayed.
```

3. View the new system ID at the Maintenance mode prompt on the impaired node: `disk show`

   In this example, the new system ID is 118065481:

```
Local System ID: 118065481
    ...
    ...
```

4. Reassign disk ownership (for FAS systems), by using the system ID information obtained from the disk show command: `disk reassign -s old system ID`

   In the case of the preceding example, the command is: `disk reassign -s 118073209`

   You can respond `Y` when prompted to continue.

5. Verify that the disks were assigned correctly: `disk show -a`

   Verify that the disks belonging to the *replacement* node show the new system ID for the *replacement* node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

```
*> disk show -a
Local System ID: 118065481

  DISK     OWNER                POOL   SERIAL NUMBER  HOME
-------    -------------        -----  -------------  -------------
disk_name   system-1  (118065481) Pool0  J8Y0TDZC       system-1
(118065481)
disk_name   system-1  (118065481) Pool0  J8Y09DXC       system-1
(118065481)
.
.
.
```

6. From the healthy node, verify that any coredumps are saved:

a. Change to the advanced privilege level: `set -privilege advanced`

   You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

b. Verify that the coredumps are saved: `system node run -node local-node-name partner savecore`

   If the command output indicates that savecore is in progress, wait for savecore to complete before issuing the giveback. You can monitor the progress of the savecore using the `system node run -node local-node-name partner savecore -s command`.</info>.

c. Return to the admin privilege level: `set -privilege admin`

7. If the *replacement* node is in Maintenance mode (showing the *> prompt), exit Maintenance mode and go to the LOADER prompt: `halt`
8. Boot the *replacement* node: `boot_ontap`
9. After the *replacement* node has fully booted, perform a switchback: `metrocluster switchback`
10. Verify the MetroCluster configuration: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify the operation of the MetroCluster configuration in Data ONTAP:
    a. Check for any health alerts on both clusters: `system health alert show`
    b. Confirm that the MetroCluster is configured and in normal mode: `metrocluster show`
    c. Perform a MetroCluster check: `metrocluster check run`
    d. Display the results of the MetroCluster check: `metrocluster check show`
    e. Run Config Advisor. Go to the Config Advisor page on the NetApp Support Site at support.netapp.com/NOW/download/tools/config_advisor/.

       After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

12. Simulate a switchover operation:
    a. From any node's prompt, change to the advanced privilege level: `set -privilege advanced`

You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (*>).

b. Perform the switchback operation with the -simulate parameter: `metrocluster switchover -simulate`

c. Return to the admin privilege level: `set -privilege admin`

Complete system restoration - FAS8300 and FAS8700

To restore your system to full operation, you must restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Install licenses for the replacement controller in ONTAP

You must install new licenses for the *replacement* node if the impaired node was using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses continue to be available to the *replacement* node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed.

Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the *replacement* node as soon as possible.

Before you begin

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

> If your system was initially running ONTAP 9.10.1 or later, use the procedure documented in Post Motherboard Replacement Process to update Licensing on a AFF/FAS system. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`

b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`

3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured        switchover
   Remote: cluster_A configured        waiting-for-switchback
   ```

   The switchback operation is complete when the clusters are in the `normal` state.:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    ------------------     ---------
    Local: cluster_B configured            normal
   Remote: cluster_A configured            normal
   ```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.
6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - FAS8300 and FAS8700

You must replace a DIMM in the controller when your storage system encounters errors such as, excessive CECC (Correctable Error Correction Codes) errors that are based on Health Monitor alerts or uncorrectable ECC errors, typically caused by a single DIMM failure preventing the storage system from booting ONTAP.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

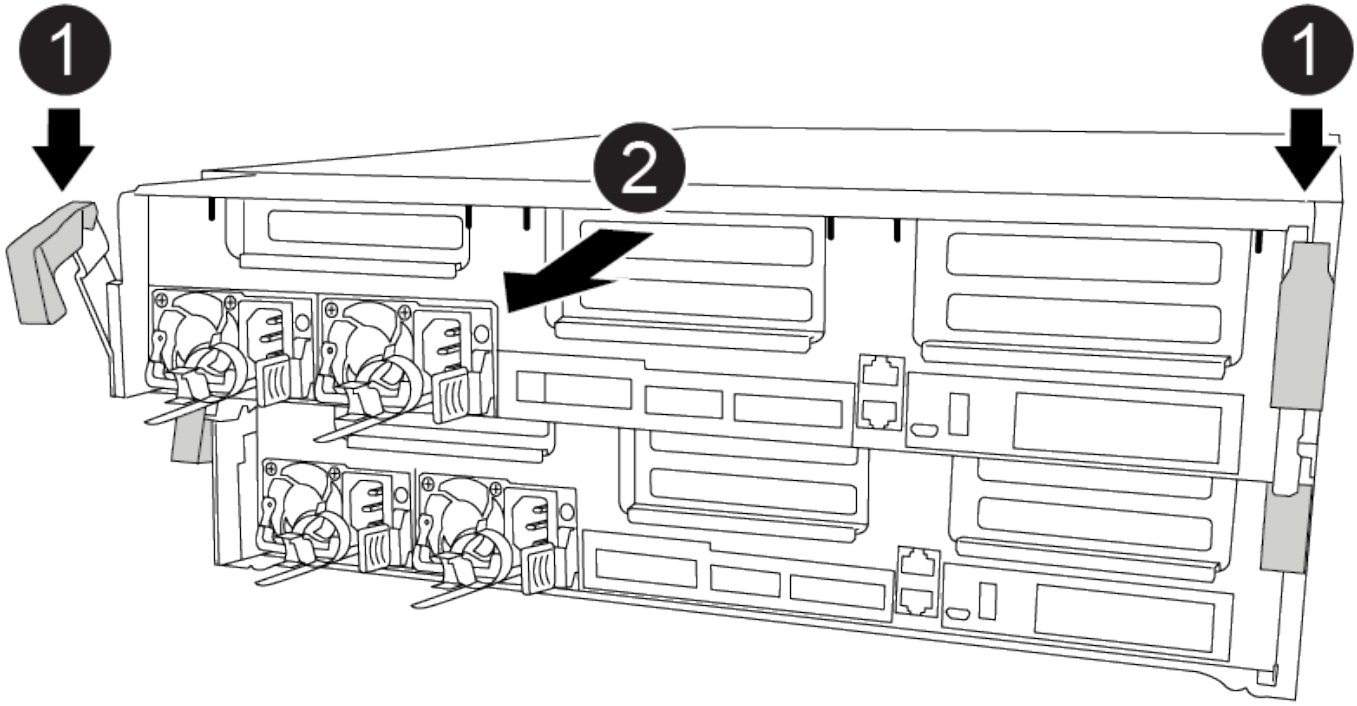
8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module



Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
7. Place the controller module on a stable, flat surface.

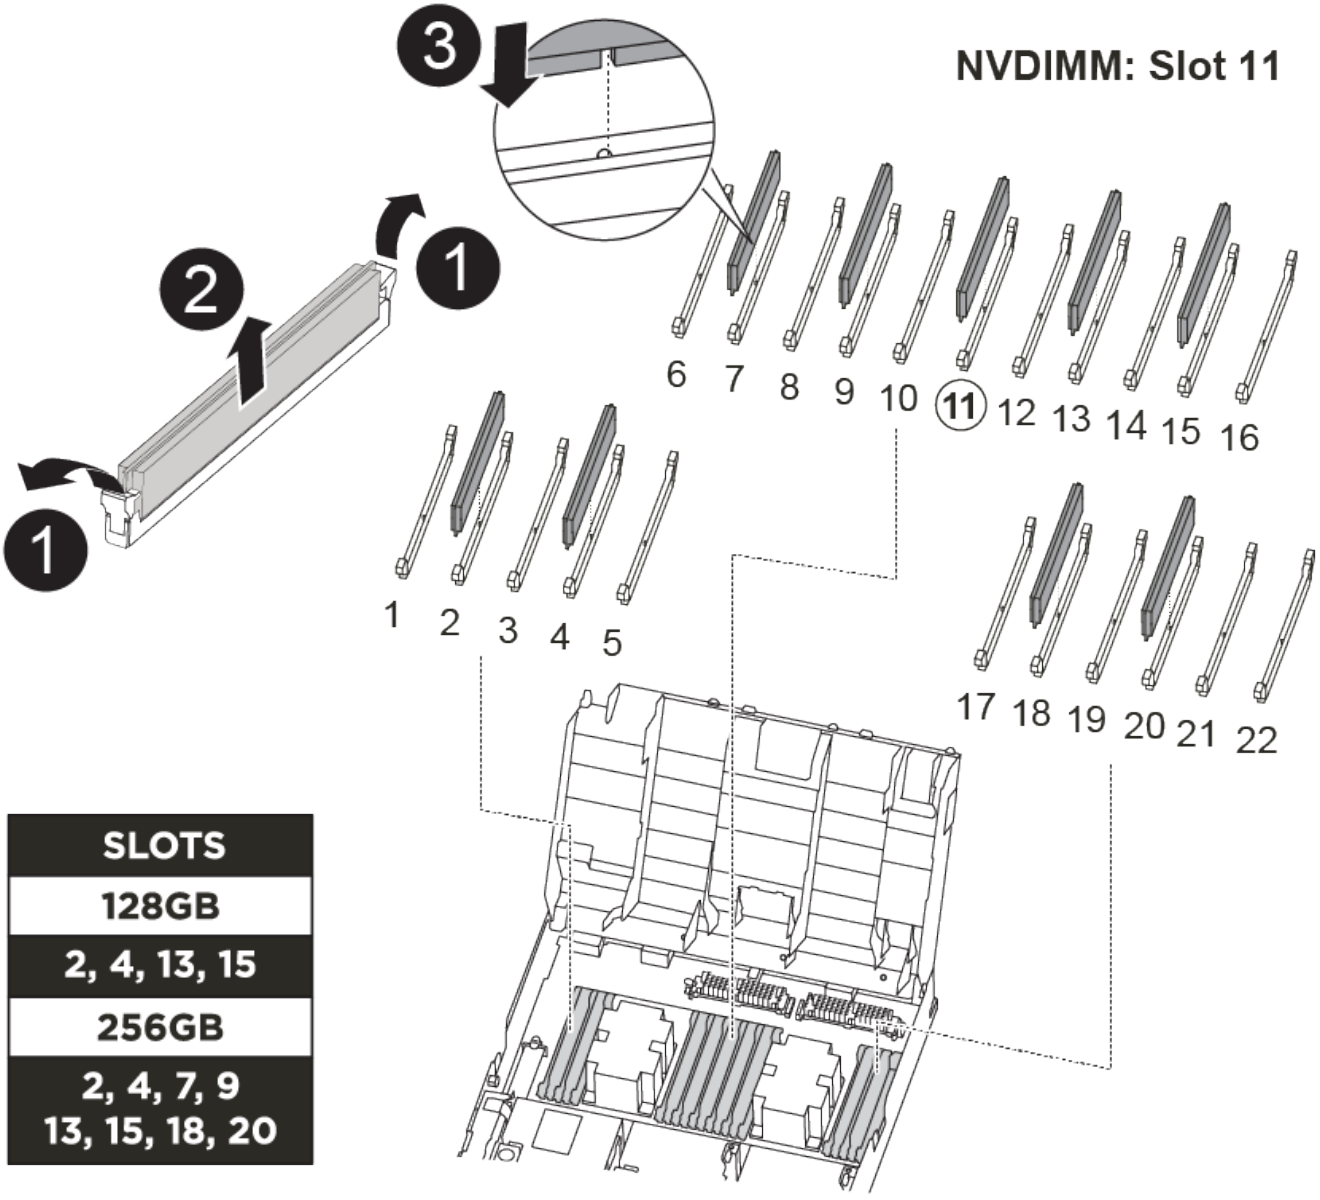
Step 3: Replace system DIMMs

Replacing a system DIMM involves identifying the target DIMM through the associated error message, locating the target DIMM using the FRU map on the air duct, and then replacing the DIMM.

You can use the following animation, illustration, or the written steps to replace a system DIMM.

> The animation and illustration shows empty slots for sockets without DIMMs. These empty sockets are populated with blanks.

Animation - Replace a system DIMM



The number and location of DIMMS in your system depends on the model of your system. Refer to FRU map on the air duct for more information.

- If you have a FAS8300 system, the system DIMMs are located in sockets 2, 4, 13, and 15.
- If you have a FAS8700 system, the system DIMMs are located in slots 2, 4, 7, 9, 13, 15, 18, and 20.

- The NVDIMM is located in slot 11.

Steps

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the DIMMs on your controller module.
3. Note the orientation of the DIMM in the socket so that you can insert the replacement DIMM in the proper orientation.
4. Eject the DIMM from its socket by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the socket.

   Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

5. Remove the replacement DIMM from the antistatic shipping bag, hold the DIMM by the corners, and align it to the slot.

   The notch among the pins on the DIMM should line up with the tab in the socket.

6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

   The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

   Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

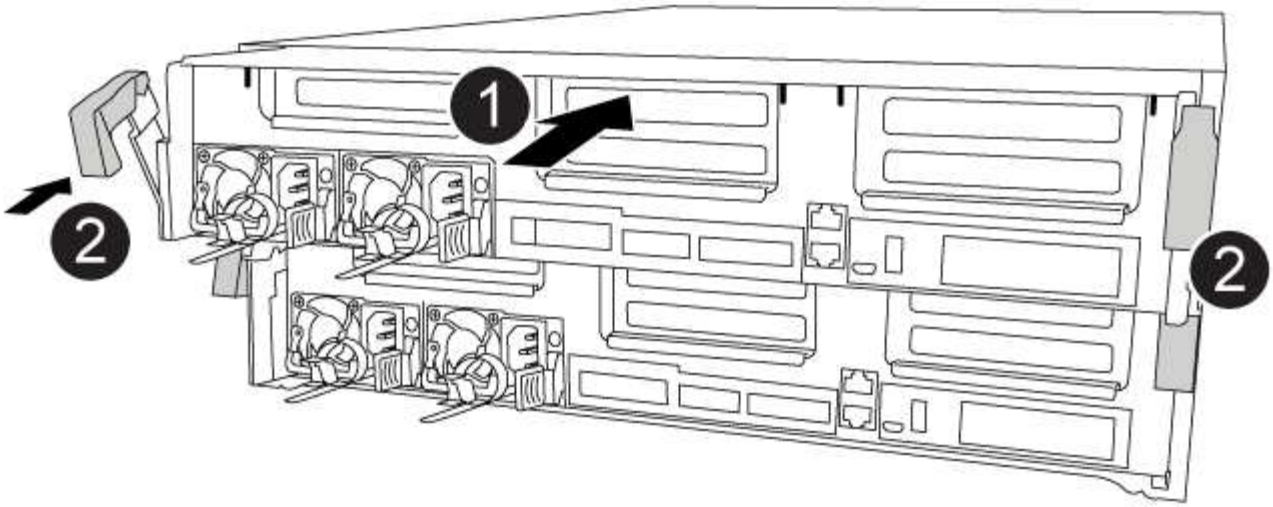
7. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
8. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis.

You can use the following animation, drawing, or the written steps to install the controller module in the chassis.

Animation - Install the controller module



Steps

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

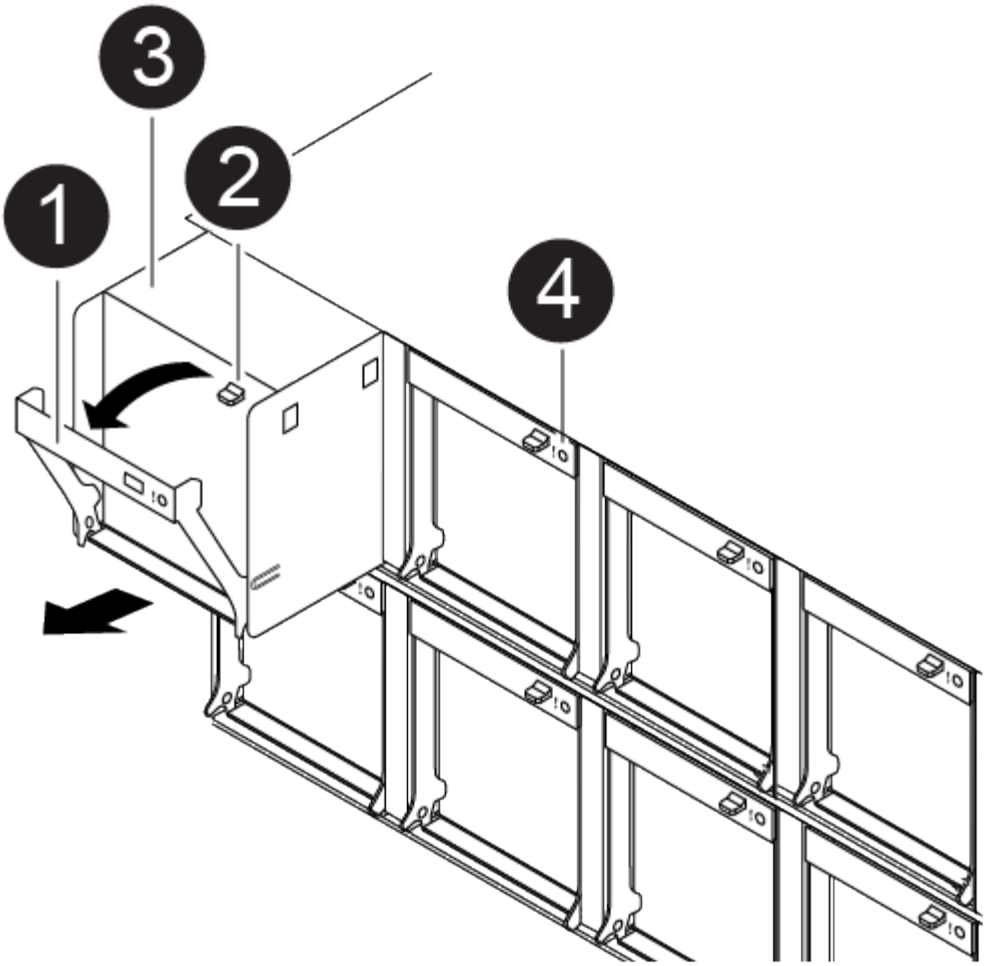
Hot-swap a fan module - FAS8300 and FAS8700

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

> You must replace the fan module within two minutes of removing it from the chassis. System airflow is disrupted and the controller module or modules shut down after two minutes to avoid overheating.

You can use the following animation, illustration, or the written steps to hot-swap a fan module.

Animation - Replace a fan



Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

   The fan module moves a little bit away from the chassis.

5. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

6. Set the fan module aside.
7. Insert the replacement fan module into the chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

   The Attention LED should not be lit after the fan is seated and has spun up to operational speed.

10. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
11. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace an NVDIMM - FAS8300 and FAS8700

You must replace the NVDIMM in the controller module when your system registers that the flash lifetime is almost at an end or that the identified NVDIMM is not healthy in general; failure to do so causes a system panic.

All other components in the system must be functioning properly; if not, you must contact technical support.

You must replace the failed component with a replacement FRU component you received from your provider.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

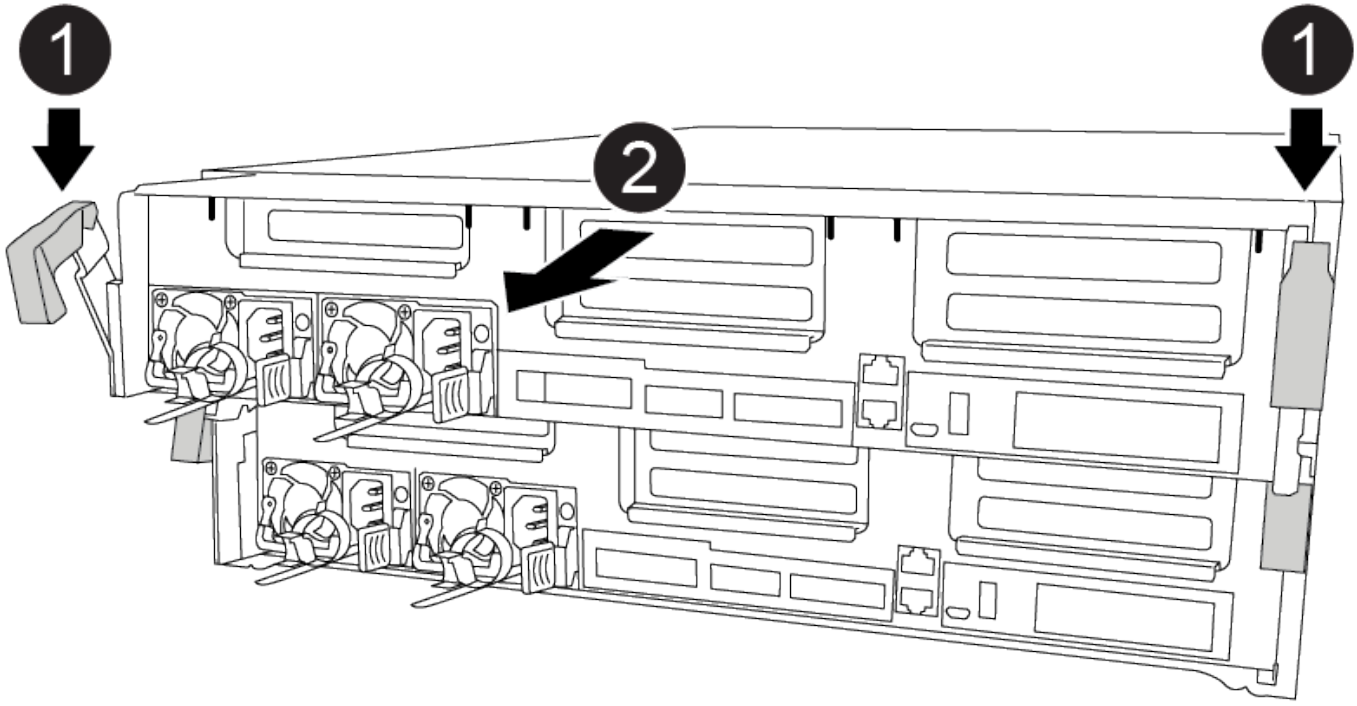
8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following , illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module



Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.

6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

7. Place the controller module on a stable, flat surface.

Step 3: Replace the NVDIMM

To replace the NVDIMM, you must locate it in the controller module using the FRU map on top of the air duct the FRU Map on the top of the slot 1 riser.

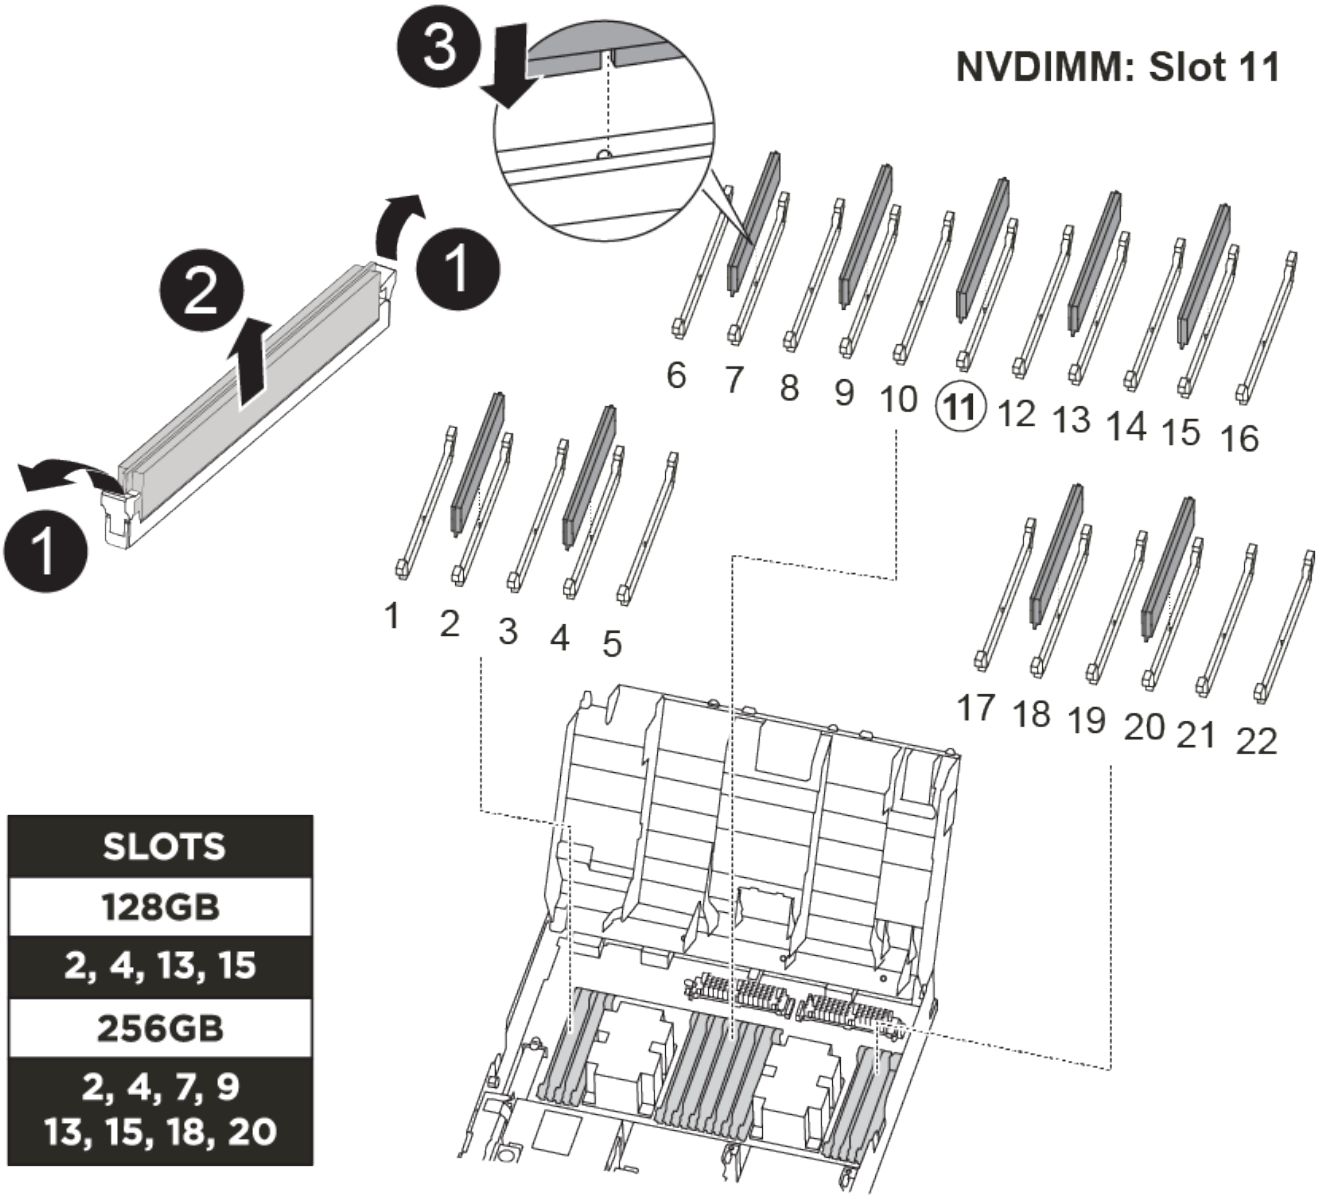
- The NVDIMM LED blinks while destaging contents when you halt the system. After the destage is complete, the LED turns off.
- Although the contents of the NVDIMM is encrypted, it is a best practice to erase the contents of the NVDIMM before replacing it. For more information, see the Statement of Volatility on the NetApp Support Site.

> You must log into the NetApp Support Site to display the *Statement of Volatility* for your system.

You can use the following animation, illustration, or the written steps to replace the NVDIMM.

> The animation and illustration show empty slots for sockets without DIMMs. These empty sockets are populated with blanks.

Animation - Replace the NVDIMM



Steps

1. Open the air duct and then locate the NVDIMM in slot 11 on your controller module.

   The NVDIMM looks significantly different than system DIMMs.

2. Eject the NVDIMM from its slot by slowly pushing apart the two NVDIMM ejector tabs on either side of the NVDIMM, and then slide the NVDIMM out of the socket and set it aside.

   Carefully hold the NVDIMM by the edges to avoid pressure on the components on the NVDIMM circuit board.

3. Remove the replacement NVDIMM from the antistatic shipping bag, hold the NVDIMM by the corners, and then align it to the slot.

   The notch among the pins on the NVDIMM should line up with the tab in the socket.

4. Locate the slot where you are installing the NVDIMM.
5. Insert the NVDIMM squarely into the slot.

   The NVDIMM fits tightly in the slot, but should go in easily. If not, realign the NVDIMM with the slot and reinsert it.

   Visually inspect the NVDIMM to verify that it is evenly aligned and fully inserted into the slot.

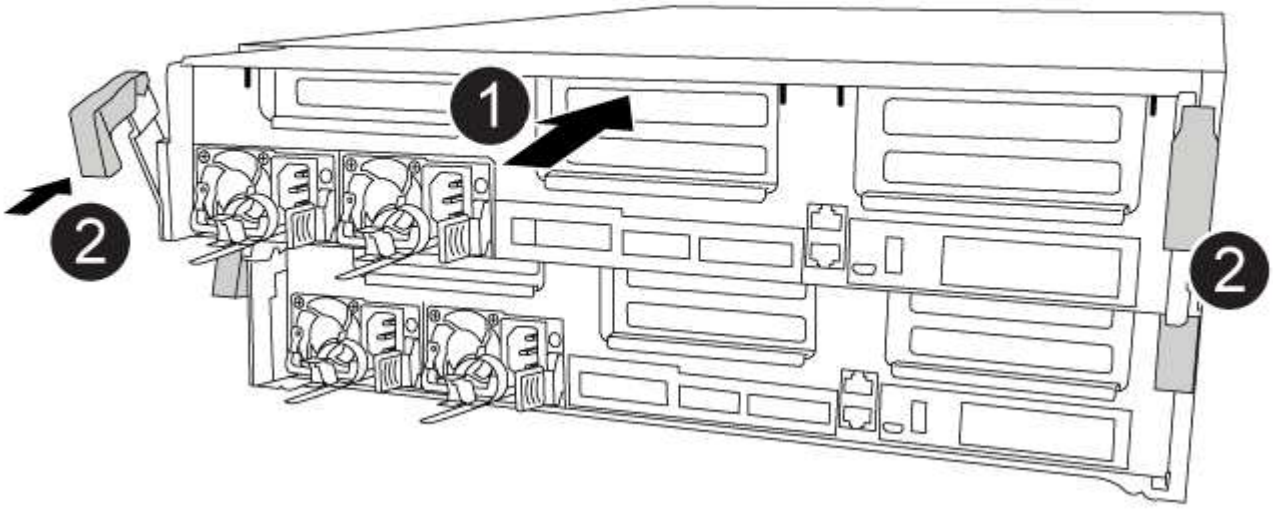
6. Push carefully, but firmly, on the top edge of the NVDIMM until the ejector tabs snap into place over the notches at the ends of the NVDIMM.
7. Close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module



Steps

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      If your system stops at the boot menu, select the option to boot to LOADER.

   f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` `*impaired_node_name*`
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

   The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the NVDIMM battery - FAS8300 and FAS8700

To replace the NVDIMM battery, you must remove the controller module, remove the battery, replace the battery, and then reinstall the controller module.

All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

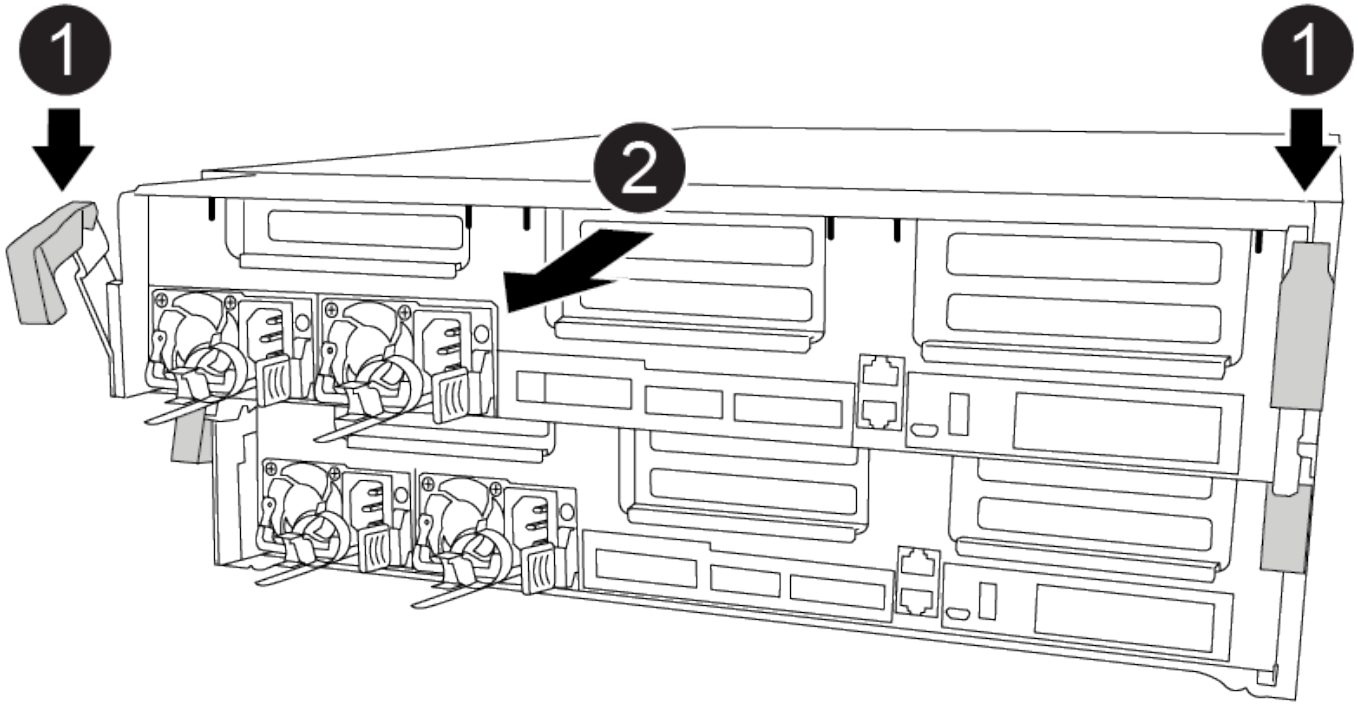
8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module



Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
7. Place the controller module on a stable, flat surface.

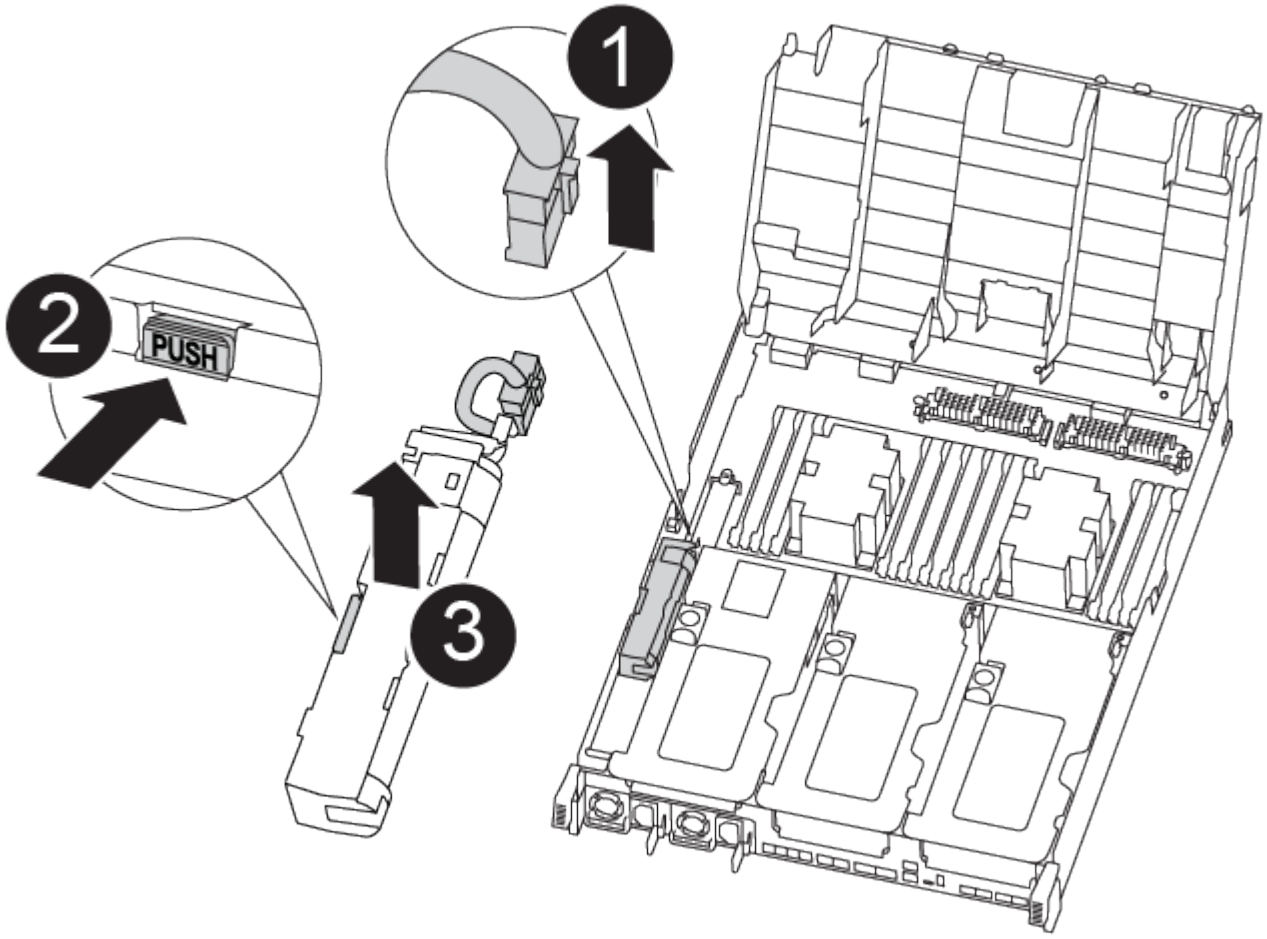
Step 3: Replace the NVDIMM battery

To replace the NVDIMM battery, you must remove the failed battery from the controller module and install the replacement battery into the controller module. See the FRU map inside the controller module to locate the NVDIMM battery.

The NVDIMM LED blinks while destaging contents when you halt the system. After the destage is complete, the LED turns off.

You can use the following animation, illustration, or the written steps to replace the NVDIMM battery.

Animation - Replace the NVDIMM battery



Steps

1. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
2. Locate the NVDIMM battery in the controller module.
3. Locate the battery plug and squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Grasp the battery and press the blue locking tab marked PUSH, and then lift the battery out of the holder

and controller module.

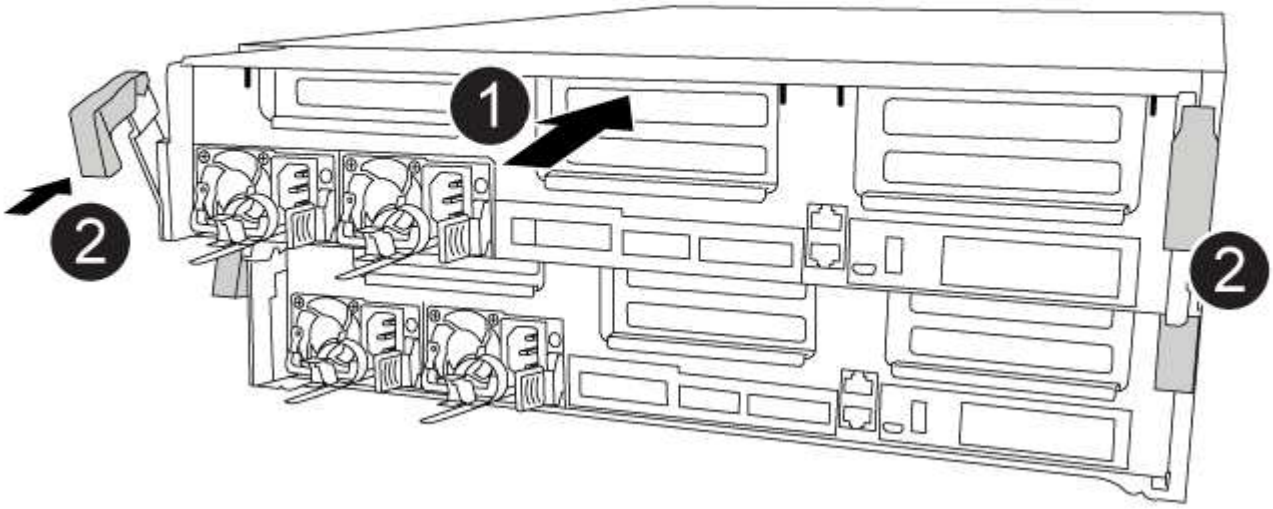
5. Remove the replacement battery from its package.
6. Align the battery module with the opening for the battery, and then gently push the battery into slot until it locks into place.
7. Plug the battery plug back into the controller module, and then close the air duct.

Step 4: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module



Steps

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Cable the management and console ports only, so that you can access the system to perform the tasks in the following sections.

   You will connect the rest of the cables to the controller module later in this procedure.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.

Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

   The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

d. If you have not already done so, reinstall the cable management device.

e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

   If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.

Step 5: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*

3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    -------------------     ---------
    Local: cluster_B configured        switchover
   Remote: cluster_A configured        waiting-for-switchback
   ```

   The switchback operation is complete when the clusters are in the `normal` state.:

   ```
   cluster_B::> metrocluster show
   Cluster              Configuration State    Mode
   --------------------    -------------------     ---------
    Local: cluster_B configured            normal
   Remote: cluster_A configured            normal
   ```

   If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a PCIe or mezzanine card - FAS8300 and FAS8700

To replace a PCIe or mezzanine card, you must disconnect the cables and any SFP and QSFP modules from the cards, replace the failed PCIe or mezzanine card, and then recable the cards.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

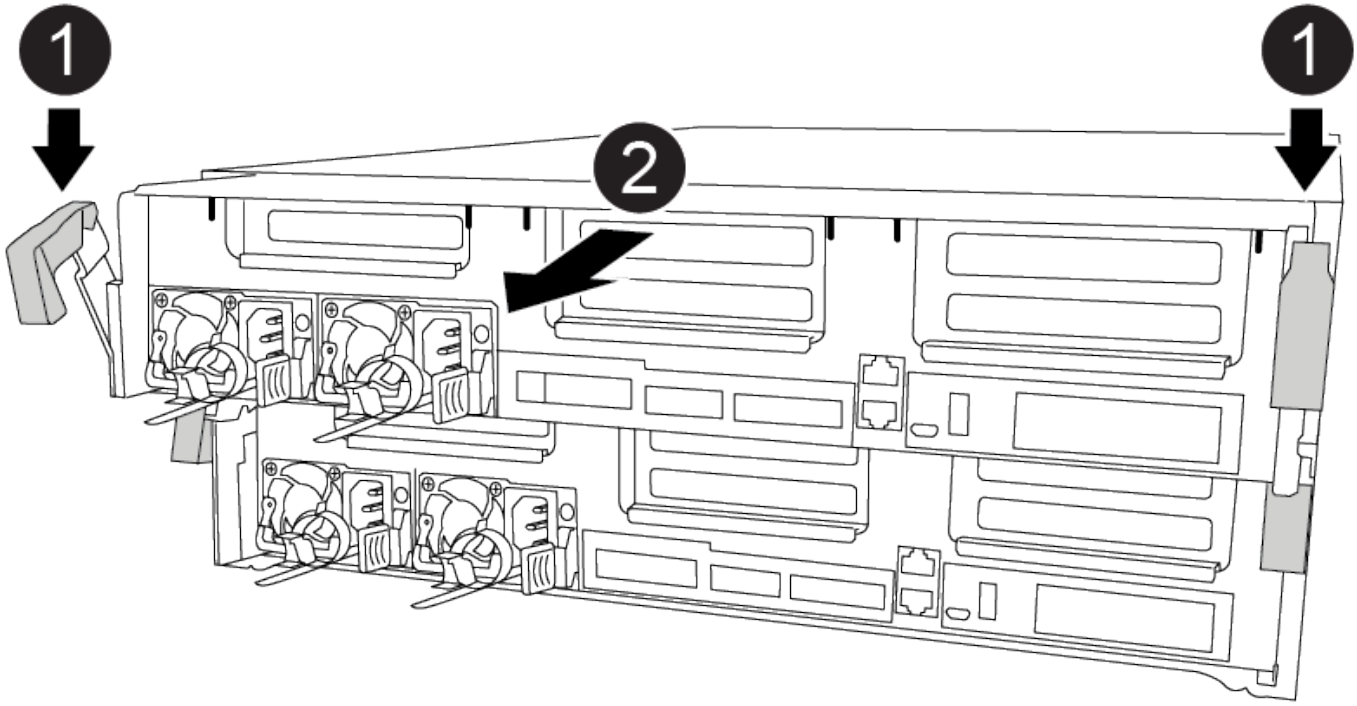
8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module



Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

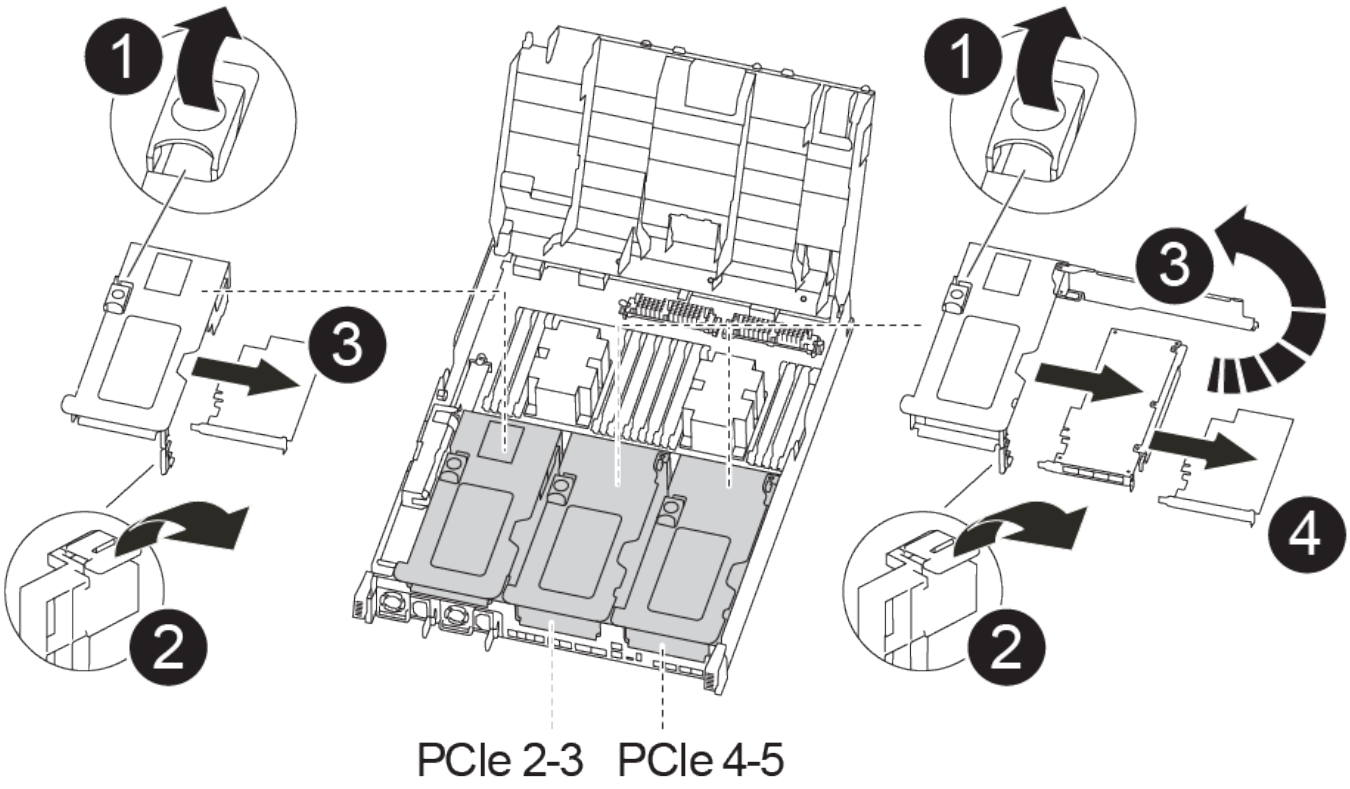
   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
7. Place the controller module on a stable, flat surface.

Step 3: Replace a PCIe card

To replace a PCIe card, you must locate the failed PCIe card, remove the riser that contains the card from the controller module, replace the card, and then reinstall the PCIe riser in the controller module.

You can use the following animation, illustration, or the written steps to replace a PCIe card.

Animation - Replace a PCIe card



Steps

1. Remove the riser containing the card to be replaced:
   a. Open the air duct by pressing the locking tabs on the sides of the air duct, slide it toward the back of the controller module, and then rotate it to its completely open position.
   b. Remove any SFP or QSFP modules that might be in the PCIe cards.
   c. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

   d. Lift the riser up straight up and set it aside on a stable flat surface,
2. Remove the PCIe card from the riser:
   a. Turn the riser so that you can access the PCIe card.
   b. Press the locking bracket on the side of the PCIe riser, and then rotate it to the open position.
   c. For risers 2 and 3 only, swing the side panel up.
   d. Remove the PCIe card from the riser by gently pushing up on the bracket and lift the card straight out of the socket.
3. Install the replacement PCIe card in the riser by aligning the card with the socket, press the card into the

socket and then close the side panel on the riser, if present.

Be sure that you properly align the card in the slot and exert even pressure on the card when seating it in the socket. The PCIe card must be fully and evenly seated in the slot.

> If you are installing a card in the bottom slot and cannot see the card socket well, remove the top card so that you can see the card socket, install the card, and then reinstall the card you removed from the top slot.

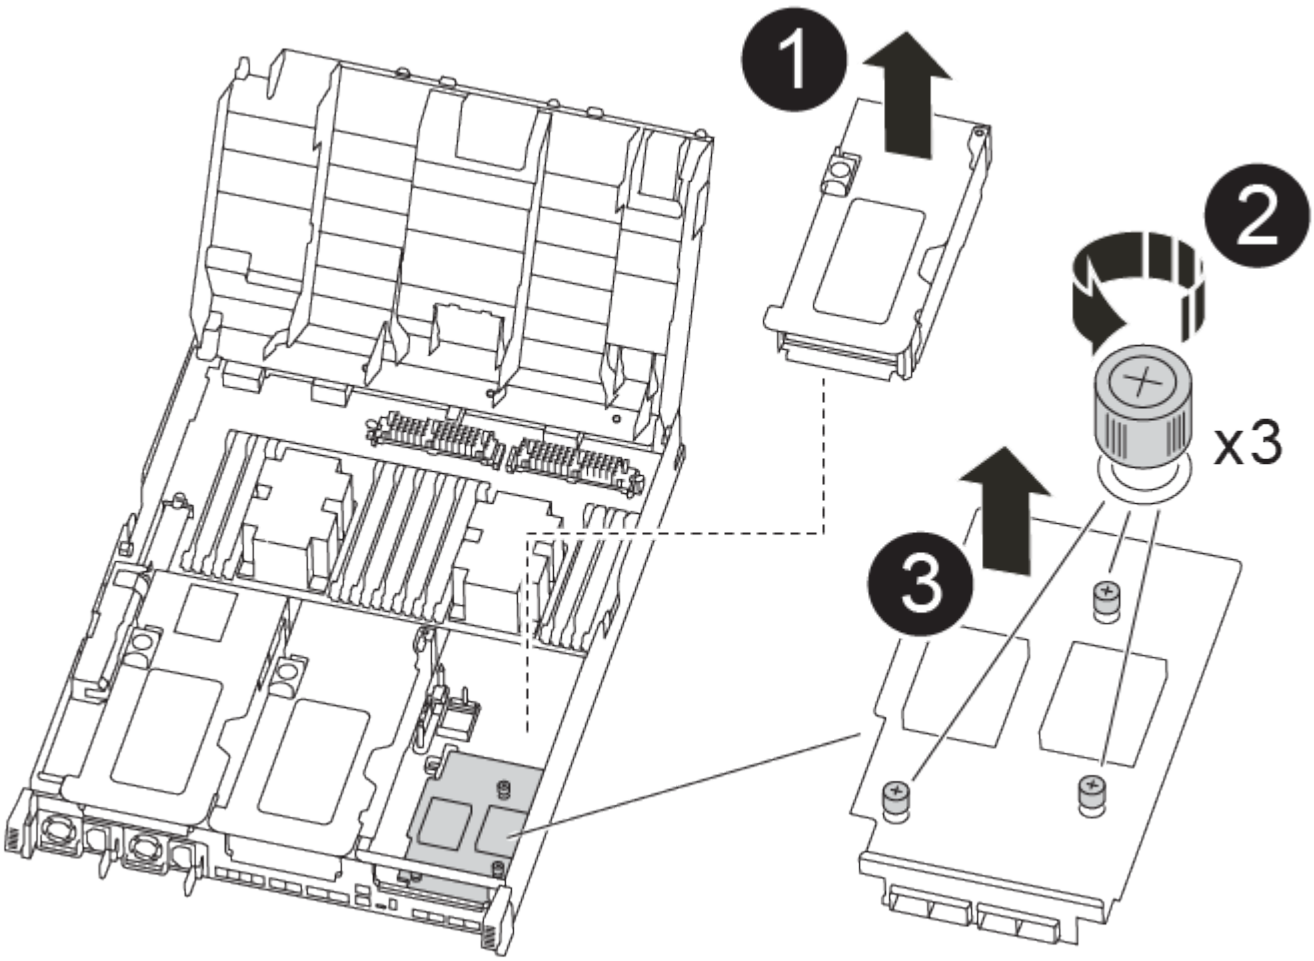
4. Reinstall the riser:
   a. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins.
   b. Push the riser squarely into the socket on the motherboard.
   c. Rotate the latch down flush with the sheet metal on the riser.

Step 4: Replace the mezzanine card

The mezzanine card is located under riser number 3 (slots 4 and 5). You must remove that riser to access the mezzanine card, replace the mezzanine card, and then reinstall riser number 3. See the FRU map on the controller module for more information.

You can use the following animation, illustration, or the written steps to replace the mezzanine card.

Animation - Replace the mezzanine card



Steps

1. Remove riser number 3 (slots 4 and 5):
   a. Open the air duct by pressing the locking tabs on the sides of the air duct, slide it toward the back of the controller module, and then rotate it to its completely open position.
   b. Remove any SFP or QSFP modules that might be in the PCIe cards.
   c. Rotate the riser locking latch on the left side of the riser up and toward air duct.

      The riser raises up slightly from the controller module.

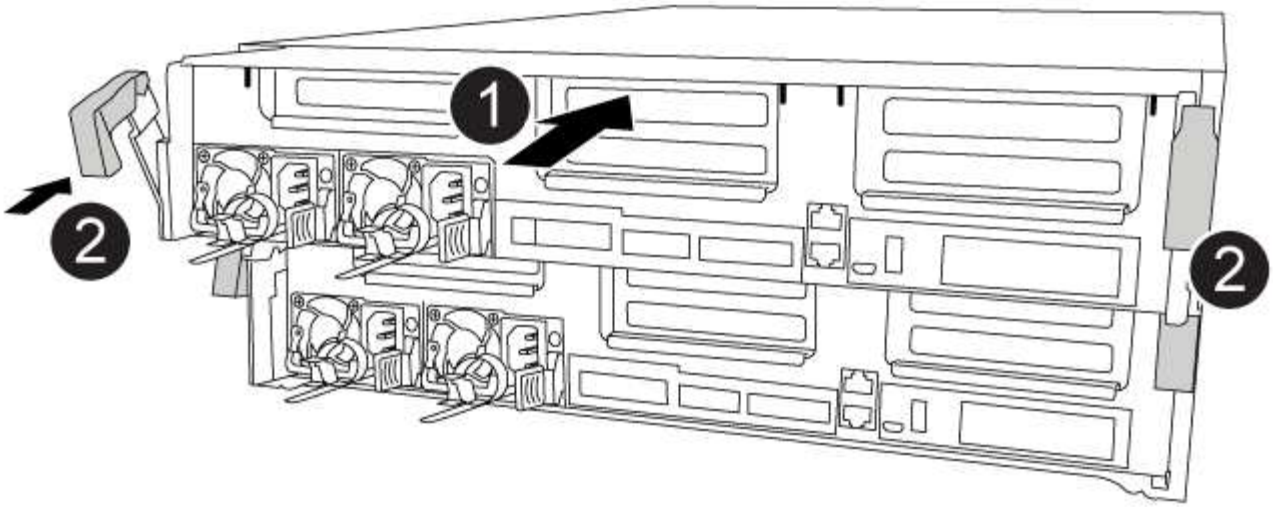
   d. Lift the riser up, and then set it aside on a stable, flat surface.
2. Replace the mezzanine card:
   a. Remove any QSFP or SFP modules from the card.
   b. Loosen the thumbscrews on the mezzanine card, and gently lift the card directly out of the socket and set it aside.
   c. Align the replacement mezzanine card over the socket and the guide pins and gently push the card into the socket.
   d. Tighten the thumbscrews on the mezzanine card.
3. Reinstall the riser:
   a. Align the riser with the pins to the side of the riser socket, lower the riser down on the pins.
   b. Push the riser squarely into the socket on the motherboard.
   c. Rotate the latch down flush with the sheet metal on the riser.

Step 5: Install the controller module

After you have replaced the component in the controller module, you must reinstall the controller module into the chassis, and then boot it to Maintenance mode.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module



Steps

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

If your system stops at the boot menu, select the option to boot to LOADER.

f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.

5. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
6. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 6: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the `enabled` state: `metrocluster node show`

```
cluster_B::>  metrocluster node show

DR                           Configuration  DR
Group Cluster Node           State          Mirroring Mode
----- ------- -------------- -------------- ---------
--------------------
1     cluster_A
              controller_A_1 configured     enabled   heal roots
completed
      cluster_B
              controller_B_1 configured     enabled   waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`
4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

   The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured        switchover
Remote: cluster_A configured        waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster              Configuration State    Mode
--------------------    -------------------     ---------
 Local: cluster_B configured            normal
Remote: cluster_A configured            normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 7: Restore the controller module to operation

You must recable the system, give back the controller module, and then reenable automatic giveback.

Steps

1. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

2. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 8: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Hot-swap a power supply - FAS8300 and FAS8700

Replacing a power supply (PSU) involves disconnecting the target PSU from the power source, unplugging the power cable, removing the old PSU and installing the replacement PSU, and then reconnecting the replacement PSU to the power source.

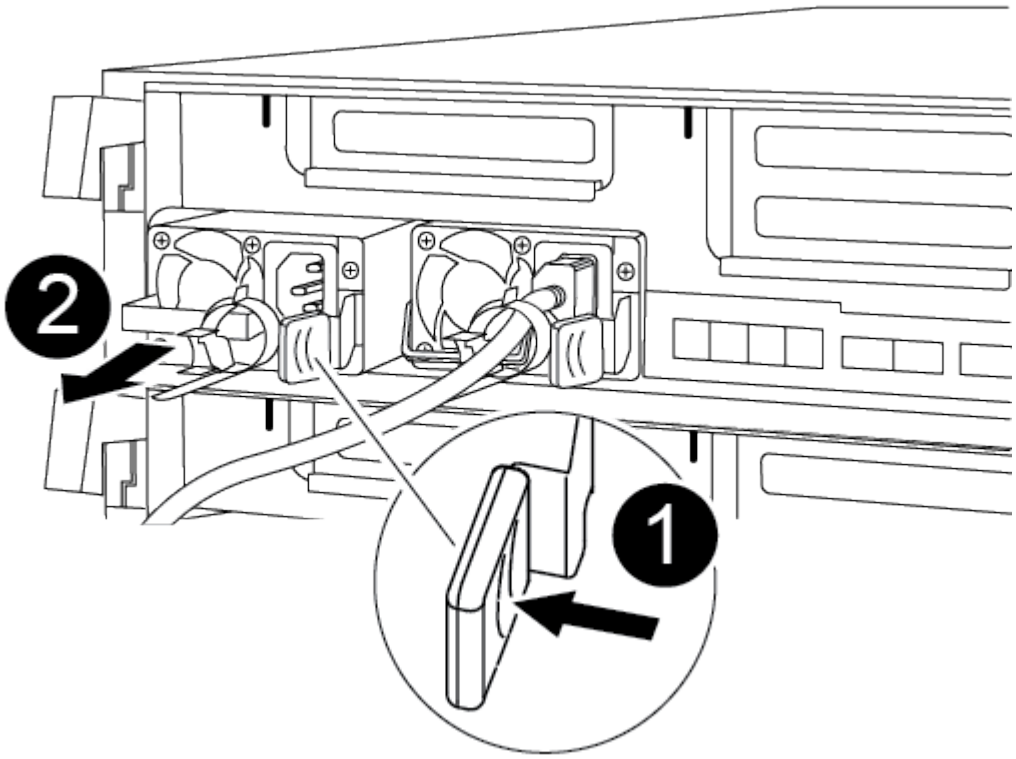
- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

You can use the following animation, illustration, or the written steps to replace the power supply.

Animation - Replace a power supply



Steps

1. If you are not already grounded, properly ground yourself.
2. Identify the power supply you want to replace, based on console error messages or through the LEDs on the power supplies.
3. Disconnect the power supply:
   a. Open the power cable retainer, and then unplug the power cable from the power supply.
   b. Unplug the power cable from the power source.
4. Remove the power supply:
   a. Rotate the cam handle so that it can be used to pull the power supply out of the chassis.
   b. Press the blue locking tab to release the power supply from the chassis.
   c. Using both hands, pull the power supply out of the chassis, and then set it aside.
5. Using both hands, support and align the edges of the power supply with the opening in the controller module, and then gently push the power supply into the controller module until the locking tab clicks into place.

   The power supplies will only properly engage with the internal connector and lock in place one way.

   > To avoid damaging the internal connector, do not use excessive force when sliding the power supply into the system.

6. Rotate the cam handle so that it is flush against the power supply.

7. Reconnect the power supply cabling:

   a. Reconnect the power cable to the power supply and the power source.

   b. Secure the power cable to the power supply using the power cable retainer.

   Once power is restored to the power supply, the status LED should be green.

8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the real-time clock battery - FAS8300 and FAS8700

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired controller

You can shut down or take over the impaired controller using different procedures, depending on the storage system hardware configuration.

Option 1: Most configurations

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:

   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* `-halt` *`true`* The *-halt true* parameter brings you to the LOADER prompt.

Option 2: Controller is in a two-node MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller…	Then…
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: `metrocluster switchover`
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
    Operation: heal-aggregates
      State: successful
Start Time: 7/25/2016 18:45:55
   End Time: 7/25/2016 18:45:56
     Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate     Size Available Used% State   #Vols  Nodes
RAID Status
--------- -------- --------- ----- ------- ------ ----------------
------------
...
aggr_b2    227.1GB   227.1GB    0% online       0 mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the -override-vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mcc1A::> metrocluster operation show
  Operation: heal-root-aggregates
      State: successful
 Start Time: 7/29/2016 20:54:41
   End Time: 7/29/2016 20:54:42
     Errors: -
```

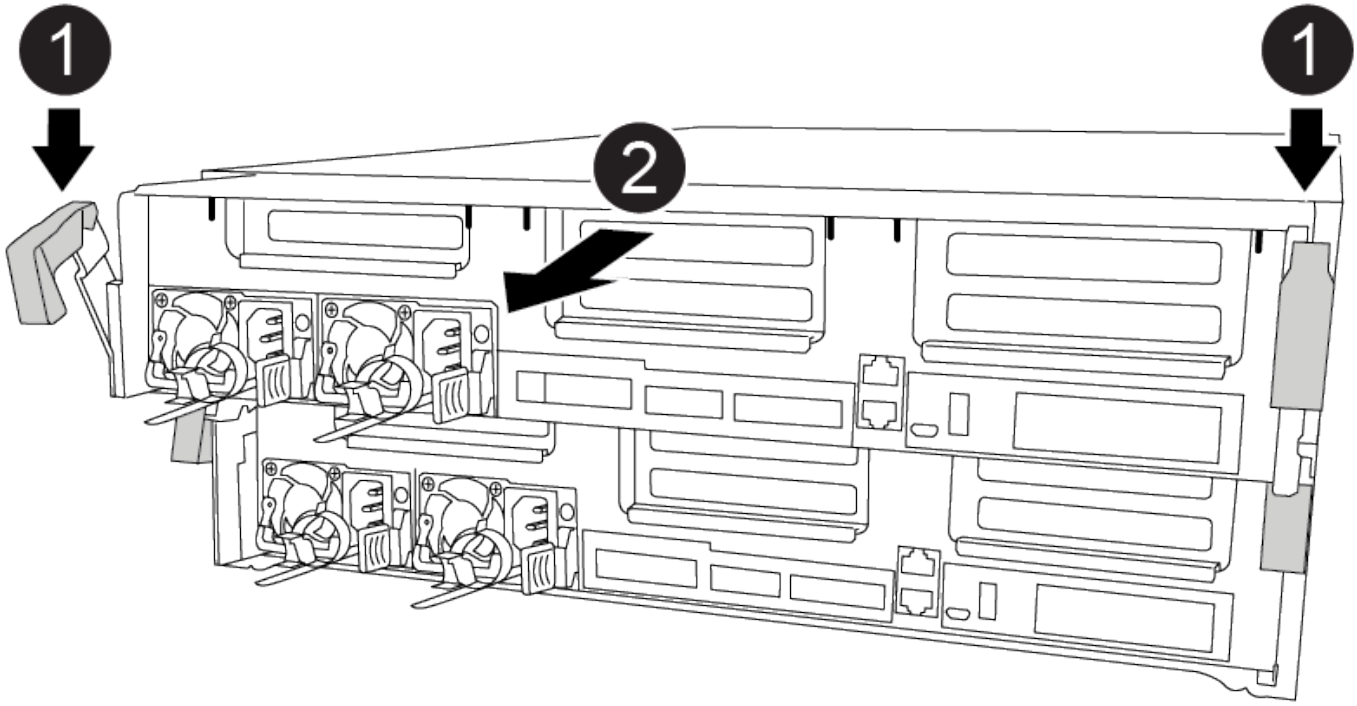
8. On the impaired controller module, disconnect the power supplies.

Step 2: Remove the controller module

To access components inside the controller module, you must remove the controller module from the chassis.

You can use the following animation, illustration, or the written steps to remove the controller module from the chassis.

Animation - Remove the controller module



Steps

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

   Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.
4. Remove the cable management device from the controller module and set it aside.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

   The controller module moves slightly out of the chassis.
6. Slide the controller module out of the chassis.

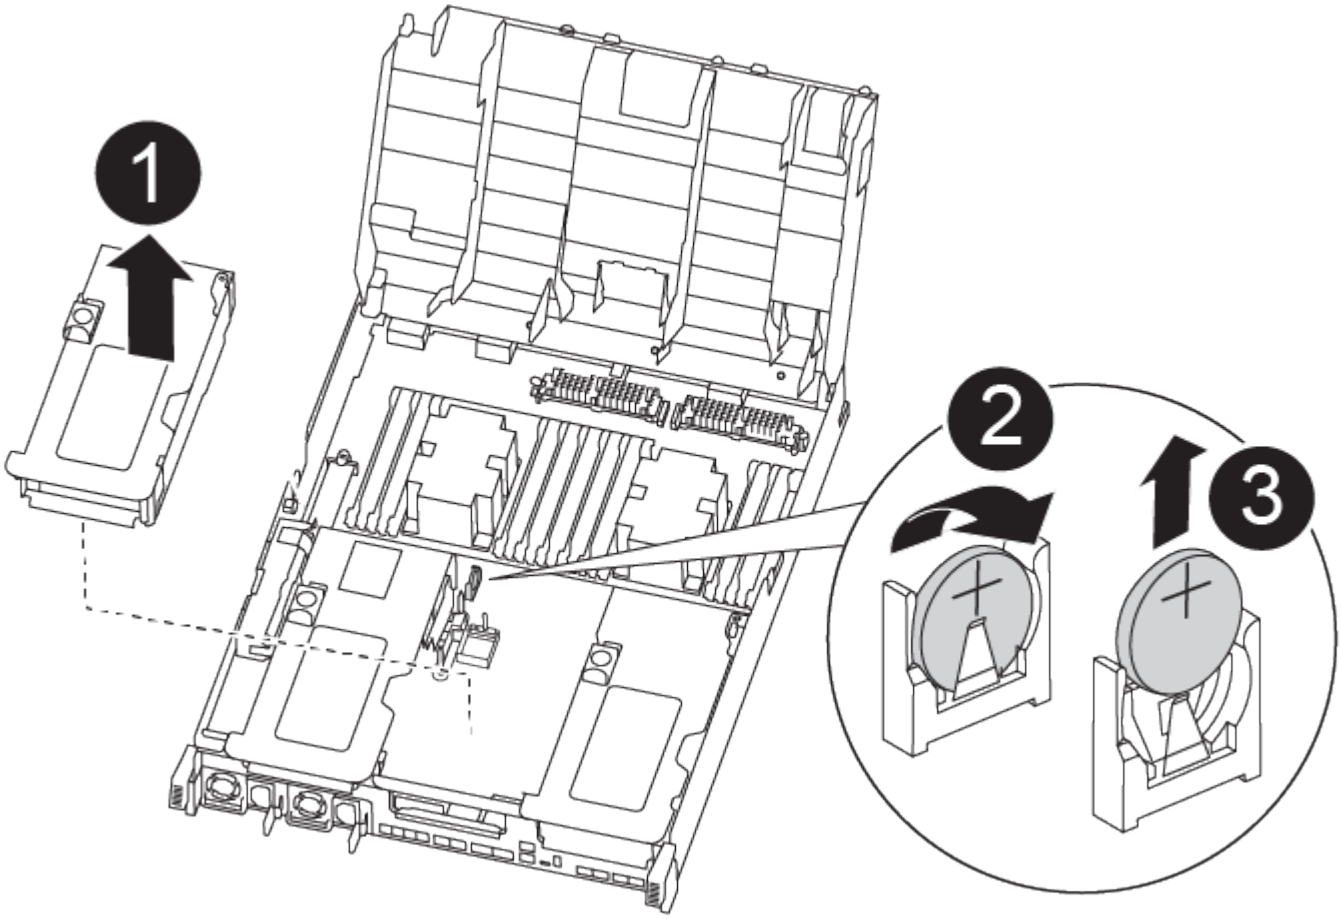
   Make sure that you support the bottom of the controller module as you slide it out of the chassis.
7. Place the controller module on a stable, flat surface.

Step 3: Replace the RTC battery

You need to locate the RTC battery inside the controller module, and then follow the specific sequence of steps. See the FRU map inside the controller module for the location of the RTC battery.

You can use the following animation, illustration, or the written steps to replace the RTC battery.

Animation - Replace the RTC battery



Steps

1. If you are not already grounded, properly ground yourself.
2. Open the air duct:
   a. Press the locking tabs on the sides of the air duct in toward the middle of the controller module.
   b. Slide the air duct toward the back of the controller module, and then rotate it upward to its completely open position.
3. Locate, remove, and then replace the RTC battery:
   a. Using the FRU map, locate the RTC battery on the controller module.
   b. Gently push the battery away from the holder, rotate it away from the holder, and then lift it out of the holder.

Note the polarity of the battery as you remove it from the holder. The battery is marked with a plus sign and must be positioned in the holder correctly. A plus sign near the holder tells you how the battery should be positioned.

c. Remove the replacement battery from the antistatic shipping bag.

d. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.

4. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.

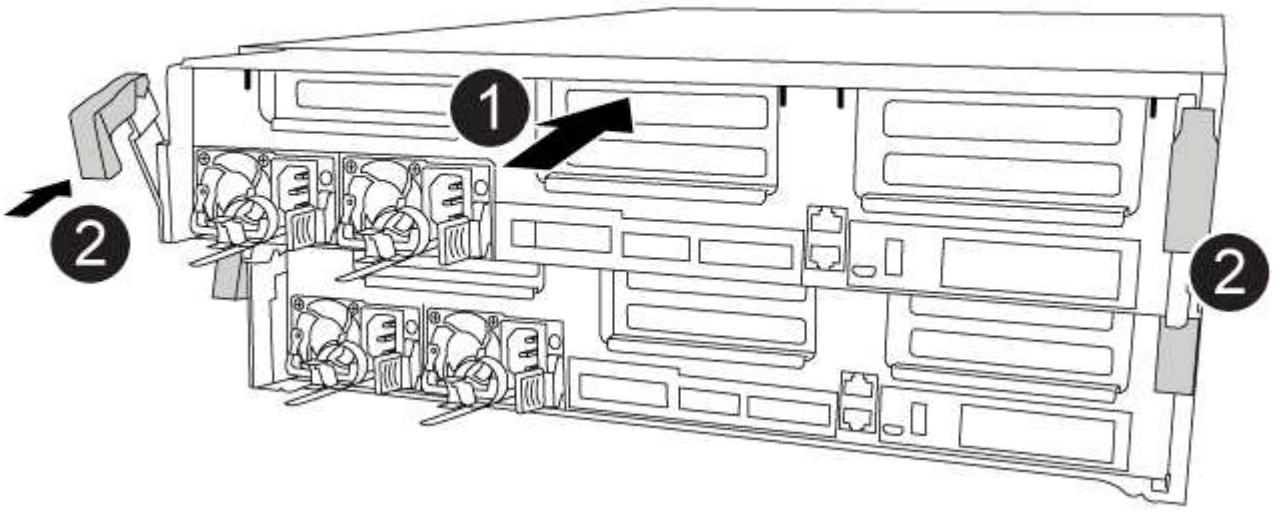
5. Close the air duct.

Step 4: Reinstall the controller module and sett time/date after RTC battery replacement

After you replace a component within the controller module, you must reinstall the controller module in the system chassis, reset the time and date on the controller, and then boot it.

You can use the following animation, illustration, or the written steps to install the controller module in the chassis.

Animation - Install the controller module



Steps

1. If you have not already done so, close the air duct or controller module cover.

2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

   Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. Complete the installation of the controller module:
   a. Using the locking latches, firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      > Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

   b. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.
   c. Plug the power cords into the power supplies, reinstall the power cable locking collar, and then connect the power supplies to the power source.

      The controller module begins to boot as soon as power is restored. Be prepared to interrupt the boot process.

   d. If you have not already done so, reinstall the cable management device.
   e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.

      > If your system stops at the boot menu, select the option to boot to LOADER.

5. 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.
6. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the controller reboot.
7. Return the controller to normal operation by giving back its storage: `storage failover giveback -ofnode` *`impaired_node_name`*
8. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Switch back aggregates in a two-node MetroCluster configuration

This task only applies to two-node MetroCluster configurations.

Steps

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

Key specifications for FAS8300

The following are select specifications for the FAS8300 storage system in a single high availability pair. Visit NetApp Hardware Universe (HWU) for the complete specifications for this storage system.

FAS8300 specifications at a glance

- Platform Configuration: FAS8300 Single Chassis HA Pair, Ethernet Bundle
- Max Raw Capacity: 15.0000 PB
- Memory: 256.0000 GB
- Form Factor: 4U chassis with 2 HA controllers
- ONTAP Version: ONTAP: 9.16.1P2
- PCIe Expansion Slots: 14
- Minimum ONTAP Version: ONTAP 9.7RC1

Scaleout maximums

- Type: NAS; HA Pairs: 12; Raw Capacity: 180.0 PB / 159.9 PiB; Max Memory: 3072 GB
- Type: SAN; HA Pairs: 6; Raw Capacity: 90.0 PB / 79.9 PiB; Max Memory: 1536 GB
- Type: HA Pair; Raw Capacity: 15.0 PB / 13.3 PiB; Max Memory: 256.0000

I/O

Onboard I/O

- Protocol: Ethernet 100 Gbps; Ports: 4
- Protocol: Ethernet 25 Gbps; Ports: 12
- Protocol: SAS 12 Gbps; Ports: 8

Total I/O

- Protocol: Ethernet 100 Gbps; Ports: 20
- Protocol: Ethernet 25 Gbps; Ports: 28
- Protocol: Ethernet 10 Gbps; Ports: 32
- Protocol: FC 32 Gbps; Ports: 32
- Protocol: NVMe/FC 32 Gbps; Ports: 32
- Ports: 0
- Protocol: SAS 12 Gbps; Ports: 40

Management ports

- Protocol: Ethernet 1 Gbps; Ports: 2
- Protocol: RS-232 115 Kbps; Ports: 4
- Protocol: USB 12 Mbps; Ports: 4

Storage networking supported

- CIFS
- FC
- iSCSI
- NFS v3
- NFS v4.0
- NFS v4.1
- NFS v4.2
- NVMe/FC
- 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: 4024 BTU/hr
- Worst-case Power: 5017 BTU/hr
- Weight: 108.5 lb
  49.2 kg
- Height: 4U
- Width: 19” IEC rack-compliant (17.6” 44.7 cm)
- Depth: 32.6”
  (34.7” with cable management bracket)
- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°F) at up to 3048m (10000 ft) elevation; 8% to 80% 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): 8.5
  Sound pressure (LpAm) (bystander positions): 67.2 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 for external shelves
[//] 2025-10-15 ontap-systems-internal/issues/1357

Key specifications for FAS8700

The following are select specifications for the FAS8700 storage system in a single high availability pair. Visit NetApp Hardware Universe (HWU) for the complete specifications for this storage system.

FAS8700 specifications at a glance

- Platform Configuration: FAS8700 Single Chassis HA Pair, Ethernet Bundle
- Max Raw Capacity: 15.0000 PB
- Memory: 512.0000 GB
- Form Factor: 4U chassis with 2 HA controllers
- ONTAP Version: ONTAP: 9.16.1P2
- PCIe Expansion Slots: 14
- Minimum ONTAP Version: ONTAP 9.7RC1

Scaleout maximums

- Type: NAS; HA Pairs: 12; Raw Capacity: 180.0 PB / 159.9 PiB; Max Memory: 6144 GB
- Type: SAN; HA Pairs: 6; Raw Capacity: 90.0 PB / 79.9 PiB; Max Memory: 3072 GB
- Type: HA Pair; Raw Capacity: 15.0 PB / 13.3 PiB; Max Memory: 512.0000

I/O

Onboard I/O

- Protocol: Ethernet 100 Gbps; Ports: 4
- Protocol: Ethernet 25 Gbps; Ports: 12
- Protocol: SAS 12 Gbps; Ports: 8

Total I/O

- Protocol: Ethernet 100 Gbps; Ports: 20
- Protocol: Ethernet 25 Gbps; Ports: 28
- Protocol: Ethernet 10 Gbps; Ports: 32
- Protocol: FC 32 Gbps; Ports: 32
- Protocol: NVMe/FC 32 Gbps; Ports: 32
- Ports: 0
- Protocol: SAS 12 Gbps; Ports: 40

Management ports

- Protocol: Ethernet 1 Gbps; Ports: 2
- Protocol: RS-232 115 Kbps; Ports: 4
- Protocol: USB 12 Mbps; Ports: 4

Storage networking supported

- CIFS
- FC
- iSCSI
- NFS v3
- NFS v4.0
- NFS v4.1
- NFS v4.2
- NVMe/FC
- 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: 5430 BTU/hr
- Worst-case Power: 5796 BTU/hr
- Weight: 108.9 lb
  49.4 kg
- Height: 4U
- Width: 19” IEC rack-compliant (17.6” 44.7 cm)
- Depth: 32.6”
  (34.7” with cable management bracket)
- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°F) at up to 3048m (10000 ft) elevation; 8% to 80% 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): 8.5
  Sound pressure (LpAm) (bystander positions): 67.2 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 for external shelves
[//] 2025-10-15 ontap-systems-internal/issues/1357

FAS9500 systems

Install and setup

Start here: Choose your installation and setup experience

You can choose from different content formats to guide you through installing and setting up your new storage system.

- 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.

Quick steps - FAS9500

This topic gives graphic instructions for a typical installation of your system from racking and cabling, through initial system bring-up. Use this content if you are familiar with installing NetApp systems.

Access the *Installation and Setup Instructions* PDF poster:

FAS9500 Installation and Setup Instructions

Video steps - FAS9500

The following video shows how to install and cable your new system.

Animation - FAS9500 install and setup instructions (ISI)

Detailed steps - FAS9500

This article gives detailed step-by-step instructions for installing a typical NetApp system. Use this article if you want more detailed installation instructions.

Step 1: Prepare for installation

To install your system, you need to create an account on the NetApp Support Site, register your system, and get license keys. You also need to inventory the appropriate number and type of cables for your system and collect specific network information.

You need to have access to the NetApp Hardware Universe for information about site requirements as well as additional information on your configured system.

What you need

You might also want to have access to the ONTAP 9 Release Notes for your version of ONTAP for more information about this system.

You need to provide the following 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

Steps

1. Unpack the contents of all boxes.
2. Record the system serial number from the controllers.



3. 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 the Hardware Universe to locate the cable and identify its use.

   NetApp Hardware Universe

Type of cable…	Part number and length	Connector type	For…
25 GbE data Cable	X66240A-05 (112-00639), 0.5m X66240A-2 (112-00598), 2m X66240A-5 (112-00600), 5m		Network cable
32 Gb FC (SFP+ Op)	X66250-2 (112-00342), 2m X66250-5 (112-00344), 5m X66250-15 (112-00346), 15m		FC optical network cable
40 GbE network cable	X66100-1 (112-00542), 1m X66100-3 (112-00543), 3m X66100-5 (112-00544), 5m		Ethernet data, cluster network

Type of cable…	Part number and length	Connector type	For…
100 GbE cable	X66211B-1 (112-00573), 1m X66211B-2 (112-00574), 2m X66211B-5 (112-00576), 5m		Network, Ethernet data, cluster network
Optical cables	X66031A (112-00436), 1m X66032A (112-00437), 2m X66033A (112-00438), 3m		FC optical 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	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

4. Review the ONTAP Configuration Guide and collect the required information listed in that guide.

Step 2: Install the hardware

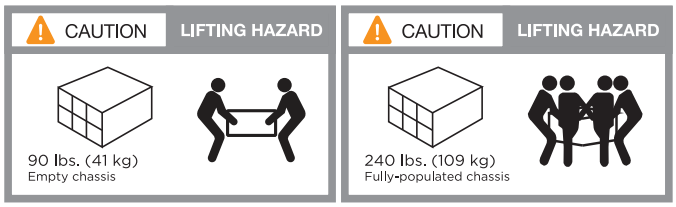
You need to install your system in a 4-post rack or NetApp system cabinet, as applicable.

Steps

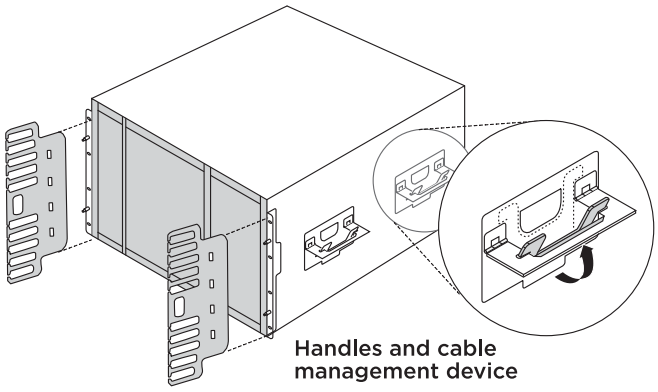
1. Install the rail kits, as needed.
2. Install and secure your system using the instructions included with the rail kit.

   You need to be aware of the safety concerns associated with the weight of the system.

   The label on the left indicates an empty chassis, while the label on the right indicates a fully-populated system.



3. Attach cable management devices (as shown).



4. Place the bezel on the front of the system.

Step 3: Cable controllers to your network

You can cable the controllers to your network by using the two-node switchless cluster method or by using the cluster interconnect network.

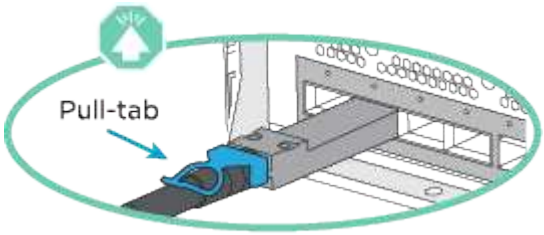
Option 1: Two-node switchless cluster

Management network, data network, and management ports on the controllers are connected to switches. The cluster interconnect ports are cabled on both controllers.

Before you begin

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all networking module ports.

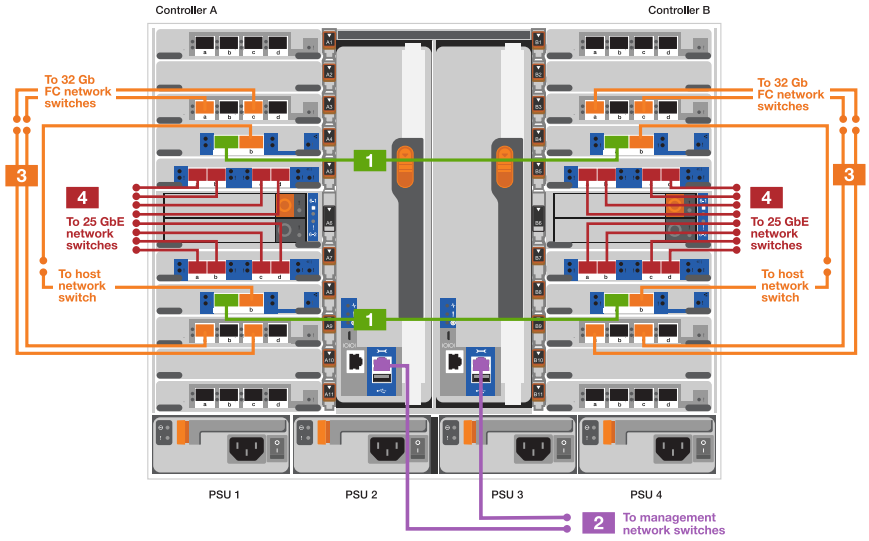


As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it over and try again.

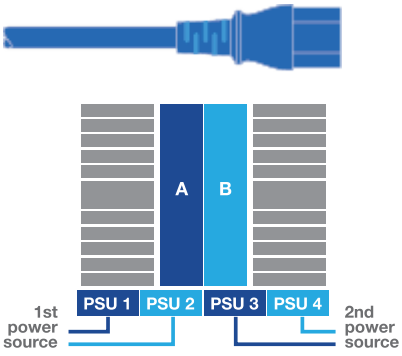
Steps

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

Animation Two-node switchless cluster cabling



Step	Perform on each controller
1	Cable cluster interconnect ports: • Slot A4 and B4 (e4a) • Slot A8 and B8 (e8a) 
2	Cable controller management (wrench) ports. 
3	Cable 32 Gb FC network switches: Ports in slot A3 and B3 (e3a and e3c) and slot A9 and B9 (e9a and e9c) to the 32 Gb FC network switches.  40GbE host network switches: Cable host-side b ports in slot A4 and B4 (e4b) and slot A8 and B8 (e8b) to the host switch. 
4	Cable 25 GbE connections: Cable ports in slot A5 and B5 (5a, 5b, 5c, and 5d) and slot A7 and B7 (7a, 7b, 7c, and 7d) to the 25 GbE network switches. 

Step	Perform on each controller
5	• Strap the cables to the cable management arms (not shown). • Connect the power cables to the PSUs and connect them to different power sources (not shown). PSU 1 and 3 provide power to all side A components, while PSU2 and PSU4 provide power to all side B components. 

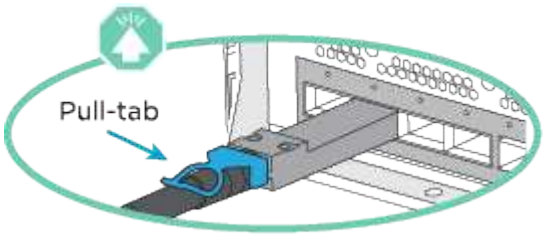
Option 2: Switched cluster

Management network, data network, and management ports on the controllers are connected to switches. The cluster interconnect and HA ports are cabled on to the cluster/HA switch.

Before you begin

You must have contacted your network administrator for information about connecting the system to the switches.

Be sure to check the direction of the cable pull-tabs when inserting the cables in the ports. Cable pull-tabs are up for all networking module ports.

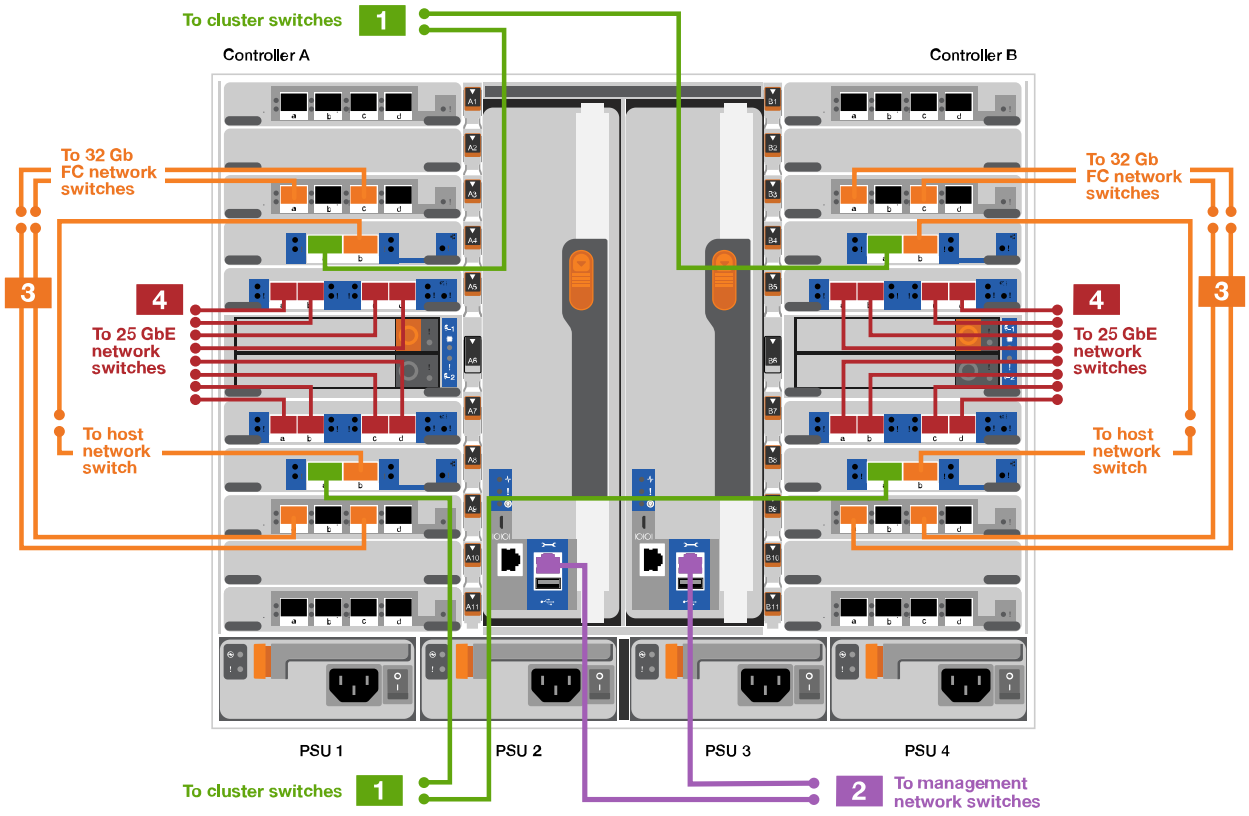


As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it over and try again.

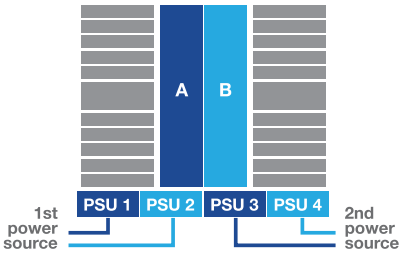
Steps

1. Use the animation or illustration to complete the cabling between the controllers and to the switches:

   Animation Switched cluster cabling



Step	Perform on each controller
1	Cable cluster interconnect a ports: • Slot A4 and B4 (e4a) to the cluster network switch. • Slot A8 and B8 (e8a) to the cluster network switch. 
2	Cable controller management (wrench) ports. 

Step	Perform on each controller
3	Cable 32 Gb FC network switches: Ports in slot A3 and B3 (e3a and e3c) and slot A9 and B9 (e9a and e9c) to the 32 Gb FC network switches.  40GbE host network switches: Cable host-side b ports in slot A4 and B4 (e4b) and slot A8 and B8 (e8b) to the host switch. 
4	Cable 25 GbE connections: Cable ports in slot A5 and B5 (5a, 5b, 5c, and 5d) and slot A7 and B7 (7a, 7b, 7c, and 7d) to the 25 GbE network switches. 
4	• Strap the cables to the cable management arms (not shown). • Connect the power cables to the PSUs and connect them to different power sources (not shown). PSU 1 and 3 provide power to all side A components, while PSU2 and PSU4 provide power to all side B components.  

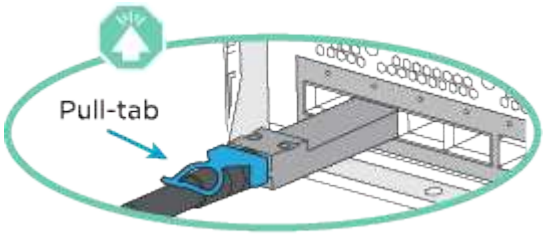
Step 4: Cable controllers to drive shelves

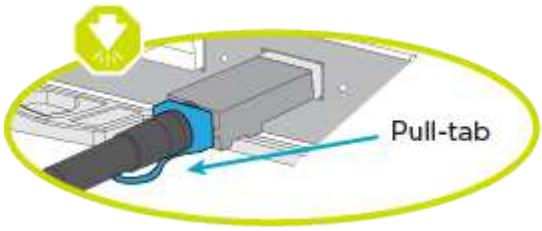
Cable either DS212C or DS224C drive shelves to your controllers.

For more SAS cabling information and worksheets, see SAS cabling rules, worksheets, and examples overview - shelves with IOM12 modules

Before you begin

- Complete the SAS cabling worksheet for your system. See SAS cabling rules, worksheets, and examples overview - shelves with IOM12 modules.
- Be sure to check the illustration arrow for the proper cable connector pull-tab orientation. The cable pull-tab for the storage modules are up, while the pull tabs on the shelves are down.



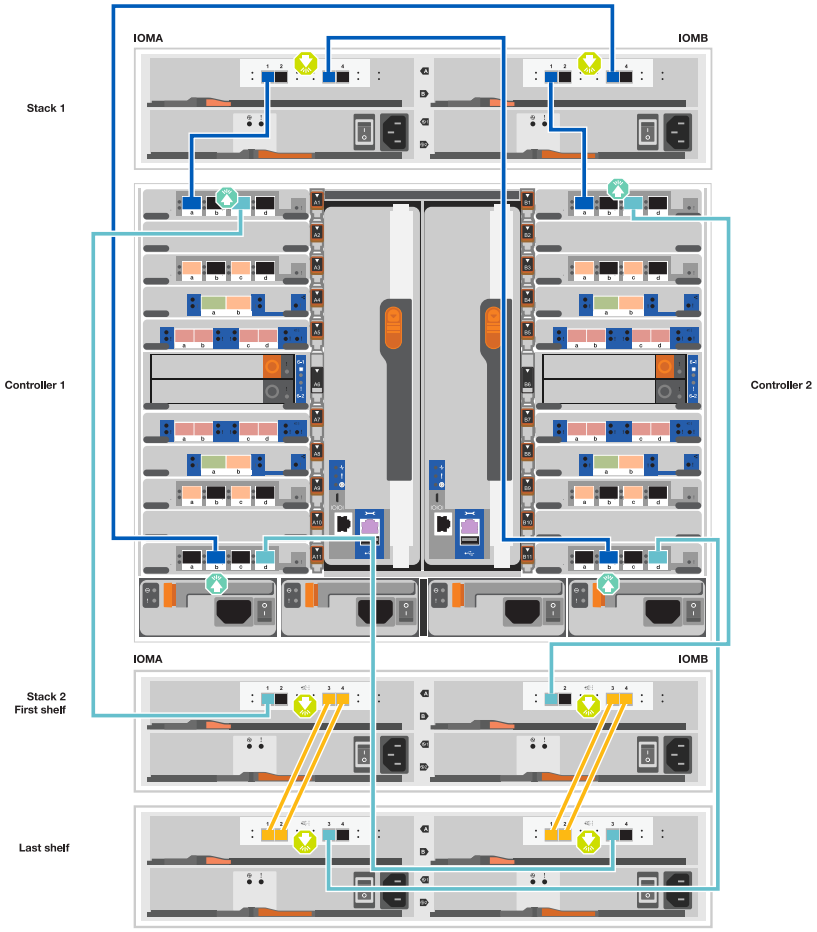


As you insert the connector, you should feel it click into place; if you do not feel it click, remove it, turn it over and try again.

Steps

1. Use the following animation or drawings to cable your controllers to three (1 stack of one drive shelf and one stack of two drive shelves) DS224C drive shelves.

   Animation Cable your drive shelves



Step	Perform on each controller
1	Connect drive shelf stack one to the controllers, using the graphic for reference. **Mini-SAS cable** 
2	Connect the drive shelves within stack two to each other, using the graphic for reference. **Mini-SAS cable** 
3	Connect drive shelf stack two to the controllers, using the graphic for reference. **Mini-SAS cable** 

Step 5: Complete system setup and configuration

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.

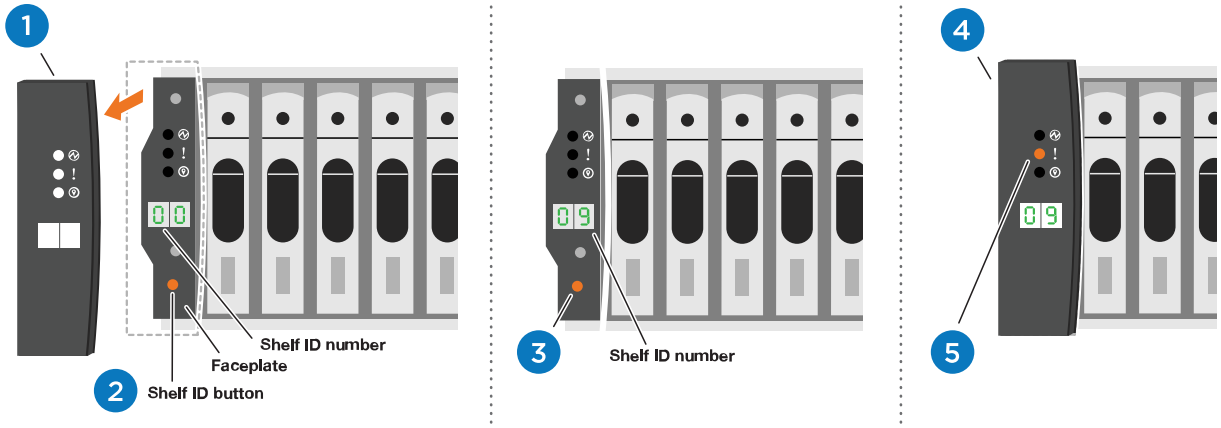
Option 1: If network discovery is enabled

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 or drawing to set one or more drive shelf IDs:

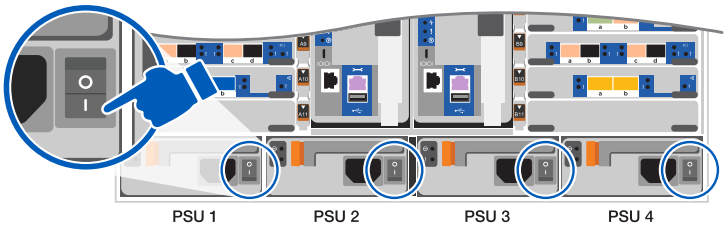
   Animation Set your shelf ID's



1	Remove the end cap.
2	Press and hold shelf ID button until first digit blinks, then push to advance to 0-9. The first digit continues to blink
2	Press and hold shelf ID button until second digit blinks, then push to advance to 0-9. The first digit stops blinking, and the second digit continues to blink.
4	Replace the end cap.
5	Wait 10 seconds for the Amber LED (!) to appear, then power-cycle the drive shelf to set shelf ID.

2. Turn on the power switches on the power supplies to both nodes.

   Animation Turn on the power to the controllers



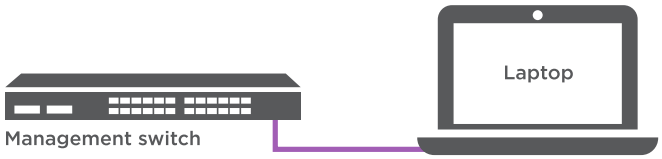
Initial booting may take up to eight minutes.

3. Make sure that your laptop has network discovery enabled.

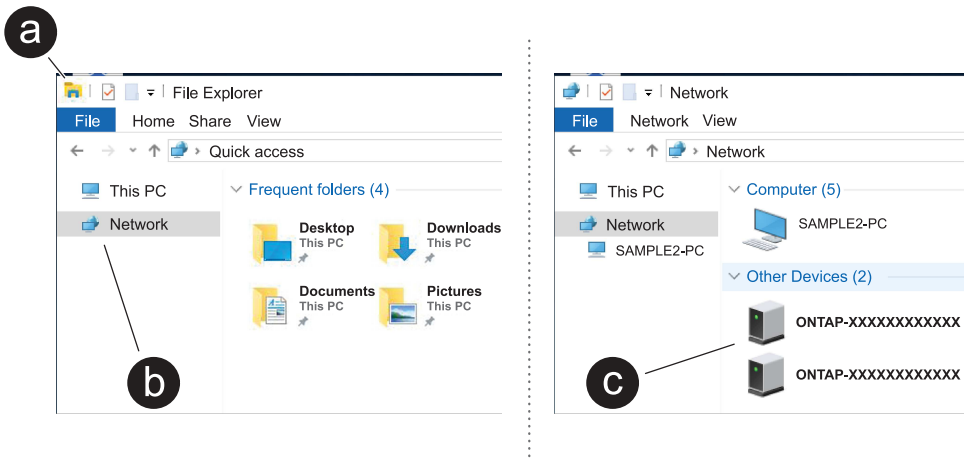
   See your laptop's online help for more information.

4. Use the following animation to connect your laptop to the Management switch.

   Animation Connect your laptop to the Management switch



5. 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.

6. Use System Manager guided setup to configure your 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 an account.

   NetApp Support Registration

b. Register your system.

   NetApp Product Registration

c. Download Active IQ Config Advisor.

   NetApp Downloads: Config Advisor

8. Verify the health of your system by running Config Advisor.
9. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Option 2: If network discovery is not enabled

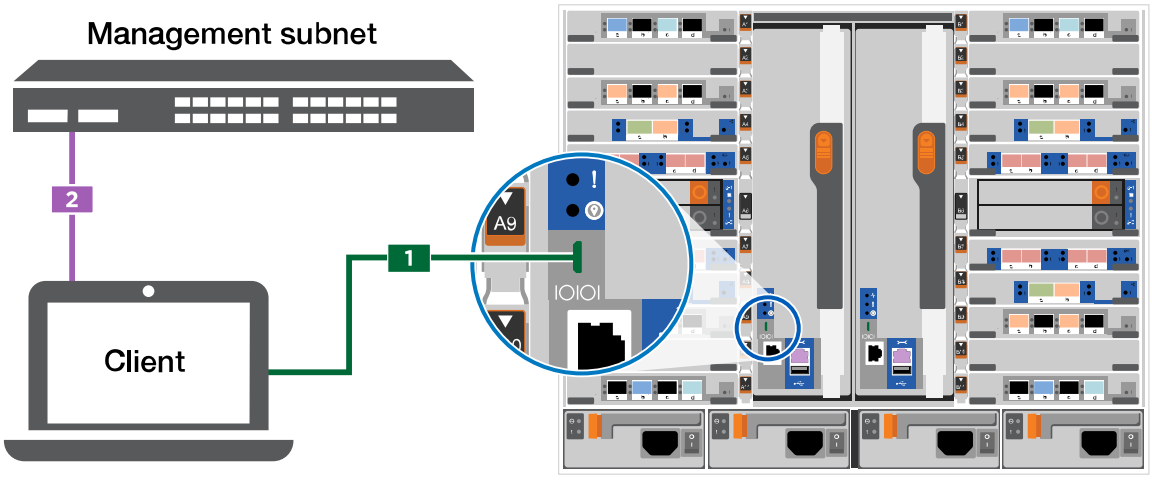
If you are not using a Windows or Mac-based laptop or console or if auto discovery is not enabled, you must complete the configuration and setup using this task.

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.

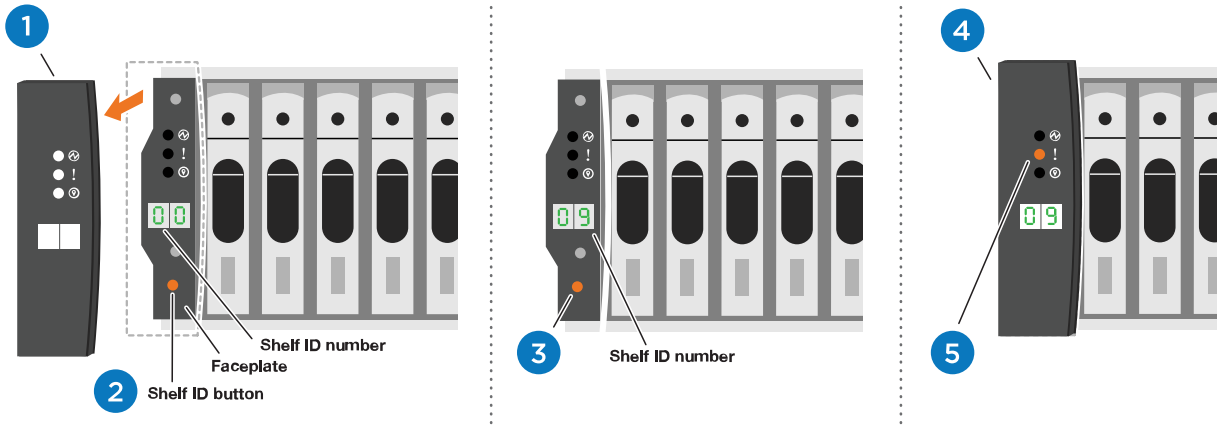
      See your laptop or console's online help for how to configure the console port.

   b. Connect the console cable to the laptop or console using the console cable that came with your system, and then connect the laptop to the switch on the management subnet.



   c. 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:

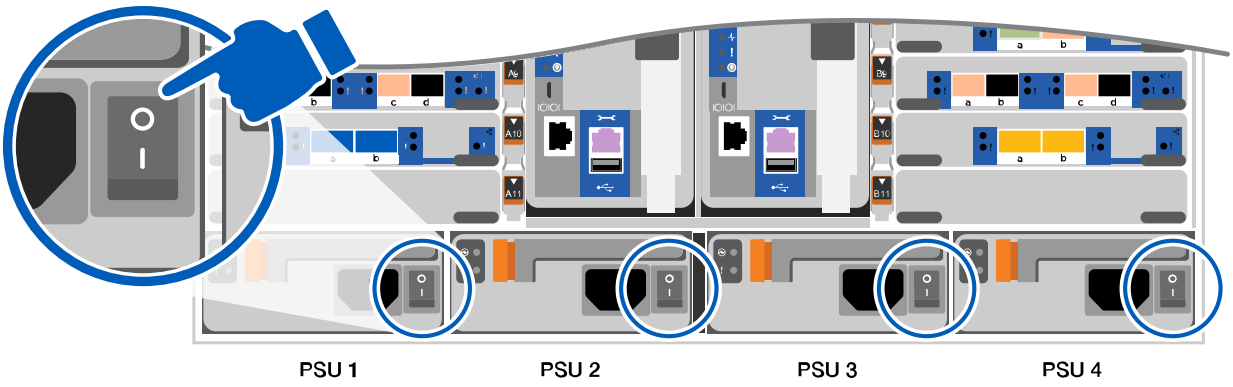
   Animation Set your shelf ID's



1	Remove the end cap.
2	Press and hold shelf ID button until first digit blinks, then push to advance to 0-9. The first digit continues to blink
2	Press and hold shelf ID button until second digit blinks, then push to advance to 0-9. The first digit stops blinking, and the second digit continues to blink.
4	Replace the end cap.
5	Wait 10 seconds for the Amber LED (!) to appear, then power-cycle the drive shelf to set shelf ID.

3. Turn on the power switches on the power supplies to both nodes.

Animation Turn on the power to the controllers



Initial booting may take up to eight minutes.

4. 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.

5. 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 .
6. Set up your account and download Active IQ Config Advisor:
   a. Log in to your existing account or create an account.
   NetApp Support Registration
   b. Register your system.
   NetApp Product Registration
   c. Download Active IQ Config Advisor.
   NetApp Downloads: Config Advisor
7. Verify the health of your system by running Config Advisor.
8. After you have completed the initial configuration, go to ONTAP 9 documentation for information about configuring additional features in ONTAP.

Maintain

Maintain FAS9500 hardware

Maintain the hardware of your FAS9500 storage system to ensure long-term reliability and optimal performance. Perform regular maintenance tasks such as replacing faulty

components, as this helps prevent downtime and data loss.

The maintenance procedures assume that the FAS9500 storage system has already been deployed as a storage node in the ONTAP environment.

System components

For the FAS9500 storage system, you can perform maintenance procedures on the following components.

Boot media - automated recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During automated recovery, the system retrieves the boot image from the partner node and automatically runs the appropriate boot menu option to install the image on your replacement boot media. The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.
Boot media - manual recovery	The boot media stores a primary and secondary set of ONTAP image files that the storage system uses to boot. During manual recovery, you boot the storage system from a USB drive and manually restore the file system image and configuration. If your storage system is running ONTAP 9.17.1 and later, use the automated boot recovery procedure.
Caching module	You must replace the controller's caching module when your system registers a single AutoSupport (ASUP) message that the module has gone offline.
Chassis	The chassis is the physical enclosure housing all the controller components such as the controller/CPU unit, power supply, and I/O.
Controller	A controller consists of a board, firmware, and software. It controls the drives and implements the ONTAP functions.
DIMM	You must replace a DIMM (dual in-line memory module) when a memory mismatch is present, or you have a failed DIMM.
DCPM	The DCPM (destage controller power module) contains the NVRAM11 battery.
Fan	The fan cools the controller.
I/O module	The I/O module (Input/Output module) is a hardware component that acts as an intermediary between the controller and various devices or systems that need to exchange data with the controller.
LED USB	The LED USB module provides connectivity to console ports and system status.
NVRAM	The NVRAM module (Non-Volatile Random Access Memory) allows the controller to retain data across power cycles or system reboots, while the NVRAM DIMM maintains NVRAM settings.

Power supply	A power supply provides a redundant power source in a controller.
Real-time clock battery	A real time clock battery preserves system date and time information if the power is off.

Boot media - automated recovery

Boot media automated recovery workflow - FAS9500

The automated recovery of the boot image involves the system automatically identifying and selecting the appropriate boot menu option. It uses the boot image on partner node to reinstall ONTAP on the replacement boot media in your FAS9500 storage system.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To get started, review the replacement requirements, shut down the controller, replace the boot media, allow the system to restore the image, and verify system functionality.

1. **Review the boot media requirements**

Review the requirements for boot media replacement.

2. **Shut down the controller**

Shut down the controller in your storage system when when you need to replace the boot media.

3. **Replace the boot media**

Remove the failed boot media from the controller module and install the replacement boot media.

4. **Restore the image on the boot media**

Restore the ONTAP image from the partner controller.

5. **Return the failed part to NetApp**

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for automated boot media recovery - FAS9500

Before replacing the boot media in your FAS9500, ensure you meet the necessary requirements for a successful replacement. This includes verifying that you have the correct replacement boot media, confirming that the e0S (e0M wrench) port on the impaired controller is not faulty, and determining whether Onboard Key Manager (OKM) or External Key Manager (EKM) is enabled.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

- You must replace the failed component with a replacement FRU component of the same capacity that you received from NetApp.
- Verify that the e0M (wrench) port on the impaired controller is connected and not faulty.

  The e0M port is used to communicate between the two controllers during the automated boot recovery process.

- For OKM, you need the cluster-wide passphrase and also the backup data.
- For EKM, you need copies of the following files from the partner node:
  - /cfcard/kmip/servers.cfg file.
  - /cfcard/kmip/certs/client.crt file.
  - /cfcard/kmip/certs/client.key file.
  - /cfcard/kmip/certs/CA.pem file.
- It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:
  - The *impaired controller* is the controller on which you are performing maintenance.
  - The *healthy controller* is the HA partner of the impaired controller.

What's next

After you've reviewed the boot media requirements, you shut down the controller.

Shut down the controller for automated boot media recovery - FAS9500

Shut down the impaired controller in your FAS9500 storage system to prevent data loss and maintain system stability during the automated boot media recovery process.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   `cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node impaired_node_name -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

What's next

After you shut down the impaired controller, you replace the boot media.

Replace the boot media for automated boot recovery - FAS9500

The boot media in your FAS9500 system stores essential firmware and configuration data. The replacement process involves removing and opening the controller module, removing the impaired boot media, installing the replacement boot media in the controller module, and then reinstalling the controller module.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

The boot media is located inside the controller module under the air duct, and is accessed by removing the controller module from the system.

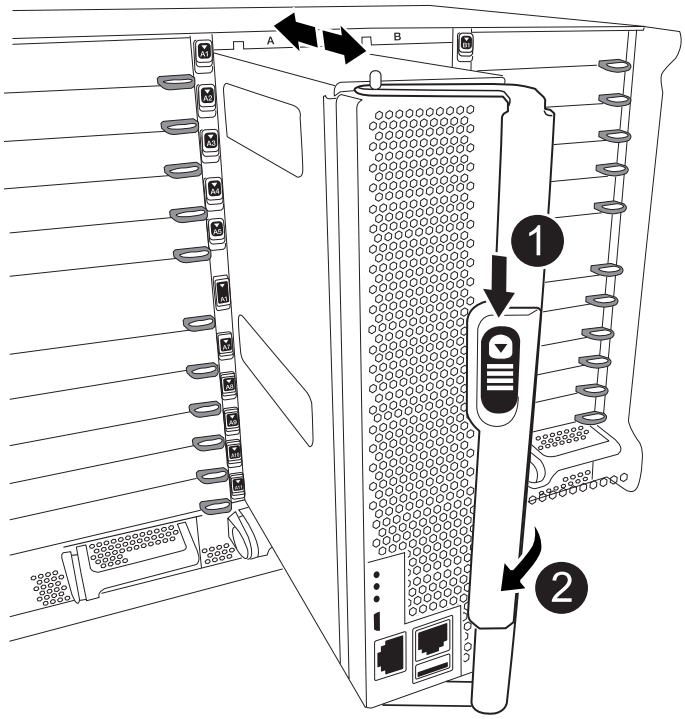
Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were

connected.

3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

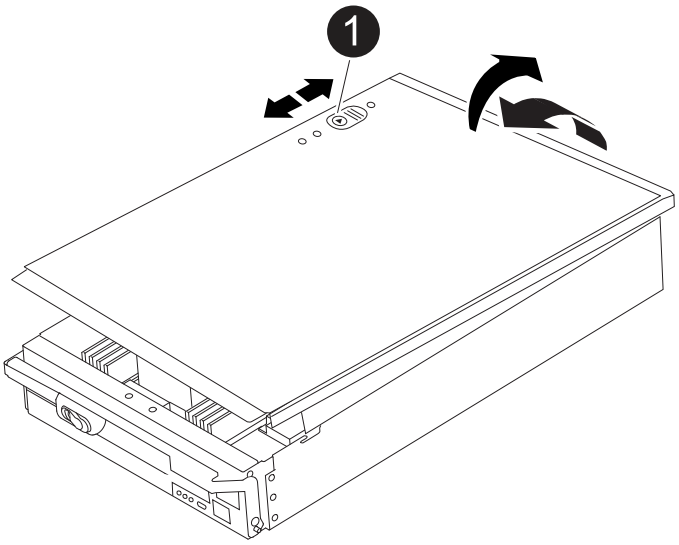


1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

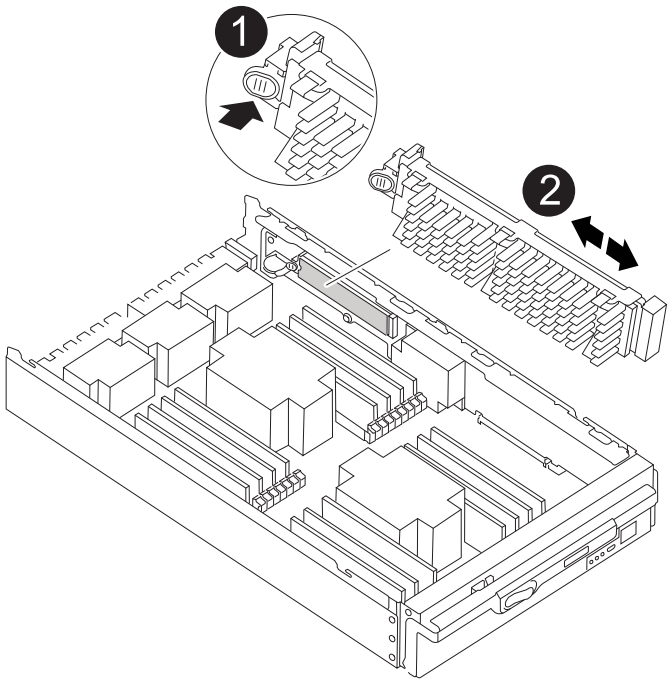
5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.



1	Controller module cover locking button

6. Replace the boot media:
   a. Lift the black air duct at the back of the controller module and then locate the boot media using the following illustration or the FRU map on the controller module:

      Animation - Replace boot media



1	Press release tab
2	Boot media

b. 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.

c. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.

d. 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.

e. Push the boot media down to engage the locking button on the boot media housing.

7. Reinstall the controller module lid by aligning the pins on the lid with the slots on the motherboard carrier, and then slide the lid into place.
8. Reinstall the controller module:
   a. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
   b. Recable the controller module, as needed.
   c. 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, and then push the cam handle to the closed position.

      The controller begins to boot as soon as it is completely installed into the chassis.

      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.

9. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:
   a. Boot to Maintenance mode: `boot_ontap maint`
   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t iniitator adapter_name`
   c. Halt to return to Maintenance mode: `halt`

What's next

After physically replacing the impaired boot media, restore the ONTAP image from the partner node.

Automated boot media recovery from the partner node - FAS9500

After installing the new boot media device in your FAS9500 system, you can start the automated boot media recovery process to restore the configuration from the partner node. During the recovery process, the system checks whether encryption is enabled and determines the type of key encryption in use. If key encryption is enabled, the system guides you through the appropriate steps to restore it.

The automated boot media recovery process is supported only in ONTAP 9.17.1 and later. If your storage system is running an earlier version of ONTAP, use the manual boot recovery procedure.

Before you begin

- Determine your key manager type:
  - Onboard Key Manager (OKM): Requires cluster-wide passphrase and backup data
  - External Key Manager (EKM): Requires the following files from the partner node:
    - `/cfcard/kmip/servers.cfg`
    - `/cfcard/kmip/certs/client.crt`
    - `/cfcard/kmip/certs/client.key`
    - `/cfcard/kmip/certs/CA.pem`

Steps

1. From the LOADER prompt, start the boot media recovery process:

   ```
   boot_recovery -partner
   ```

   The screen displays the following message:

   ```
   Starting boot media recovery (BMR) process. Press Ctrl-C to abort…
   ```

2. Monitor the boot media install recovery process.

   The process completes and displays the `Installation complete` message.

3. The system checks for encryption and displays one of the following messages:

If you see this message…	Do this…
`key manager is not configured. Exiting.`	Encryption is not installed on the system. a. Wait for the login prompt to display. b. Log into the node and give back the storage: `storage failover giveback -ofnode` *`impaired_node_name`* c. Go to re-enabling automatic giveback if it was disabled.
`key manager is configured.`	Encryption is installed. Go to restoring the key manager.

> If the system cannot identify the key manager configuration, it displays an error message and prompts you to confirm whether key manager is configured and which type (onboard or external). Answer the prompts to proceed.

1. Restore the key manager using the appropriate procedure for your configuration:

Onboard Key Manager (OKM)

The system displays the following message and begins running BootMenu Option 10:

```
key manager is configured.
Entering Bootmenu Option 10...

This option must be used only in disaster recovery procedures. Are
you sure? (y or n):
```

a. Enter `y` at the prompt to confirm you want to start the OKM recovery process.

b. Enter the passphrase for onboard key management when prompted.

c. Enter the passphrase again when prompted to confirm.

d. Enter the backup data for onboard key manager when prompted.

Show example of passphrase and backup data prompts

```
Enter the passphrase for onboard key management:
-----BEGIN PASSPHRASE-----
<passphrase_value>
-----END PASSPHRASE-----
Enter the passphrase again to confirm:
-----BEGIN PASSPHRASE-----
<passphrase_value>
-----END PASSPHRASE-----
Enter the backup data:
-----BEGIN BACKUP-----
<passphrase_value>
-----END BACKUP-----
```

e. Monitor the recovery process as it restores the appropriate files from the partner node.

When the recovery process is complete, the node reboots. The following messages indicate a successful recovery:

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.keydb file.

Successfully recovered keymanager secrets.
```

f. After the node reboots, verify that the system is back online and operational.

g. Return the impaired controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

h. After the partner node is fully up and serving data, synchronize the OKM keys across the cluster:

```
security key-manager onboard sync
```

Go to re-enabling automatic giveback if it was disabled.

External Key Manager (EKM)

The system displays the following message and begins running BootMenu Option 11:

```
key manager is configured.
Entering Bootmenu Option 11...
```

a. Enter the EKM configuration settings when prompted:

i. Enter the client certificate contents from the `/cfcard/kmip/certs/client.crt` file:

Show example of client certificate contents

```
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----
```

ii. Enter the client key file contents from the `/cfcard/kmip/certs/client.key` file:

Show example of client key file contents

```
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----
```

iii. Enter the KMIP server CA(s) file contents from the `/cfcard/kmip/certs/CA.pem` file:

Show example of KMIP server file contents

```
-----BEGIN CERTIFICATE-----
<KMIP_certificate_CA_value>
-----END CERTIFICATE-----
```

iv. Enter the server configuration file contents from the `/cfcard/kmip/servers.cfg` file:

Show example of server configuration file contents

```
xxx.xxx.xxx.xxx:5696.host=xxx.xxx.xxx.xxx
xxx.xxx.xxx.xxx:5696.port=5696
xxx.xxx.xxx.xxx:5696.trusted_file=/cfcard/kmip/certs/CA.pem
xxx.xxx.xxx.xxx:5696.protocol=KMIP1_4
1xxx.xxx.xxx.xxx:5696.timeout=25
xxx.xxx.xxx.xxx:5696.nbio=1
xxx.xxx.xxx.xxx:5696.cert_file=/cfcard/kmip/certs/client.cr
t
xxx.xxx.xxx.xxx:5696.key_file=/cfcard/kmip/certs/client.key
xxx.xxx.xxx.xxx:5696.ciphers="TLSv1.2:kRSA:!CAMELLIA:!IDEA:
!RC2:!RC4:!SEED:!eNULL:!aNULL"
xxx.xxx.xxx.xxx:5696.verify=true
xxx.xxx.xxx.xxx:5696.netapp_keystore_uuid=<id_value>
```

v. If prompted, enter the ONTAP Cluster UUID from the partner node. You can check the cluster UUID from the partner node using the `cluster identify show` command.

Show example of ONTAP Cluster UUID prompt

```
Notice: bootarg.mgwd.cluster_uuid is not set or is empty.
Do you know the ONTAP Cluster UUID? {y/n} y
Enter the ONTAP Cluster UUID: <cluster_uuid_value>


System is ready to utilize external key manager(s).
```

vi. If prompted, enter the temporary network interface and settings for the node:

- The IP address for the port
- The netmask for the port
- The IP address of the default gateway

Show example of temporary network setting prompts

```
In order to recover key information, a temporary network
interface needs to be
configured.

Select the network port you want to use (for example,
'e0a')
e0M

Enter the IP address for port : xxx.xxx.xxx.xxx
Enter the netmask for port : xxx.xxx.xxx.xxx
Enter IP address of default gateway: xxx.xxx.xxx.xxx
Trying to recover keys from key servers....
[discover_versions]
[status=SUCCESS reason= message=]
```

b. Verify the key restoration status:

- If you see `kmip2_client: Successfully imported the keys from external key server: xxx.xxx.xxx.xxx:5696` in the output, the EKM configuration has been successfully restored. The process restores the appropriate files from the partner node and reboots the node. Proceed to the next step.
- If the key is not successfully restored, the system halts and displays error and warning messages. Rerun the recovery process from the LOADER prompt: `boot_recovery -partner`

Show example of key recovery error and warning messages

```
ERROR: kmip_init: halting this system with encrypted
mroot...
WARNING: kmip_init: authentication keys might not be
available.
*****************************************************
*                 A T T E N T I O N                 *
*                                                   *
*       System cannot connect to key managers.      *
*                                                   *
*****************************************************
ERROR: kmip_init: halting this system with encrypted
mroot...
.
Terminated

Uptime: 11m32s
System halting...

LOADER-B>
```

c. After the node reboots, verify that the system is back online and operational.

d. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -ofnode impaired_node_name
```

Go to re-enabling automatic giveback if it was disabled.

5. If automatic giveback was disabled, reenable it:

```
storage failover modify -node local -auto-giveback true
```

6. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

What's next

After you've restored the ONTAP image and the node is up and serving data, you return the failed part to NetApp.

Return the failed boot media to NetApp - FAS9500

If a component in your FAS9500 system fails, return the failed part to NetApp. See the Part Return and Replacements page for further information.

Boot media - manual recovery

Boot media manual recovery workflow - FAS9500

Get started with replacing the boot media in your FAS9500 storage system by reviewing the replacement requirements, checking encryption status, shutting down the controller, replacing the boot media, booting the recovery image, restoring encryption, and verifying the system functionality.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.



Review the boot media requirements

Review the requirements for replacing the boot media.



Check encryption key support and status

Determine whether the system has security key manager enabled or encrypted disks.



Shut down the controller

Shut down the controller when when you need to replace the boot media.



Replace the boot media

Remove the failed boot media from the System Management module and install the replacement boot media, and then transfer an ONTAP image using a USB flash drive.



Boot the recovery image

Boot the ONTAP image from the USB drive, restore the file system, and verify the environmental variables.



Restore encryption

Restore the onboard key manager configuration or the external key manager from the ONATP boot menu.



Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Requirements for manual boot media recovery - FAS9500

Before replacing the boot media in your FAS9500 system, ensure you meet the necessary requirements for a successful replacement. This includes making sure you have a USB flash drive with the appropriate amount of storage and verifying that you have the correct replacement boot device.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

USB flash drive

- Ensure you have a USB flash drive formatted to FAT32.
- The USB must have sufficient storage capacity to hold the `image_xxx.tgz` file.

File preparation

Copy the `image_xxx.tgz` file to the USB flash drive. This file will be used when you transfer the ONTAP image using the USB flash drive.

Component replacement

Replace the failed component with the replacement component provided by NetApp.

Controller identification

It is critical to apply the commands to the correct controller when you are replacing the impaired boot media:

- The *impaired controller* is the controller on which you are performing maintenance.
- The *healthy controller* is the HA partner of the impaired controller.

What's next?

After you've reviewed the requirements to replace the boot media, you need to check encryption key support and status on the boot media.

Check encryption key support and status - FAS9500

To ensure data security on your storage system, you need to verify the encryption key support and status on your boot media. Check if your ONTAP version supports NetApp Volume Encryption (NVE), and before you shut down the controller check if the key manager is active.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Step 1: Check NVE support and download the correct ONTAP image

Determine whether your ONTAP version supports NetApp Volume Encryption (NVE) so you can download the correct ONTAP image for the boot media replacement.

Steps

1. Check if your ONTAP version supports encryption:

   ```
   version -v
   ```

   If the output includes `10no-DARE`, NVE is not supported on your cluster version.

2. Download the appropriate ONTAP image based on NVE support:
   - If NVE is supported: Download the ONTAP image with NetApp Volume Encryption
   - If NVE is not supported: Download the ONTAP image without NetApp Volume Encryption

Download the ONTAP image from the NetApp Support Site to your HTTP or FTP server or a local folder. You will need this image file during the boot media replacement procedure.

Step 2: Verify key manager status and back up configuration

Before shutting down the impaired controller, verify the key manager configuration and back up the necessary information.

Steps

1. Determine which key manager is enabled on your system:

ONTAP version	Run this command
ONTAP 9.14.1 or later	`security key-manager keystore show` • If EKM is enabled, `EKM` is listed in the command output. • If OKM is enabled, `OKM` is listed in the command output. • If no key manager is enabled, `No key manager keystores configured` is listed in the command output.
ONTAP 9.13.1 or earlier	`security key-manager show-key-store` • If EKM is enabled, `external` is listed in the command output. • If OKM is enabled, `onboard` is listed in the command output. • If no key manager is enabled, `No key managers configured` is listed in the command output.

2. Depending on whether a key manager is configured on your system, do one of the following:

   If no key manager is configured:

   You can safely shut down the impaired controller and proceed to the shutdown procedure.

   If a key manager is configured (EKM or OKM):

   a. Enter the following query command to display the status of the authentication keys in your key manager:

      ```
      security key-manager key query
      ```

   b. Review the output and check the value in the `Restored` column. This column indicates whether the authentication keys for your key manager (either EKM or OKM) have been successfully restored.

3. Complete the appropriate procedure based on your key manager type:

External Key Manager (EKM)

Complete these steps based on the value in the `Restored` column.

If all keys show `true` in the Restored column:

You can safely shut down the impaired controller and proceed to the shutdown procedure.

If any keys show a value other than `true` in the Restored column:

a. Restore the external key management authentication keys to all nodes in the cluster:

   `security key-manager external restore`

   If the command fails, contact NetApp Support.

b. Verify that all authentication keys are restored:

   `security key-manager key query`

   Confirm that the `Restored` column displays `true` for all authentication keys.

c. If all keys are restored, you can safely shut down the impaired controller and proceed to the shutdown procedure.

Onboard Key Manager (OKM)

Complete these steps based on the value in the `Restored` column.

If all keys show `true` in the Restored column:

a. Back up the OKM information:

   i. Switch to advanced privilege mode:

      `set -priv advanced`

      Enter `y` when prompted to continue.

   ii. Display the key management backup information:

      `security key-manager onboard show-backup`

   iii. Copy the backup information to a separate file or your log file.

      You will need this backup information if you need to manually recover OKM during the replacement procedure.

   iv. Return to admin mode:

      `set -priv admin`

b. You can safely shut down the impaired controller and proceed to the shutdown procedure.

If any keys show a value other than `true` in the Restored column:

a. Synchronize the onboard key manager:

```
security key-manager onboard sync
```

Enter the 32-character alphanumeric onboard key management passphrase when prompted.

> This is the cluster-wide passphrase you created when you initially configured the Onboard Key Manager. If you do not have this passphrase, contact NetApp Support.

b. Verify all authentication keys are restored:

```
security key-manager key query
```

Confirm that the `Restored` column displays `true` for all authentication keys and the `Key Manager` type shows `onboard`.

c. Back up the OKM information:

i. Switch to advanced privilege mode:

```
set -priv advanced
```

Enter `y` when prompted to continue.

ii. Display the key management backup information:

```
security key-manager onboard show-backup
```

iii. Copy the backup information to a separate file or your log file.

You will need this backup information if you need to manually recover OKM during the replacement procedure.

iv. Return to admin mode:

```
set -priv admin
```

d. You can safely shut down the impaired controller and proceed to the shutdown procedure.

Shut down the controller for manual boot media recovery - FAS9500

Shut down or take over the impaired controller using one of the following options.

After completing the NVE or NSE tasks, you need to complete the shutdown of the impaired node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   `system node autosupport invoke -node * -type all -message MAINT=<# of hours>h`

   The following AutoSupport message suppresses automatic case creation for two hours:

   `cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      `storage failover modify -node impaired_node_name -auto-giveback false`

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Replace the boot media and prepare for manual boot recovery - FAS9500

You must unplug the controller module, remove and open the controller module, locate and replace the boot media in the controller, and then transfer the image to the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your

system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

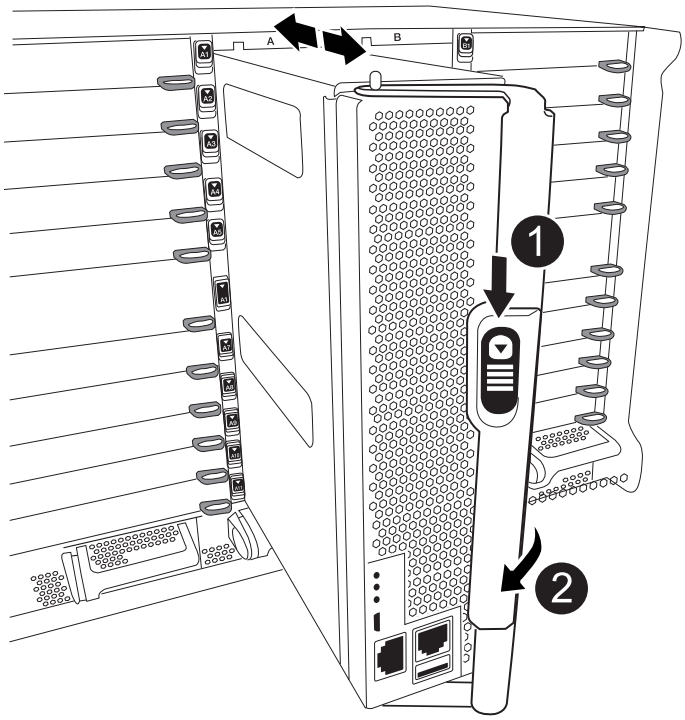
Step 1: Remove the controller module

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.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

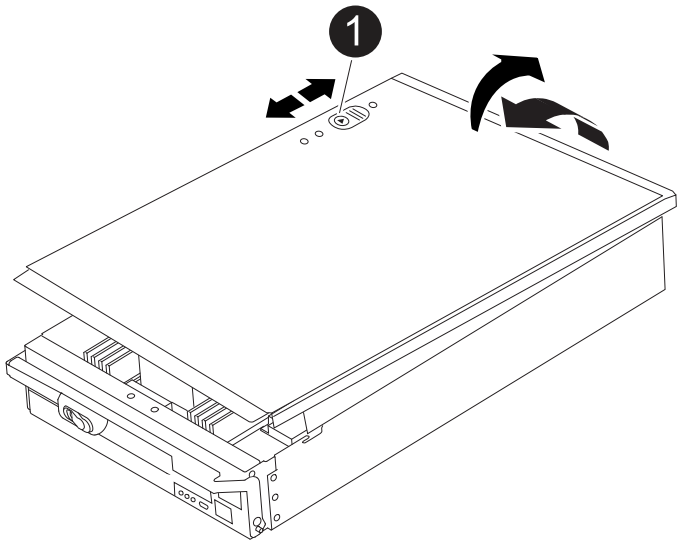


1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.



1	Controller module cover locking button

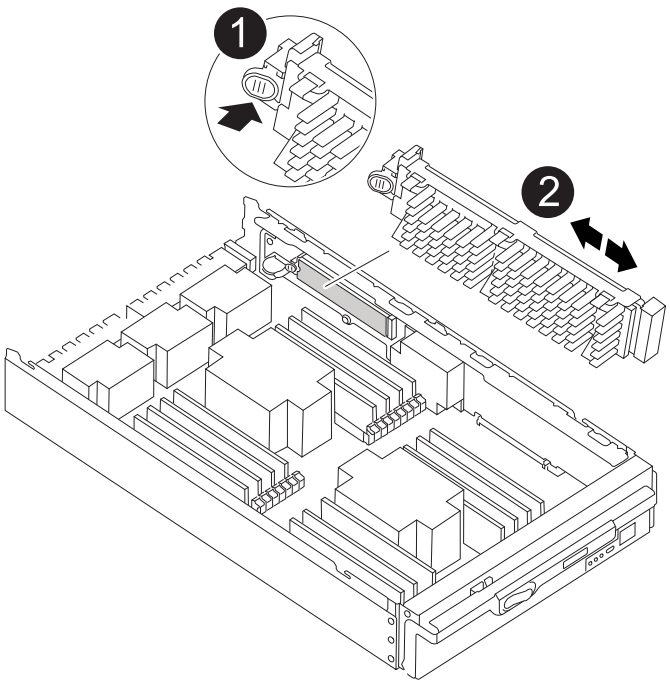
Step 2: Replace the boot media

You must locate the boot media in the controller and follow the directions to replace it.

Steps

1. Lift the black air duct at the back of the controller module and then locate the boot media using the following illustration or the FRU map on the controller module:

   Animation - Replace boot media



1	Press release tab
2	Boot media

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. Align the edges of the replacement boot media with the boot media socket, and then gently push it into the socket.
4. Check the boot media to make sure that it is seated squarely and completely in the socket.

   If necessary, remove the boot media and reseat it into the socket.

5. Push the boot media down to engage the locking button on the boot media housing.
6. Reinstall the controller module lid by aligning the pins on the lid with the slots on the motherboard carrier, and then slide the lid into place.

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.

Before you begin

- You must have a USB flash drive, formatted to FAT32, with at least 4GB capacity.
- Download a copy of the same image version of ONTAP as what the impaired controller was running. You can download the appropriate image from the Downloads section on the NetApp Support Site. Use the `version -v` command to display if your version of ONTAP supports NVE. If the command output displays `<10no- DARE>`, your version of ONTAP does not support NVE.
  - If NVE is supported by your version of ONTAP, download the image with NetApp Volume Encryption, as indicated in the download button.
  - If NVE is not supported, download the image without NetApp Volume Encryption, as indicated in the download button.
- If your system is 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. If you have not done so, download and copy the appropriate service image from the NetApp Support Site to the USB flash drive.
   a. Download the service image from the Downloads link on the page, to your work space on your laptop.
   b. Unzip the service image.

> If you are extracting the contents using Windows, do not use WinZip to extract the netboot image. Use another extraction tool, such as 7-Zip or WinRAR.

The USB flash drive should have the appropriate ONTAP image of what the impaired controller is running.

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.
3. Recable the controller module, as needed.
4. 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.

5. 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, and then push the cam handle to the closed position.

   The controller begins to boot as soon as it is completely installed into the chassis.

6. 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.

7. If the controller is in a stretch or fabric-attached MetroCluster, you must restore the FC adapter configuration:

   a. Boot to Maintenance mode: `boot_ontap maint`
   b. Set the MetroCluster ports as initiators: `ucadmin modify -m fc -t initiator adapter_name`
   c. Halt to return to Maintenance mode: `halt`

   The changes will be implemented when the system is booted.

Manual boot media recovery from a USB drive - FAS9500

After installing the new boot media device in your system, you can boot the recovery image from a USB drive and restore the configuration from the partner node.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Before you begin

- Ensure your console is connected to the impaired controller.
- Verify you have a USB flash drive with the recovery image.
- Determine if your system uses encryption. You will need to select the appropriate option in step 3 based on whether encryption is enabled.

Steps

1. From the LOADER prompt on the impaired controller, boot the recovery image from the USB flash drive:

   `boot_recovery`

The recovery image is downloaded from the USB flash drive.

2. When prompted, enter the name of the image or press **Enter** to accept the default image displayed in brackets.
3. Restore the var file system using the procedure for your ONTAP version:

ONTAP 9.16.0 or earlier

Complete the following steps on the impaired controller and partner controller:

a. **On the impaired controller:** Press `Y` when you see `Do you want to restore the backup configuration now?`

b. **On the impaired controller:** If prompted, press `Y` to overwrite */etc/ssh/ssh_host_ecdsa_key*.

c. **On the partner controller:** Set the impaired controller to advanced privilege level:

```
set -privilege advanced
```

d. **On the partner controller:** Run the restore backup command:

```
system node restore-backup -node local -target-address
impaired_node_IP_address
```

> If you see any message other than a successful restore, contact NetApp Support.

e. **On the partner controller:** Return to admin level:

```
set -privilege admin
```

f. **On the impaired controller:** Press `Y` when you see `Was the restore backup procedure successful?`

g. **On the impaired controller:** Press `Y` when you see `...would you like to use this restored copy now?`

h. **On the impaired controller:** Press `Y` when prompted to reboot, then press `Ctrl-C` when you see the Boot Menu.

i. **On the impaired controller:** Do one of the following:

- If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.
- If the system uses encryption, go to Restore encryption.

ONTAP 9.16.1 or later

Complete the following steps on the impaired controller:

a. Press `Y` when prompted to restore the backup configuration.

   After the restore procedure is successful, this message displays: `syncflash_partner: Restore from partner complete`

b. Press `Y` when prompted to confirm that the restore backup was successful.

c. Press `Y` when prompted to use the restored configuration.

d. Press `Y` when prompted to reboot the node.

e. Press `Y` when prompted to reboot again, then press `Ctrl-C` when you see the Boot Menu.

f. Do one of the following:

- If the system does not use encryption, select *Option 1 Normal Boot* from the Boot Menu.

- If the system uses encryption, go to Restore encryption.

4. Connect the console cable to the partner controller.
5. Return the controller to normal operation by giving back its storage:

```
storage failover giveback -fromnode local
```

6. If you disabled automatic giveback, reenable it:

```
storage failover modify -node local -auto-giveback true
```

7. If AutoSupport is enabled, restore automatic case creation:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Restore encryption - FAS9500

Restore encryption on the replacement boot media.

If your storage system is running ONTAP 9.17.1 or later, use the automated boot recovery procedure. If your system is running an earlier version of ONTAP, you must use the manual boot recovery procedure.

Complete the appropriate steps to restore encryption on your system based on your key manager type. If you are unsure which key manager your system uses, check the settings you captured at the beginning of the boot media replacement procedure.

Onboard Key Manager (OKM)

Restore the Onboard Key Manager (OKM) configuration from the ONTAP boot menu.

Before you begin

Ensure you have the following information available:

- Cluster-wide passphrase entered while enabling onboard key management
- Backup information for the Onboard Key Manager
- Verification that you have the correct passphrase and backup data using the How to verify onboard key management backup and cluster-wide passphrase procedure

Steps

On the impaired controller:

1. Connect the console cable to the impaired controller.
2. From the ONTAP boot menu, select the appropriate option:

ONTAP version	Select this option
ONTAP 9.8 or later	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---------------------------
0123456789012345678901234567890123456789012345678901234567890123
1234567890123456789012345678901234567890123456789012345678901234
2345678901234567890123456789012345678901234567890123456789012345
3456789012345678901234567890123456789012345678901234567890123456
4567890123456789012345678901234567890123456789012345678901234567
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
```

```
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
01234567890123456789012345678901234567890123456789012345678901
23
12345678901234567890123456789012345678901234567890123456789012
34
23456789012345678901234567890123456789012345678901234567890123
45
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA
AAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAAA
AA

---------------------------END
BACKUP---------------------------
```

b. Press Enter twice at the end of the input.

   The recovery process completes and displays the following message:

   ```
   Successfully recovered keymanager secrets.
   ```

Show example prompt

```
Trying to recover keymanager secrets....
Setting recovery material for the onboard key manager
Recovery secrets set successfully
Trying to delete any existing km_onboard.wkeydb file.

Successfully recovered keymanager secrets.

*******************************************************************
********************
* Select option "(1) Normal Boot." to complete recovery
process.
*
* Run the "security key-manager onboard sync" command to
synchronize the key database after the node reboots.
*******************************************************************
********************
```

Do not proceed if the displayed output is anything other than `Successfully recovered keymanager secrets`. Perform troubleshooting to correct the error.

6. Select option `1` from the boot menu to continue booting into ONTAP.

Show example prompt

```
*******************************************************************
******************
* Select option "(1) Normal Boot." to complete the recovery
process.
*
*******************************************************************
******************


(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

7. Confirm that the controller's console displays the following message:

   `Waiting for giveback…(Press Ctrl-C to abort wait)`

   On the partner controller:

8. Giveback the impaired controller:

   `storage failover giveback -fromnode local -only-cfo-aggregates true`

   On the impaired controller:

9. After booting with only the CFO aggregate, synchronize the key manager:

   `security key-manager onboard sync`

10. Enter the cluster-wide passphrase for the Onboard Key Manager when prompted.

Show example prompt

```
Enter the cluster-wide passphrase for the Onboard Key Manager:

All offline encrypted volumes will be brought online and the
corresponding volume encryption keys (VEKs) will be restored
automatically within 10 minutes. If any offline encrypted
volumes are not brought online automatically, they can be
brought online manually using the "volume online -vserver
<vserver> -volume <volume_name>" command.
```

If the sync is successful, the cluster prompt is returned with no additional messages. If the sync fails, an error message appears before returning to the cluster prompt. Do not continue until the error is corrected and the sync runs successfully.

11. Verify that all keys are synced:

    ```
    security key-manager key query -restored false
    ```

    The command should return no results. If any results appear, repeat the sync command until no results are returned.

    On the partner controller:

12. Giveback the impaired controller:

    ```
    storage failover giveback -fromnode local
    ```

13. Restore automatic giveback if you disabled it:

    ```
    storage failover modify -node local -auto-giveback true
    ```

14. If AutoSupport is enabled, restore automatic case creation:

    ```
    system node autosupport invoke -node * -type all -message MAINT=END
    ```

External Key Manager (EKM)

Restore the External Key Manager configuration from the ONTAP boot menu.

Before you begin

Gather the following files from another cluster node or from your backup:

- `/cfcard/kmip/servers.cfg` file or the KMIP server address and port
- `/cfcard/kmip/certs/client.crt` file (client certificate)
- `/cfcard/kmip/certs/client.key` file (client key)
- `/cfcard/kmip/certs/CA.pem` file (KMIP server CA certificates)

Steps

On the impaired controller:

1. Connect the console cable to the impaired controller.
2. Select option `11` from the ONTAP boot menu.

   Show example boot menu

   ```
   (1)  Normal Boot.
   (2)  Boot without /etc/rc.
   (3)  Change password.
   (4)  Clean configuration and initialize all disks.
   (5)  Maintenance mode boot.
   (6)  Update flash from backup config.
   (7)  Install new software first.
   (8)  Reboot node.
   (9)  Configure Advanced Drive Partitioning.
   (10) Set Onboard Key Manager recovery secrets.
   (11) Configure node for external key management.
   Selection (1-11)? 11
   ```

3. Confirm you have gathered the required information when prompted:

   Show example prompt

   ```
   Do you have a copy of the /cfcard/kmip/certs/client.crt file?
   {y/n}
   Do you have a copy of the /cfcard/kmip/certs/client.key file?
   {y/n}
   Do you have a copy of the /cfcard/kmip/certs/CA.pem file? {y/n}
   Do you have a copy of the /cfcard/kmip/servers.cfg file? {y/n}
   ```

4. Enter the client and server information when prompted:
   a. Enter the client certificate (client.crt) file contents, including the BEGIN and END lines.
   b. Enter the client key (client.key) file contents, including the BEGIN and END lines.
   c. Enter the KMIP server CA(s) (CA.pem) file contents, including the BEGIN and END lines.
   d. Enter the KMIP server IP address.
   e. Enter the KMIP server port (press Enter to use the default port 5696).

Show example

```
Enter the client certificate (client.crt) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the client key (client.key) file contents:
-----BEGIN RSA PRIVATE KEY-----
<key_value>
-----END RSA PRIVATE KEY-----

Enter the KMIP server CA(s) (CA.pem) file contents:
-----BEGIN CERTIFICATE-----
<certificate_value>
-----END CERTIFICATE-----

Enter the IP address for the KMIP server: 10.10.10.10
Enter the port for the KMIP server [5696]:

System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
kmip_init: configuring ports
Running command '/sbin/ifconfig e0M'
..
..
kmip_init: cmd: ReleaseExtraBSDPort e0M
```

The recovery process completes and displays the following message:

```
Successfully recovered keymanager secrets.
```

Show example

```
System is ready to utilize external key manager(s).
Trying to recover keys from key servers....
Performing initialization of OpenSSL
Successfully recovered keymanager secrets.
```

5. Select option `1` from the boot menu to continue booting into ONTAP.

Show example prompt

```
*********************************************************************
***********
* Select option "(1) Normal Boot." to complete the recovery
process.
*
*********************************************************************
***********

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 1
```

6. Restore automatic giveback if you disabled it:

   ```
   storage failover modify -node local -auto-giveback true
   ```

7. If AutoSupport is enabled, restore automatic case creation:

   ```
   system node autosupport invoke -node * -type all -message MAINT=END
   ```

Return the failed boot media to NetApp - FAS9500

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

Replace the chassis - FAS9500

Before you begin

To replace the chassis, you must remove the power supplies, fans, controller modules, I/O modules, DCPM modules, and USB LED module from the impaired chassis, remove the impaired chassis from the equipment rack or system cabinet, install the replacement chassis in its place, and then install the components into the replacement chassis.

All other components in the system must be functioning properly; if not, you must contact technical support.

About this task

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shutdown the impaired controller - FAS9500

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
  - Local administrator credentials for ONTAP.
  - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
  - Perform additional system health checks.
  - Upgrade ONTAP to a recommended release for the system.
  - Resolve any Active IQ Wellness Alerts and Risks.
    Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

   ```
   system node autosupport invoke -node * -type all -message "MAINT=2h Replace chassis"
   ```

5. Identify the SP/BMC address of all cluster nodes:

   ```
   system service-processor show -node * -fields address
   ```

6. Exit the cluster shell:

   ```
   exit
   ```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
true -ignore-quorum-warnings true -inhibit-takeover true
```

> For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict-sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

```
Warning: Are you sure you want to halt node <node_name>? {y|n}:
```

10. Wait for each controller to halt and display the LOADER prompt.

Move and replace hardware - FAS9500

To replace the chassis, you must remove the components from the impaired chassis and install them in the replacement chassis.

Step 1: Remove the power supplies

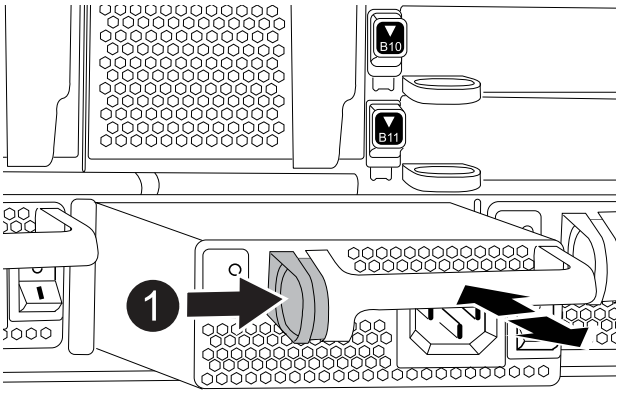
Removing the power supplies when replacing a chassis involves turning off, disconnecting, and then removing the four power supplies from the rear of the impaired chassis.

Steps

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. Press and hold the terra cotta locking button on the power supply handle, and then pull the power supply out of the chassis.

> When removing a power supply, always use two hands to support its weight.

Animation - Remove/install PSU



1	Terra cotta locking button

4. Repeat the preceding steps for any remaining power supplies.

Step 2: Remove the fans

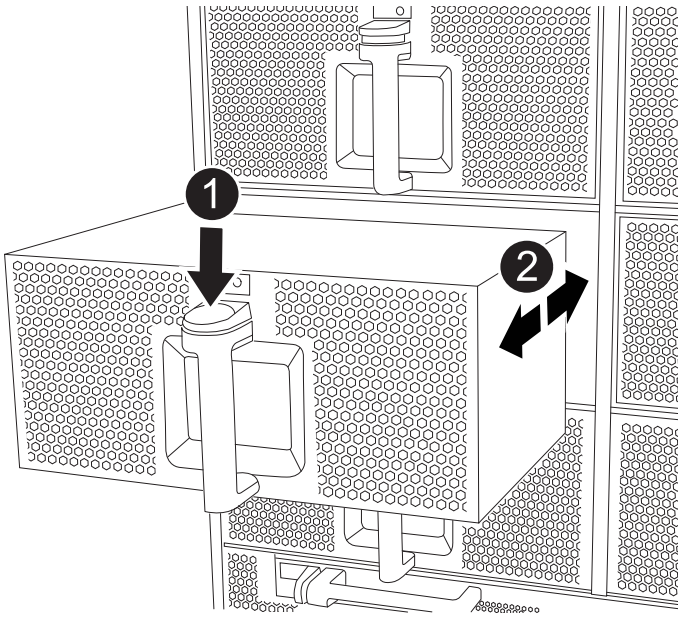
You must remove the six fan modules, located on in the front of the chassis, when replacing the chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Press the terra cotta locking button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

> The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

Animation - Remove/install fan



1	Terra cotta locking button
2	Slide fan in/out of chassis

4. Set the fan module aside.
5. Repeat the preceding steps for any remaining fan modules.

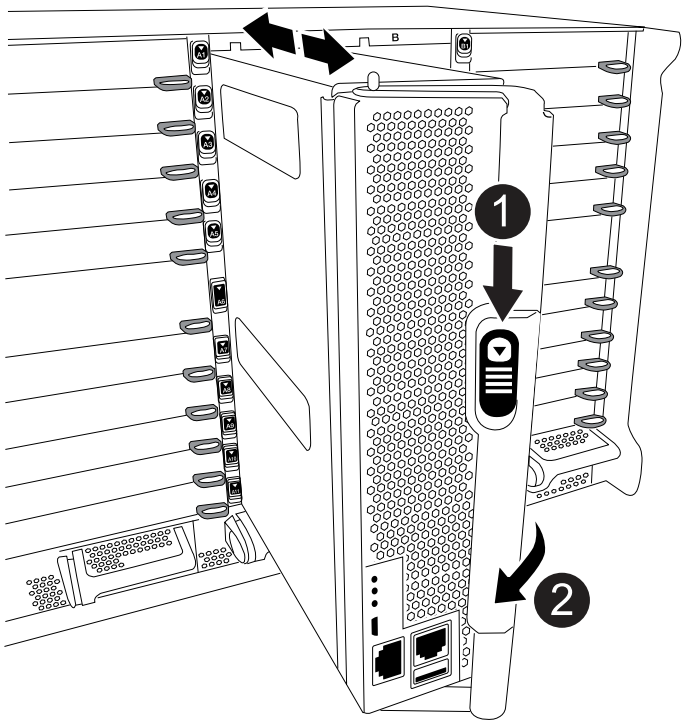
Step 3: Remove the controller module

To replace the chassis, you must remove the controller module or modules from the impaired chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta locking button on the cam handle downward until it unlocks.

   Animation - Remove controller module



1	Cam handle locking button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Set the controller module aside in a safe place and keep track of which chassis slot it came from, so that it can be installed into the same slot in the replacement chassis..
6. Repeat these steps if you have another controller module in the chassis.

Step 4: Remove the I/O modules

To remove I/O modules from the impaired chassis, including the NVRAM modules, follow the specific sequence of steps. You do not have to remove the Flash Cache module, if present, from the NVRAM module when moving it to a replacement chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling associated with the target I/O module.

   Make sure that you label the cables so that you know where they came from.

3. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam locking button.

      The cam locking button moves away from the chassis.

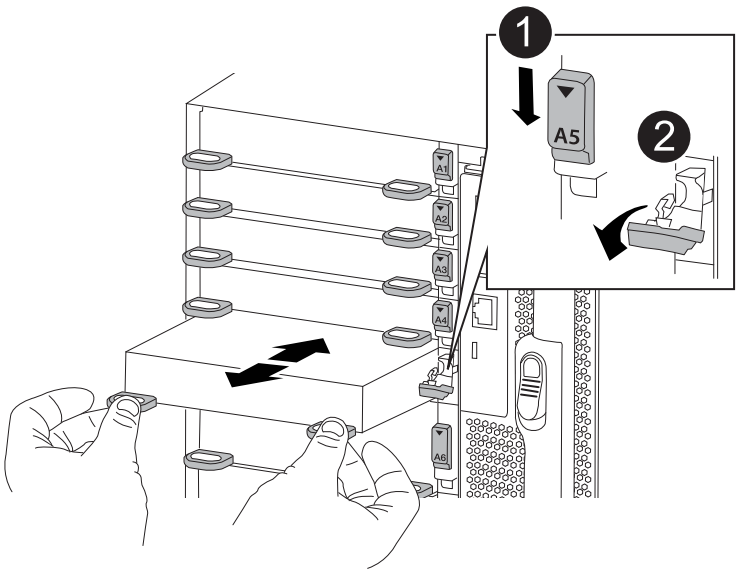
   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

      Animation - Remove/install I/O module



1	Lettered and numbered I/O cam latch

2	I/O cam latch completely unlocked

4. Set the I/O module aside.
5. Repeat the preceding step for the remaining I/O modules in the impaired chassis.

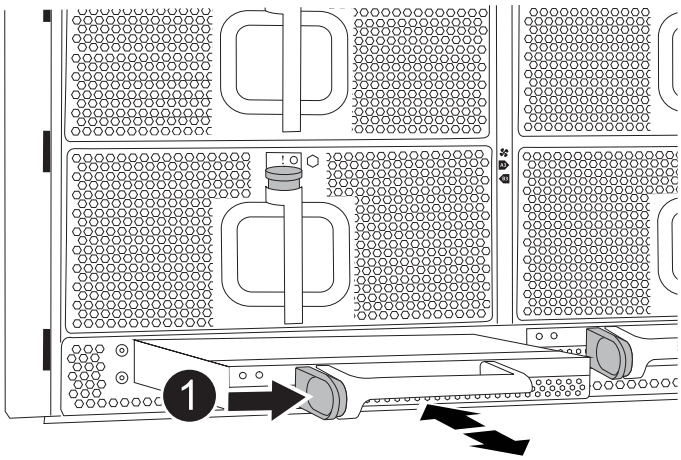
Step 5: Remove the De-stage Controller Power Module

Remove the two de-stage controller power modules from the front of the impaired chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. Press the terra cotta locking button on the module handle, and then slide the DCPM out of the chassis.

   Animation - Remove/install DCPM



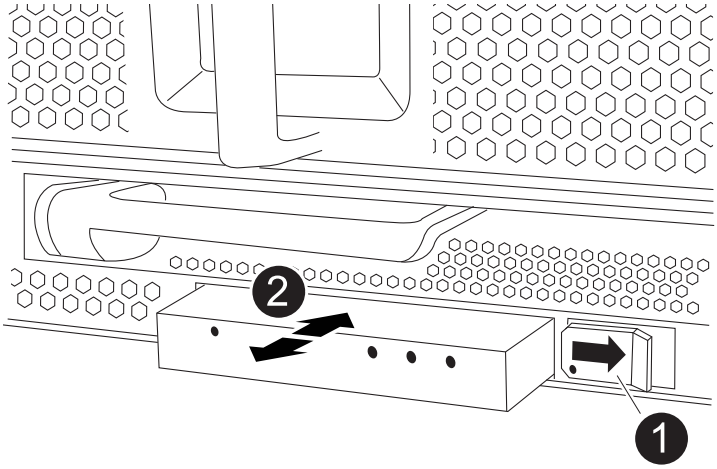
1	DCPM terra cotta locking button

3. Set the DCPM aside in a safe place and repeat this step for the remaining DCPM.

Step 6: Remove the USB LED module

Remove the USB LED modules.

Animation - Remove/install USB module



1	Eject the module.
2	Slide out of chassis.

Steps

1. Locate the USB LED module on the front of the impaired chassis, directly under the power supply bays.
2. Press the black locking button on the right side of the module to release the module from the chassis, and then slide it out of the impaired chassis.
3. Set the module aside in a safe place.

Step 7: Replace a chassis from within the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the replacement chassis.

Steps

1. Remove the screws from the chassis mount points.

   If the system is in a system cabinet, you might need to remove the rear tie-down bracket.

2. With the help of two or three people, slide the impaired 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 impaired chassis.
7. Secure the rear of the chassis to the equipment rack or system cabinet.
8. If you are using the cable management brackets, remove them from the impaired chassis, and then install them on the replacement chassis.

Step 8: Install the de-stage controller power module when replacing the chassis

Once the replacement chassis is installed into the rack or system cabinet, you must reinstall the de-stage controller power modules into it.

Steps

1. If you are not already grounded, properly ground yourself.
2. Align the end of the DCPM with the chassis opening, and then gently slide it into the chassis until it clicks into place.

   The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

3. Repeat this step for the remaining DCPM.

Step 9: Install fans into the chassis

To install the fan modules when replacing the chassis, you must perform a specific sequence of tasks.

Steps

1. If you are not already grounded, properly ground yourself.
2. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

   When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

3. Repeat these steps for the remaining fan modules.
4. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.

Step 10: Install I/O modules

To install I/O modules, including the NVRAM/Flash Cache modules from the impaired chassis, follow the specific sequence of steps.

You must have the chassis installed so that you can install the I/O modules into the corresponding slots in the replacement chassis.

Steps

1. If you are not already grounded, properly ground yourself.
2. After the replacement chassis is installed in the rack or cabinet, install the I/O modules into their corresponding slots in the replacement chassis by gently sliding the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage, and then push the I/O cam latch all the way up to lock the module in place.
3. Recable the I/O module, as needed.
4. Repeat the preceding step for the remaining I/O modules that you set aside.

   If the impaired chassis has blank I/O panels, move them to the replacement chassis at this time.

Step 11: Install the power supplies

Installing the power supplies when replacing a chassis involves installing the power supplies into the replacement chassis, and connecting to the power source.

Steps

1. If you are not already grounded, properly ground yourself.
2. Make sure the power supplies rockers are in the off position.
3. 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 until it locks into place.

   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.

4. 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.

5. Repeat the preceding steps for any remaining power supplies.

Step 12 Install the USB LED modules

Install the USB LED modules in the replacement chassis.

Steps

1. Locate the USB LED module slot on the front of the replacement chassis, directly under the DCPM bays.
2. Align the edges of the module with the USB LED bay, and gently push the module all the way into the chassis until it clicks into place.

Step 13: Install the controller

After you install the controller module and any other components into the replacement chassis, boot the system.

Steps

1. If you are not already grounded, properly ground yourself.
2. Connect the power supplies to different power sources, and then turn them on.
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 console to the controller module, and then reconnect the management port.
5. With the cam handle in the open position, slide the controller module into the chassis and firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle until it clicks into the locked position.

Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors.

The controller module begins to boot as soon as it is fully seated in the chassis.

6. Repeat the preceding steps to install the second controller into the replacement chassis.
7. Boot each controller.

Restore and verify the configuration - FAS9500

To complete the chassis replacement, you must complete specific tasks.

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
      - non-ha
3. Confirm that the setting has changed: `ha-config show`
4. If you have not already done so, recable the rest of your system.

Step 2: Bring up the system

1. If you have not done so, plug the power cables back into the PSUs.
2. Turn on the PSUs by toggling the rocker switched to **ON**, and wait for the controllers to power up completely.
3. Check the front and the back of the chassis and controllers for any fault lights after power up.
4. Connect to the SP or BMC IP address of the nodes via SSH. This will be the same address used to shut down the nodes.
5. Perform additional health checks as described in How_to_perform_a_cluster_health_check_with_a_script_in_ONTAP
6. Turn AutoSupport back on (end the maintenance window message):
   `system node autosupport invoke -node * -type all -message MAINT=end`

As a best practice, you should do the following:

- Resolve any Active IQ Wellness Alerts and Risks (Active IQ will take time to process post-power up AutoSupports - expect a delay in results)
- Run Active IQ Config Advisor
- Check system health using How_to_perform_a_cluster_health_check_with_a_script_in_ONTAP

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

Replace the controller module - FAS9500

To replace the impaired controller module, you must shut down the impaired controller, move the internal components to the replacement controller module, install the replacement controller module, and reboot the replacement controller.

Before you begin

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 has a V_StorageAttach license, you must refer to the additional required steps before performing this procedure.
- If your system is in an HA pair, the healthy node must be able to take over the node that is being replaced (referred to in this procedure as the “impaired node”).
- 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 node in a four or eight node MetroCluster configuration is the same as that in an HA pair. No MetroCluster-specific steps are required because the failure is restricted to an HA pair and storage failover commands can be used to provide nondisruptive operation during the replacement.

- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- In this procedure, the boot device is moved from the impaired node to the replacement node so that the replacement node 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 node is the node that is being replaced.

- The replacement node is the new node that is replacing the impaired node.
- The healthy node is the surviving node.

- You must always capture the node'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 node - FAS9500

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

Replace the controller module hardware - FAS9500

To replace the controller module hardware, you must remove the impaired node, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

The following animation shows the whole process of moving components from the impaired to the replacement controller.

Animation - Replace controller module, complete process

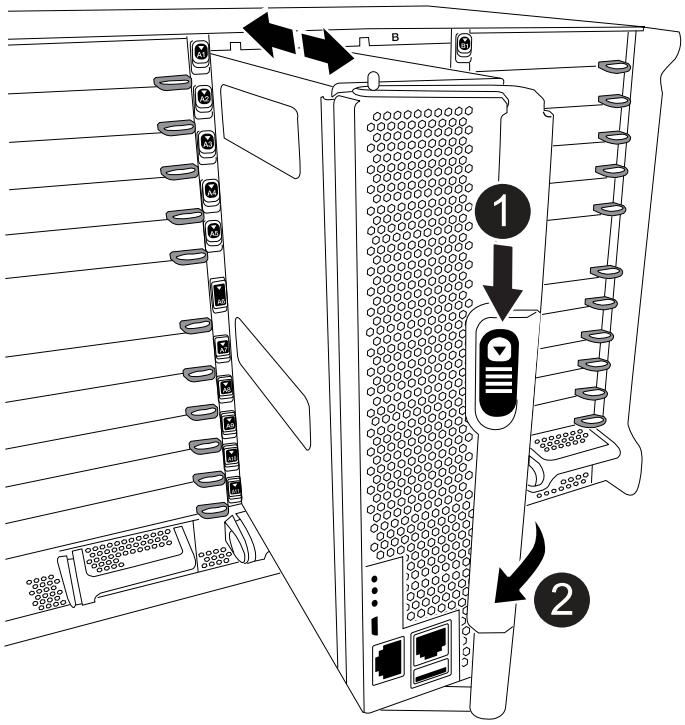
Step 1: Remove the controller module

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.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove controller module

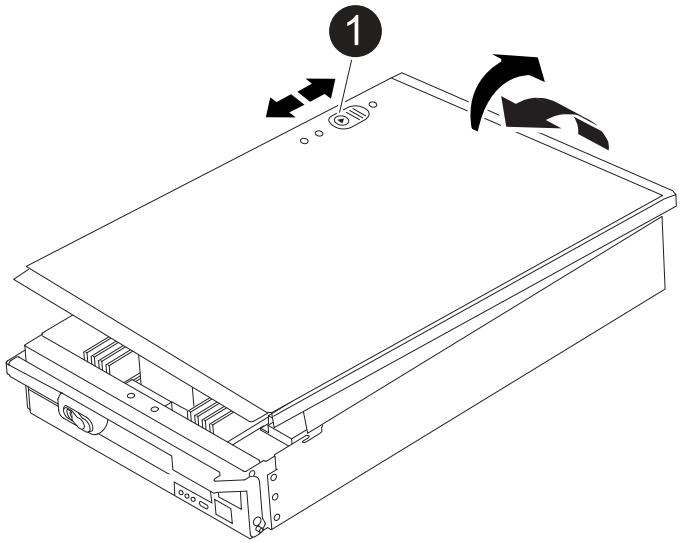


1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.



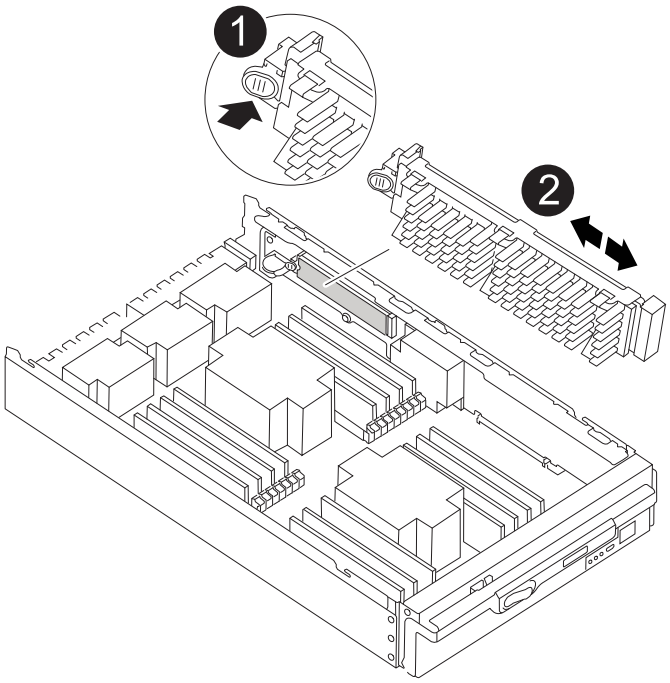
1	Controller module cover locking button

Step 2: Move the boot media

You must locate the boot media and follow the directions to remove it from the old controller and insert it in the new controller.

Steps

1. Locate the boot media using the following illustration or the FRU map on the controller module:



1	Press release tab
2	Boot media

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 3: Move the system DIMMs

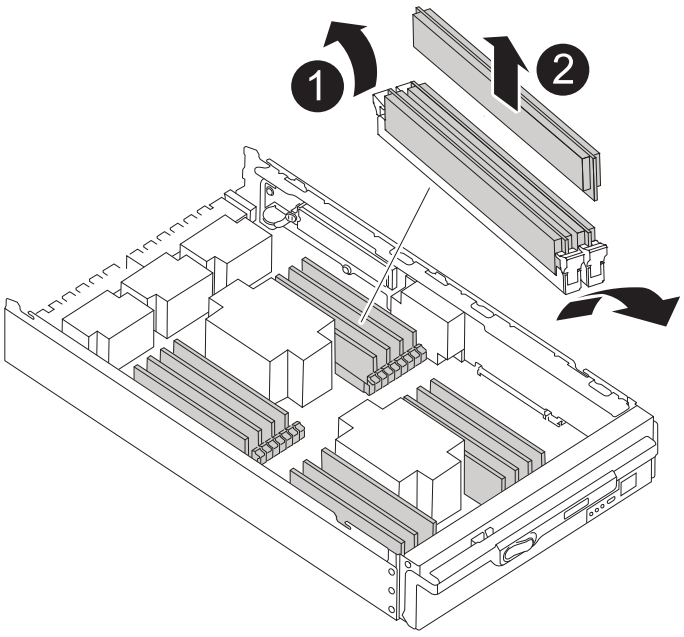
To move the DIMMs, locate and move them from the old controller into the replacement controller and follow the specific sequence of steps.

> The VER2 controller has fewer DIMM sockets. There is no reduction in the number of DIMMs supported or change in the DIMM socket numbering. When moving the DIMMs to the new controller module, install the DIMMs into the same socket number/location as the impaired controller module. See the FRU map diagram on the VER2 controller module for DIMM socket locations.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller module.
3. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
4. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.

> Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.



1	DIMM ejector tabs
2	DIMM

5. Locate the slot where you are installing the DIMM.
6. Make sure that the DIMM ejector tabs on the connector are in the open position, and then insert the DIMM squarely into the slot.

The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.

Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

7. 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. 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.
9. Repeat these steps for the remaining DIMMs.

Step 4: Install the controller

After you install the components into the replacement controller module, you must install the replacement controller module into the system chassis and boot the operating system.

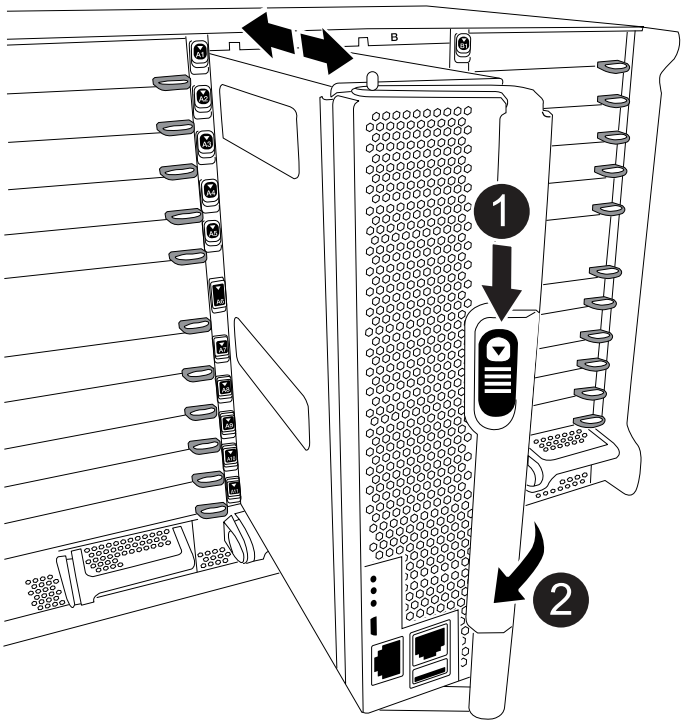
For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

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.

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.

   Animation - Install controller module



1	Cam handle release button
2	Cam handle

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:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

The controller module begins to boot as soon as it is fully seated in the chassis. Be prepared to interrupt the boot process.

   c. Rotate the controller module cam handle to the locked position.
   d. Interrupt the boot process by pressing `Ctrl-C` when you see Press Ctrl-C for Boot Menu.
   e. Select the option to boot to LOADER.

After completing the hardware replacement, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary..

Step 1: Set and verify the system time after replacing the controller module

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 replacement controller module, verify that all components display the same `HA` state: `ha-config show`

If your system is in…	The HA state for all components should be…
An HA pair	ha
A MetroCluster FC configuration with four or more nodes	mcc

If your system is in…	The HA state for all components should be…
A MetroCluster IP configuration	mccip

2. If the displayed system state of the controller module does not match your system configuration, set the `HA` state for the controller module: `ha-config modify controller ha-state`
3. If the displayed system state of the chassis does not match your system configuration, set the `HA` state for the chassis: `ha-config modify chassis ha-state`

Recable the system - FAS9500

Continue the replacement procedure by recabling the storage and network conigurations.

Step 1: Recable the system

You must recable the controller module's storage and network connections.

Steps

1. Recable the system.
2. Verify that the cabling is correct by using Active IQ Config Advisor.
   a. Download and install Config Advisor.
   b. Enter the information for the target system, and then click Collect Data.
   c. 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.
   d. Check other cabling by clicking the appropriate tab, and then examining the output from Config Advisor.

> The system ID and disk assignment information reside in the NVRAM module, which is in a module separate from the controller module and not impacted by the controller module replacement.

Step 2: Reassign disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. You must confirm the system ID change when you boot the *replacement* node and then verify that the change was implemented.

This procedure applies only to systems running ONTAP in an HA pair.

1. If the *replacement* node 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* node, boot the node, 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* node console and then, from the healthy node, 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 node, 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 node, verify that any coredumps are saved:
   a. Change to the advanced privilege level: `set -privilege advanced`

      You can respond `Y` when prompted to continue into advanced mode. The advanced mode prompt appears (*>).

   b. 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 node:
   a. From the healthy node, give back the replaced node's storage: `storage failover giveback -ofnode replacement_node_name`

      The *replacement* node takes back its storage and completes booting.

      If you are prompted to override the system ID due to a system ID mismatch, you should enter `y`.

      > If the giveback is vetoed, you can consider overriding the vetoes.

      For more information, see the Manual giveback commands topic to override the veto.

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: storage failover show

The output from the `storage failover show` command should not include the System ID changed on partner message.

7. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the *replacement* node should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> storage disk show -ownership

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -
1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -
1873775277 Pool0
.
.
.
```

8. If the system is in a MetroCluster configuration, monitor the status of the node: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each node will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

9. If the node is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a node on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The *replacement* node is the current owner of the disks on the disaster site.

     For more information, see Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration topic.

10. If your system is in a MetroCluster configuration, verify that each node is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

11. Verify that the expected volumes are present for each node: `vol show -node node-name`
12. If you disabled automatic takeover on reboot, enable it from the healthy node: `storage failover modify -node replacement-node-name -onreboot true`

Complete system restoration - FAS9500

To complete the replacement procedure and restore your system to full operation, you must recable the storage, 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: Install licenses for the replacement node 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.

Before you begin

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 ONTAP platforms. If you are unsure of the initial ONTAP release for your system, see NetApp Hardware Universe for more information.

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.

- 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 the node is in a MetroCluster configuration and all nodes at a site have been replaced, license keys must be installed on the *replacement* node or nodes prior to switchback.

Steps

1. If you need new license keys, obtain replacement license keys on the NetApp Support site in the My Support section under Software licenses.

   > The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, you should contact technical support.

2. Install each license key: `system license add -license-code license-key, license-key...`
3. Remove the old licenses, if desired:
   a. Check for unused licenses: `license clean-up -unused -simulate`
   b. If the list looks correct, remove the unused licenses: `license clean-up -unused`

Step 2: Verify LIFs and registering the serial number

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, and register the serial number of the *replacement* node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

   If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Register the system serial number with NetApp Support.
   - If AutoSupport is enabled, send an AutoSupport message to register the serial number.
   - If AutoSupport is not enabled, call NetApp Support to register the serial number.
3. Check the health of your cluster. See the How to perform a cluster health check with a script in ONTAP KB article for more information.
4. If an AutoSupport maintenance window was triggered, end it by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 3: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace a DIMM - FAS9500

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.

Before you begin

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 node

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

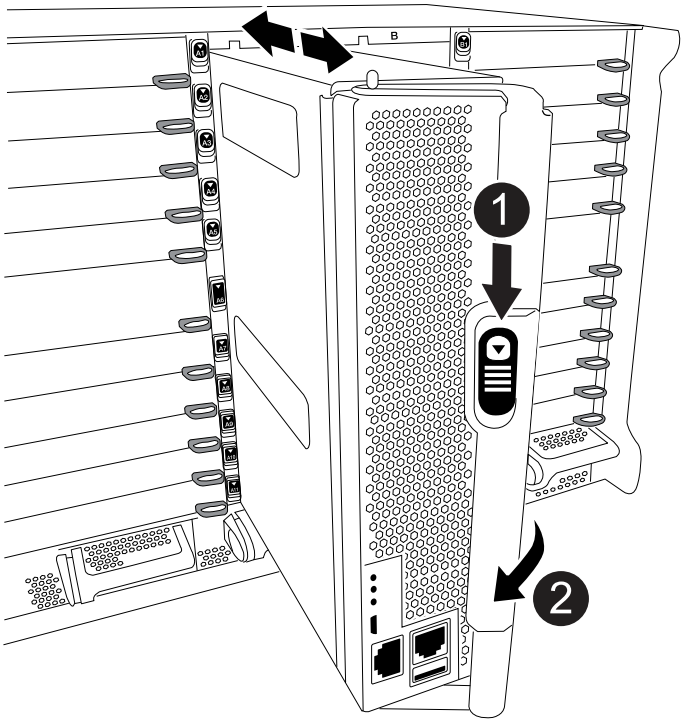
Step 2: Remove the controller module

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. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove the controller

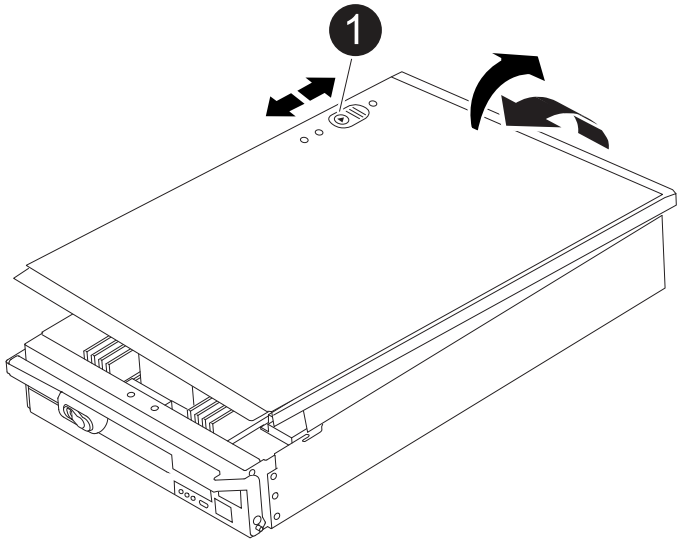


1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.



1	Controller module cover locking button

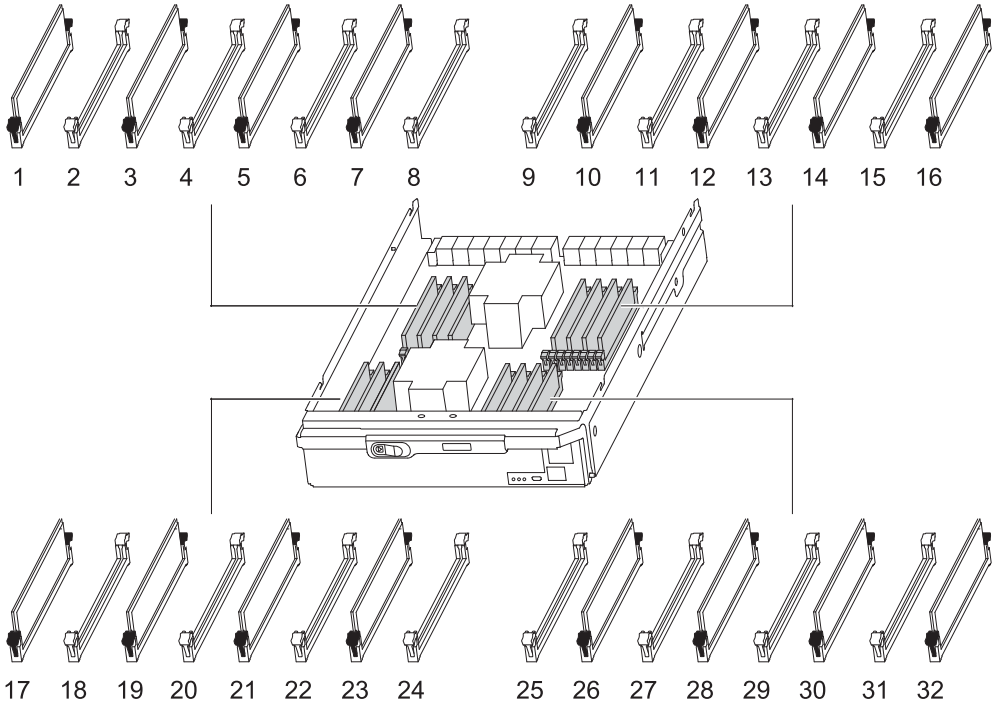
Step 3: Replace the DIMMs

To replace the DIMMs, locate them inside the controller and follow the specific sequence of steps.

> The VER2 controller has fewer DIMM sockets. There is no reduction in the number of DIMMs supported or change in the DIMM socket numbering. When moving the DIMMs to the new controller module, install the DIMMs into the same socket number/location as the impaired controller module. See the FRU map diagram on the VER2 controller module for DIMM socket locations.

Steps

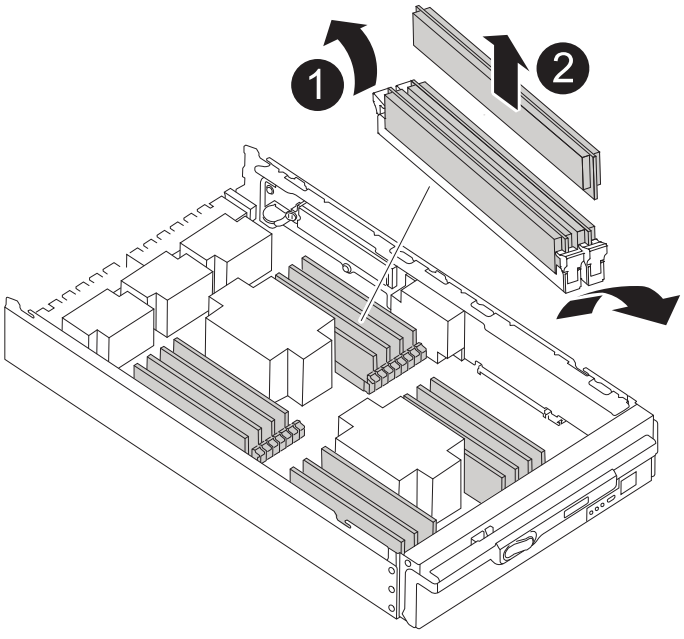
1. If you are not already grounded, properly ground yourself.
2. Locate the DIMMs on your controller module.



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.

Animation - Replace DIMMs



1	DIMM ejector tabs

2	DIMM

4. 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.

5. 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.

6. 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.
7. Close the controller module cover.

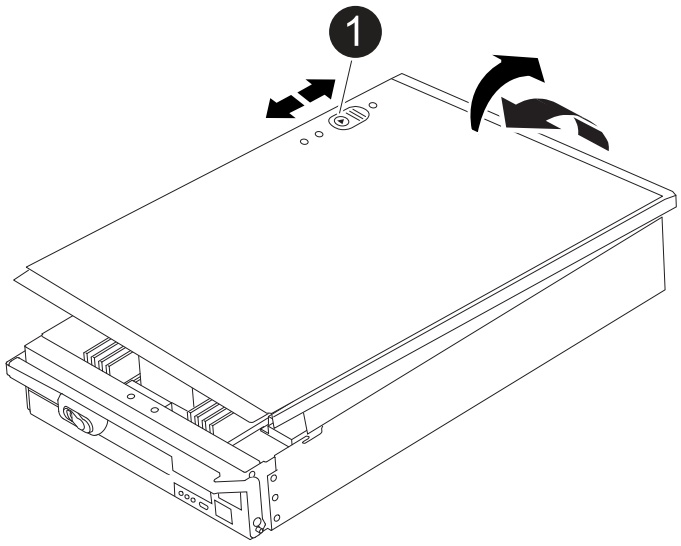
Step 4: Install the controller

After you install the components into the controller module, you must install the controller module back 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.

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.

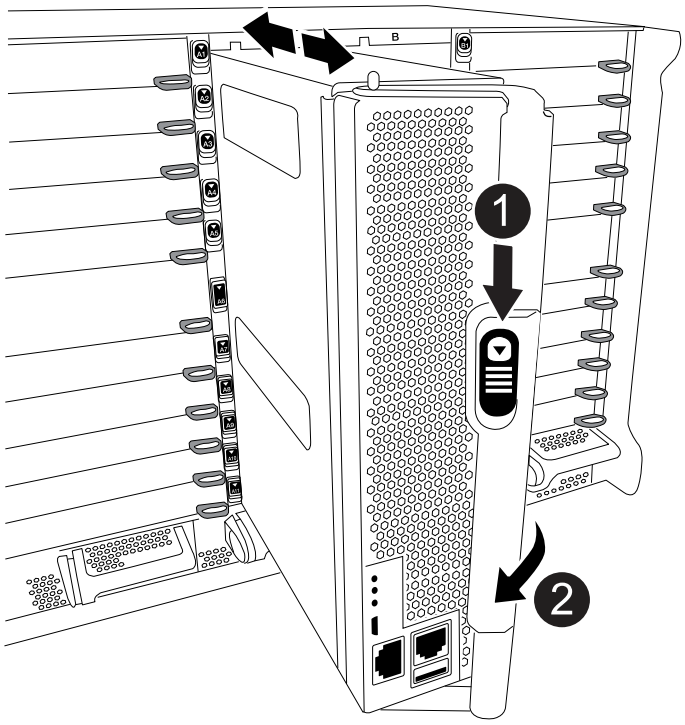


1	Controller module cover locking button

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.

Animation - Install controller



1	Cam handle release button
2	Cam handle

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:
   a. If you have not already done so, reinstall the cable management device.
   b. Firmly push the controller module into the chassis until it meets the midplane and is fully seated.

      The locking latches rise when the controller module is fully seated.

      Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

      The controller module begins to boot as soon as it is fully seated in the chassis.

   c. Rotate the locking latches upward, tilting them so that they clear the locking pins, and then lower them into the locked position.

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Replace the Destage Control Power Module containing the NVRAM11 battery - FAS9500

To hot-swap a destage controller power module (DCPM), which contains the NVRAM11 battery, you must locate the failed DCPM module, remove it from the chassis, and install the replacement DCPM module.

You must have a replacement DCPM module in-hand before removing the failed module from the chassis and it must be replaced within five minutes of removal. Once the DCPM module is removed from the chassis, there is no shutdown protection for the controller module that owns the DCPM module, other than failover to the other controller module.

Step 1: Replace the DCPM module

To replace the DCPM module in your system, you must remove the failed DCPM module from the system and then replace it with a new DCPM module.

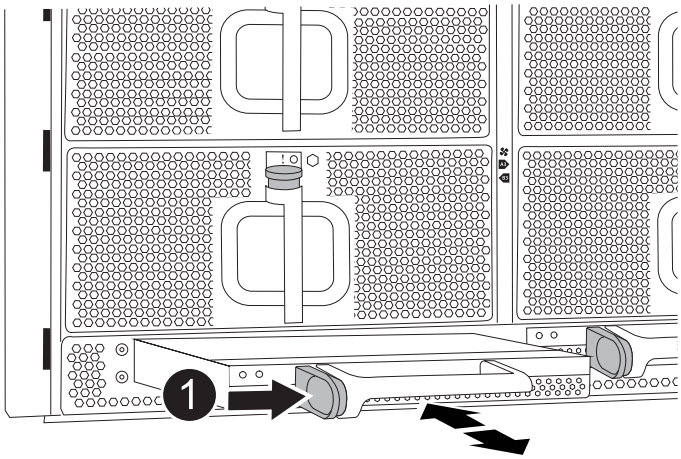
1. If you are not already grounded, properly ground yourself.
2. Remove the bezel on the front of the system and set it aside.
3. Locate the failed DCPM module in the front of the system by looking for the Attention LED on the module.

   The LED will be steady amber if the module is faulty.

   > The DCPM module must be replaced in the chassis within five minutes of removal or the associated controller will shut down.

4. Press the terra cotta locking button on the module handle, and then slide the DCPM module out of the chassis.

   Animation - Remove/install DCPM



1	DCPM module terra cotta locking button

5. Align the end of the DCPM module with the chassis opening, and then gently slide it into the chassis until it clicks into place.

   The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

   The Amber LED flashes four times upon insertion and the green LED also flashes if the battery is providing a voltage. If it does not flash, it will likely need to be replaced.

Step 2: Dispose of batteries

You must dispose of batteries according to the local regulations regarding battery recycling or disposal. If you cannot properly dispose of batteries, you must return the batteries to NetApp, as described in the RMA instructions that are shipped with the kit.

Safety Information and Regulatory Notices

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.

Swap out a fan - FAS9500

To swap out a fan module without interrupting service, you must perform a specific sequence of tasks.

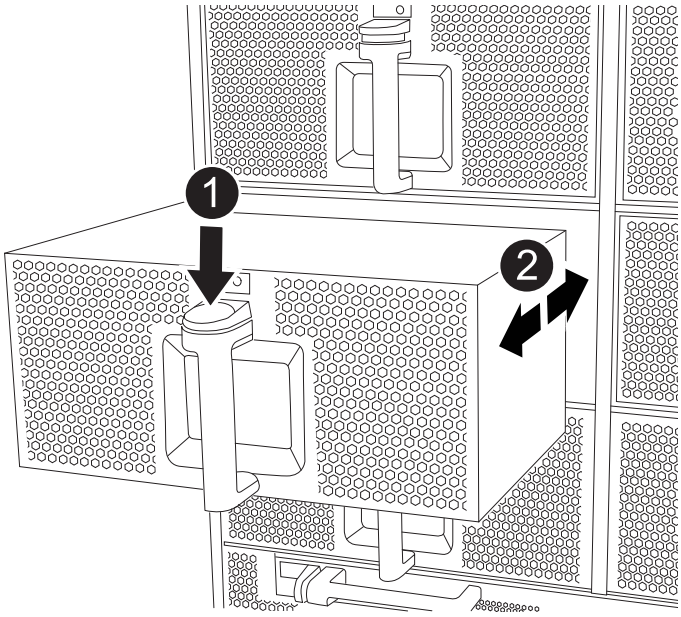
It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
3. Identify the fan module that you must replace by checking the console error messages and looking at the Attention LED on each fan module.
4. Press the terra cotta button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.

   The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.

   Animation - Remove/install fan



1	Terra cotta release button
2	Slide fan in/out of chassis

5. Set the fan module aside.
6. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

   When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

7. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.
8. Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

I/O module

Add an I/O module - FAS9500

You can add an I/O module to your system by either replacing a NIC or storage adapter with a new one in a fully-populated system, or by adding a new NIC or storage adapter into an empty chassis slot in your system.

Before you begin

- Check the NetApp Hardware Universe to make sure that the new I/O module is compatible with your system and version of ONTAP you're running.
- If multiple slots are available, check the slot priorities in NetApp Hardware Universe and use the best one available for your I/O module.
- To non-disruptively add an I/O module, you must take over the target controller, remove the slot blanking cover in the target slot or remove an existing I/O module, add the new or replacement I/O module, and then

giveback the target controller.

- Make sure that all other components are functioning properly.

Step 1: Shut down the impaired controller module

Shut down or take over the impaired controller using one of the following options.

Option 1: Most systems

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

Before you begin

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 command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the impaired controller:

   ```
   storage failover modify -node impaired-node -auto-giveback-of false
   ```

   When you see *Do you want to disable auto-giveback?*, enter `y`.

   a. If the impaired controller cannot be brought up or is already taken over, you must take the HA interconnect link down from the healthy controller before booting up the impaired controller. This will prevent the impaired controller from performing automatic giveback.

      ```
      system ha interconnect link off -node healthy-node -link 0

      system ha interconnect link off -node healthy-node -link 1
      ```

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Option 2: Controller is in a MetroCluster

Do not use this procedure if your system is in a two-node MetroCluster configuration.

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next Step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Add the new I/O modules

If the storage system has empty slots, install the new I/O module into one of the available slots. If all slots are occupied, remove an existing I/O module to make space and then install the new one.

Add I/O module to an empty slot

You can add a new I/O module into a storage system with available empty slots.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the target slot blanking cover:
   a. Depress the lettered and numbered cam latch.
   b. Rotate the cam latch down until it is the open position.
   c. Remove the blanking cover.
3. Install the I/O module:
   a. Align the I/O module with the edges of the slot.
   b. Slide the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
4. If the replacement I/O module is a NIC, cable the module to the data switches.

   Make sure that any unused I/O slots have blanks installed to prevent possible thermal issues.

5. Reboot the controller from the LOADER prompt: *bye*

   This reinitializes the PCIe cards and other components and reboots the node.

6. Give back the node from the partner node. `storage failover giveback -ofnode target_node_name`
7. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto -giveback true`
8. If you are using slots 3 and/or 7 for networking, use the `storage port modify -node <node name> -port <port name> -mode network` command to convert the slot for networking use.
9. Repeat these steps for controller B.
10. If you installed a storage I/O module, install and cable your SAS shelves, as described in Hot-adding a SAS shelf.

Add I/O module to a fully-populated system

You can add an I/O module to a fully-populated system by removing an existing I/O module and installing a new one in its place.

About this task

Make sure you understand the following scenarios for adding a new I/O module to a fully-populated system:

Scenario	Action required
NIC to NIC (same number of ports)	The LIFs will automatically migrate when its controller module is shut down.
NIC to NIC (different number of ports)	Permanently reassign the selected LIFs to a different home port. See Migrating a LIF for more information.
NIC to storage I/O module	Use System Manager to permanently migrate the LIFs to different home ports, as described in Migrating a LIF.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling on the target I/O module.
3. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam latch.

      The cam latch moves away from the chassis.

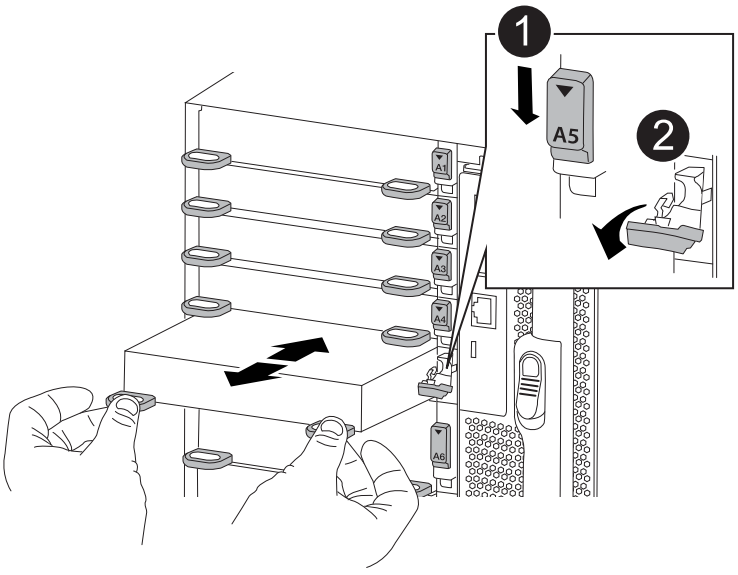
   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

      Animation - Replace an I/O module



1	Lettered and numbered I/O cam latch

2	I/O cam latch completely unlocked

4. Install the I/O module into the target slot:
   a. Align the I/O module with the edges of the slot.
   b. Slide the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin.
   c. Push the I/O cam latch all the way up to lock the module in place.
5. Repeat the remove and install steps to replace additional modules for controller A.
6. If the replacement I/O module is a NIC, cable the module or modules to the data switches.

   This reinitializes the PCIe cards and other components and reboots the node.

7. Reboot the controller from the LOADER prompt:
   a. Check the version of BMC on the controller: `system service-processor show`
   b. Update the BMC firmware if needed: `system service-processor image update`
   c. Reboot the node: `bye`

      This reinitializes the PCIe cards and other components and reboots the node.

      If you encounter an issue during reboot, see BURT 1494308 - Environment shutdown might be triggered during I/O module replacement

8. Give back the node from the partner node. `storage failover giveback -ofnode target_node_name`
9. Enable automatic giveback if it was disabled: `storage failover modify -node local -auto -giveback true`
10. If you added:

If I/O module is a…	Then…
NIC module in slots 3 or 7	Use the `storage port modify -node *<node name> -port *<port name> -mode network` command for each port.
Storage module	Install and cable your SAS shelves, as described in Hot-adding a SAS shelf.

11. Repeat these steps for controller B.

Replace an I/O module - FAS9500

To replace an I/O module, you must perform a specific sequence of tasks.

- You can use this procedure with all versions of ONTAP supported by your system.
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired node

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

Before you begin

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 command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the impaired controller:

   `storage failover modify -node impaired-node -auto-giveback-of false`

   When you see *Do you want to disable auto-giveback?*, enter `y`.

   a. If the impaired controller cannot be brought up or is already taken over, you must take the HA interconnect link down from the healthy controller before booting up the impaired controller. This will prevent the impaired controller from performing automatic giveback.

      `system ha interconnect link off -node healthy-node -link 0`

      `system ha interconnect link off -node healthy-node -link 1`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name` When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Replace I/O modules

To replace an I/O module, locate it within the chassis and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug any cabling associated with the target I/O module.

   Make sure that you label the cables so that you know where they came from.

3. Remove the target I/O module from the chassis:
   a. Depress the lettered and numbered cam button.

      The cam button moves away from the chassis.

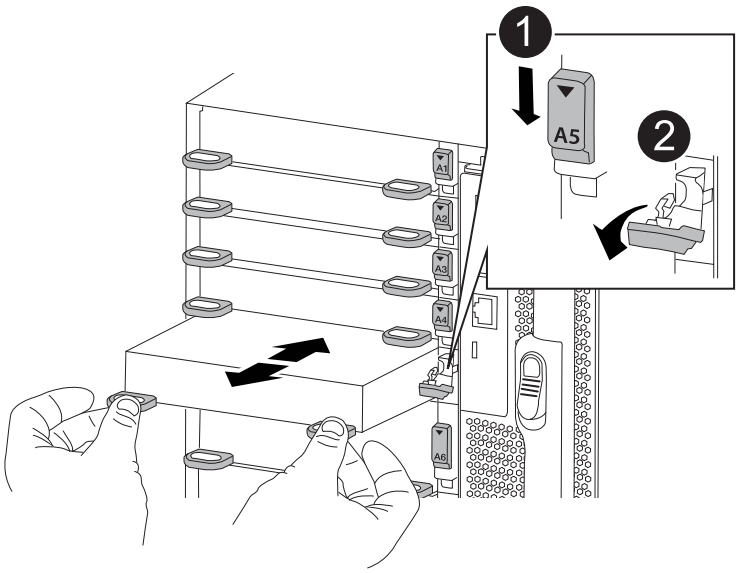
   b. Rotate the cam latch down until it is in a horizontal position.

      The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

   c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

      Make sure that you keep track of which slot the I/O module was in.

      Animation - Remove/install I/O module



1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked

4. Set the I/O module aside.
5. Install the replacement I/O module into the chassis by gently sliding the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.

6. Recable the I/O module, as needed.

Step 3: Reboot the controller after I/O module replacement

After you replace an I/O module, you must reboot the controller module.

If the new I/O module is not the same model as the failed module, you must first reboot the BMC.

Steps

1. Reboot the BMC if the replacement module is not the same model as the old module:
   a. From the LOADER prompt, change to advanced privilege mode: `priv set advanced`
   b. Reboot the BMC: `sp reboot`
2. From the LOADER prompt, reboot the node: `bye`

   This reinitializes the PCIe cards and other components and reboots the node.
3. If your system is configured to support 10 GbE cluster interconnect and data connections on 40 GbE NICs, convert these ports to 10 GbE connections by using the `nicadmin convert` command from Maintenance mode. See Convert 40GbE NIC ports into multiple 10GbE ports for 10GbE connectivity for more information.

   Be sure to exit Maintenance mode after completing the conversion.
4. Return the node to normal operation: `storage failover giveback -ofnode impaired_node_name`
5. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

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 an LED USB module - FAS9500

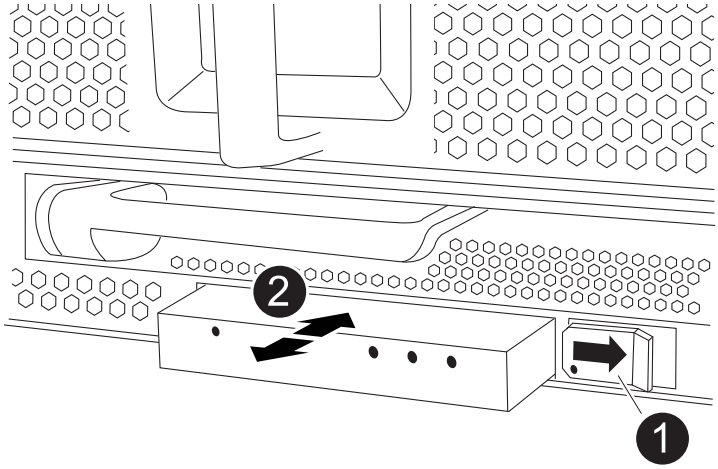
The LED USB module provides connectivity to console ports and system status. Replacement of this module does not require tools and does not interrupt service.

Step 1: Replace the LED USB module

Steps

1. Remove the old LED USB module:

   Animation - Remove/install LED-USB module



1	Locking button
2	USB LED module

a. With the bezel removed, locate the LED USB module at the front of the chassis, on the bottom left side.

b. Slide the latch to partially eject the module.

c. Pull the module out of the bay to disconnect it from the midplane. Do not leave the slot empty.

2. Install the new LED USB module:

a. Align the module to the bay with the notch in the corner of the module positioned near the slider latch on the chassis. The bay will prevent you from installing the module upside down.

b. Push the module into the bay until it is fully seated flush with the chassis.

There is an audible click when the module is secure and connected to the midplane.

Step 2: Return the failed component

1. 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 NVRAM module and/or NVRAM DIMMs - FAS9500

The NVRAM module consists of the NVRAM11 and DIMMs. You can replace a failed NVRAM module or the DIMMs inside the NVRAM module. To replace a failed NVRAM module, you must remove it from the chassis, move the DIMMs to the replacement module, and install the replacement NVRAM module into the chassis.

To replace and NVRAM DIMM, you must remove the NVRAM module from the chassis, replace the failed DIMM in the module, and then reinstall the NVRAM module.

About this task

Because the system ID is derived from the NVRAM module, if replacing the module, disks belonging to the system are reassigned to a new system ID.

Before you begin

- All disk shelves must be working properly.
- If your system is in an HA pair, the partner controller must be able to take over the controller associated with the NVRAM module that is being replaced.
- This procedure uses the following terminology:
  - The impaired controller is the controller on which you are performing maintenance.
  - The healthy controller is the HA partner of the impaired controller.
- This procedure includes steps for automatically reassigning disks to the controller module associated with the new NVRAM module. You must reassign the disks when directed to in the procedure. Completing the disk reassignment before giveback can cause issues.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You cannot change any disks or disk shelves as part of this procedure.

Step 1: Shut down the impaired controller

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

Before you begin

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 command: `system node autosupport invoke -node * -type all -message MAINT=number_of_hours_downh`

   The following AutoSupport command suppresses automatic case creation for two hours: `cluster1:*> system node autosupport invoke -node * -type all -message MAINT=2h`

2. Disable automatic giveback from the console of the impaired controller:

   `storage failover modify -node impaired-node -auto-giveback-of false`

   When you see *Do you want to disable auto-giveback?*, enter `y`.

   a. If the impaired controller cannot be brought up or is already taken over, you must take the HA interconnect link down from the healthy controller before booting up the impaired controller. This will prevent the impaired controller from performing automatic giveback.

      `system ha interconnect link off -node healthy-node -link 0`

      `system ha interconnect link off -node healthy-node -link 1`

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
System prompt or password prompt (enter system password)	Halt or take over the impaired controller from the healthy controller: `storage failover takeover -ofnode` *`impaired_node_name`* When the impaired controller shows Waiting for giveback…, press Ctrl-C, and then respond `y`.

Step 2: Replace the NVRAM module

To replace the NVRAM module, located it in slot 6 in the chassis and follow the specific sequence of steps.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the target NVRAM module from the chassis:
   a. Depress the lettered and numbered cam latch.

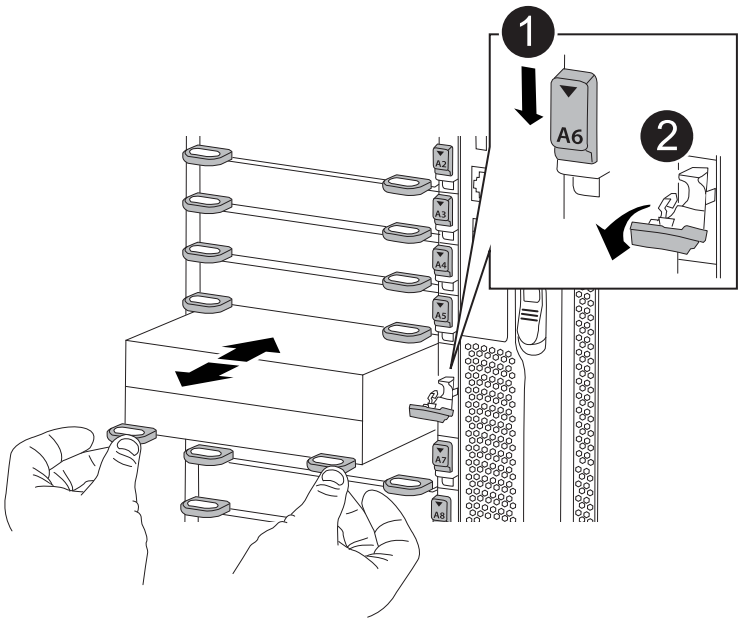
      The cam latch moves away from the chassis.

   b. Rotate the cam latch down until it is in a horizontal position.

      The NVRAM module disengages from the chassis and moves out a few inches.

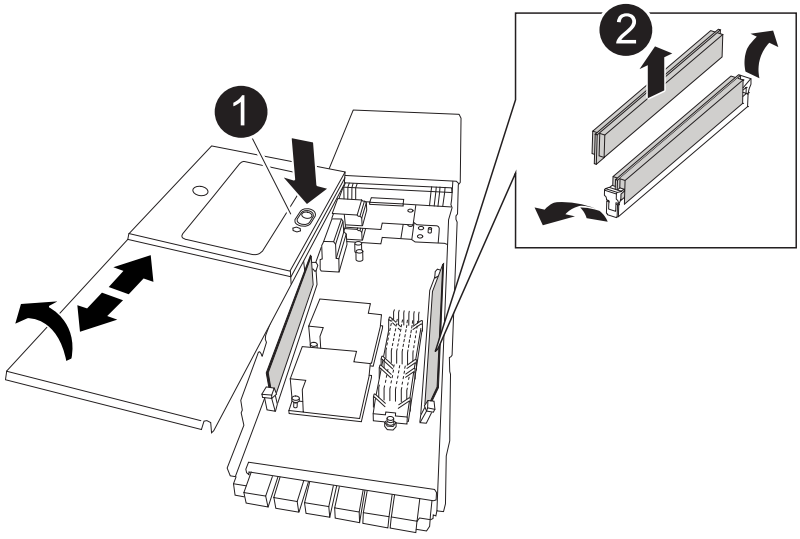
   c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.

      Animation - Replace the NVRAM module



1	Lettered and numbered I/O cam latch
2	I/O latch completely unlocked

3. Set the NVRAM module on a stable surface and remove the cover from the NVRAM module by pushing down on the blue locking button on the cover, and then, while holding down the blue button, slide the lid off the NVRAM module.



1	Cover locking button
2	DIMM and DIMM ejector tabs

4. Remove the DIMMs, one at a time, from the old NVRAM module and install them in the replacement NVRAM module.
5. Close the cover on the module.
6. Install the replacement NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 6.
   b. Gently slide the module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.

Step 3: Replace a NVRAM DIMM

To replace NVRAM DIMMs in the NVRAM module, you must remove the NVRAM module, open the module, and then replace the target DIMM.

Steps

1. If you are not already grounded, properly ground yourself.
2. Remove the target NVRAM module from the chassis:
   a. Depress the lettered and numbered cam latch.

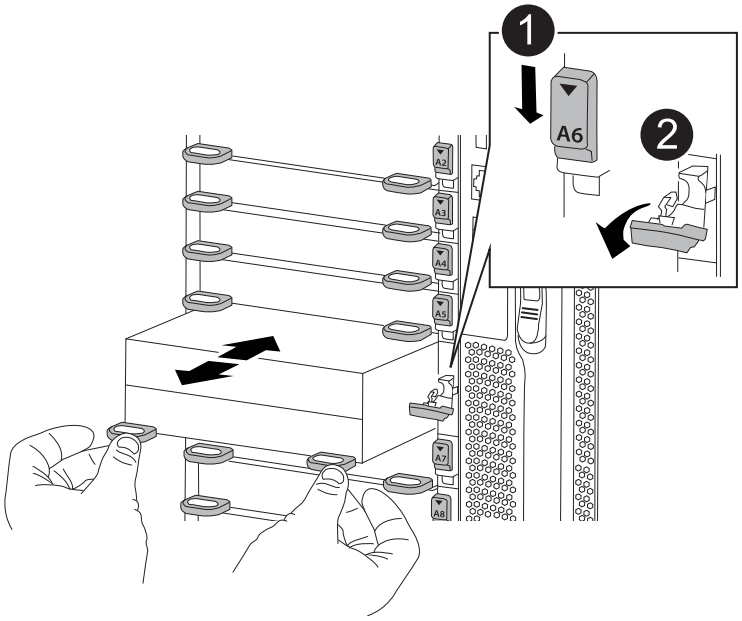
The cam latch moves away from the chassis.

b. Rotate the cam latch down until it is in a horizontal position.

The NVRAM module disengages from the chassis and moves out a few inches.

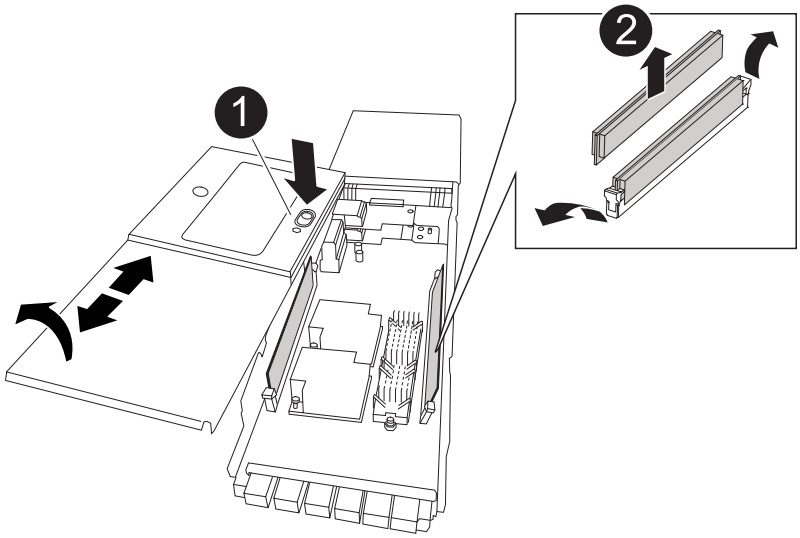
c. Remove the NVRAM module from the chassis by pulling on the pull tabs on the sides of the module face.

Animation - Replace the NVRAM module



1	Lettered and numbered I/O cam latch
2	I/O latch completely unlocked

3. Set the NVRAM module on a stable surface and remove the cover from the NVRAM module by pushing down on the blue locking button on the cover, and then, while holding down the blue button, slide the lid off the NVRAM module.



1	Cover locking button
2	DIMM and DIMM ejector tabs

4. Locate the DIMM to be replaced inside the NVRAM module, and then remove it by pressing down on the DIMM locking tabs and lifting the DIMM out of the socket.
5. Install the replacement DIMM by aligning the DIMM with the socket and gently pushing the DIMM into the socket until the locking tabs lock in place.
6. Close the cover on the module.
7. Install the NVRAM module into the chassis:
   a. Align the module with the edges of the chassis opening in slot 6.
   b. Gently slide the module into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the module in place.

Step 4: Reboot the controller after FRU replacement

After you replace the FRU, you must reboot the controller module.

To boot ONTAP from the LOADER prompt, enter `bye`.

Step 5: Reassigning disks

You must confirm the system ID change when you boot the replacement controller and then verify that the change was implemented.

Disk reassignment is only needed when replacing the NVRAM module and does not apply to NVRAM DIMM replacement.

Steps

1. If the 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 and entering y if you are prompted to override the system ID due to a system ID mismatch.
3. Wait until the Waiting for giveback… message is displayed on the console of the controller with the replacement module and then, from the healthy controller, verify that the new partner system ID has been automatically assigned: `storage failover show`

   In the command output, you should see a message that the system ID has changed on the impaired controller, showing the correct old and new IDs. In the following example, 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. 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.

      For more information, see the Manual giveback commands topic to override the veto.

   b. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

      The output from the `storage failover show` command should not include the System ID changed on partner message.

5. Verify that the disks were assigned correctly: `storage disk show -ownership`

   The disks belonging to the replacement controller should show the new system ID. In the following example, the disks owned by node1 now show the new system ID, 151759706:

```
node1:> storage disk show -ownership

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID
Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------
---------  ---
1.0.0  aggr0_1  node1 node1  -        151759706  151759706  -
151759706 Pool0
1.0.1  aggr0_1  node1 node1           151759706  151759706  -
151759706 Pool0
.
.
.
```

6. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

   The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

7. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

   This is required if both of the following are true:

   - The MetroCluster configuration is in a switchover state.
   - The replacement controller is the current owner of the disks on the disaster site.

     See Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration for more information.

8. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id            cluster node           configuration-state
-----------            ---------------------- --------------
-------------------
1 node1_siteA          node1mcc-001           configured
1 node1_siteA          node1mcc-002           configured
1 node1_siteB          node1mcc-003           configured
1 node1_siteB          node1mcc-004           configured

4 entries were displayed.
```

9. Verify that the expected volumes are present for each controller: `vol show -node node-name`
10. If storage encryption is enabled, you must restore functionality.
11. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

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 - FAS9500

Swapping out a power supply involves turning off, disconnecting, and removing the 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.

About this task

- The power supplies are redundant and hot-swappable.
- This procedure is written for replacing one power supply at a time.

> It is a best practice to replace the power supply within two minutes of removing it from the chassis. The system continues to function, but ONTAP sends messages to the console about the degraded power supply until the power supply is replaced.

- There are four power supplies in the system.
- Power supplies are auto-ranging.

> Do not mix PSUs with different efficiency ratings. Always replace like for like.

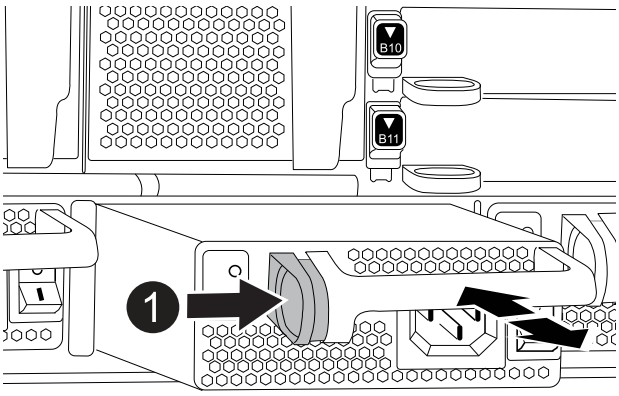
Steps

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.
4. Press and hold the terra cotta button on the power supply handle, and then pull the power supply out of the chassis.

   When removing a power supply, always use two hands to support its weight.

   Animation - Remove/install PSU



1	Locking button

5. Make sure that the on/off switch of the new power supply is in the Off position.
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 until it locks into place.

   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. Reconnect the power supply cabling:
   a. Reconnect the power cable to the power supply.
   b. Secure the power cable to the power supply using the power cable retainer.

      Once power is restored to the power supply, the status LED should be green.

8. Turn on the power to the new power supply, and then verify the operation of the power supply activity LEDs.

   The green power LED lights when the PSU is fully inserted into the chassis and the amber attention LED flashes initially, but turns off after a few moments.

9. 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 - FAS9500

You replace the real-time clock (RTC) battery in the controller module so that your system's services and applications that depend on accurate time synchronization continue to function.

- You can use this procedure with all versions of ONTAP supported by your system
- All other components in the system must be functioning properly; if not, you must contact technical support.

Step 1: Shut down the impaired node

To shut down the impaired controller, you must determine the status of the controller and, if necessary, take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from priv advanced mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

  Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see Synchronize a node with the cluster.

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

   ```
   system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
   ```

   The following AutoSupport message suppresses automatic case creation for two hours:

   ```
   cluster1:> system node autosupport invoke -node * -type all -message MAINT=2h
   ```

2. Disable automatic giveback:
   a. Enter the following command from the console of the healthy controller:

      ```
      storage failover modify -node impaired_node_name -auto-giveback false
      ```

   b. Enter `y` when you see the prompt *Do you want to disable auto-giveback?*
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying…	Then…
The LOADER prompt	Go to the next step.
Waiting for giveback…	Press Ctrl-C, and then respond `y` when prompted.

If the impaired controller is displaying…	Then…
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: `storage failover takeover -ofnode impaired_node_name -halt true` The *-halt true* parameter brings you to the LOADER prompt.

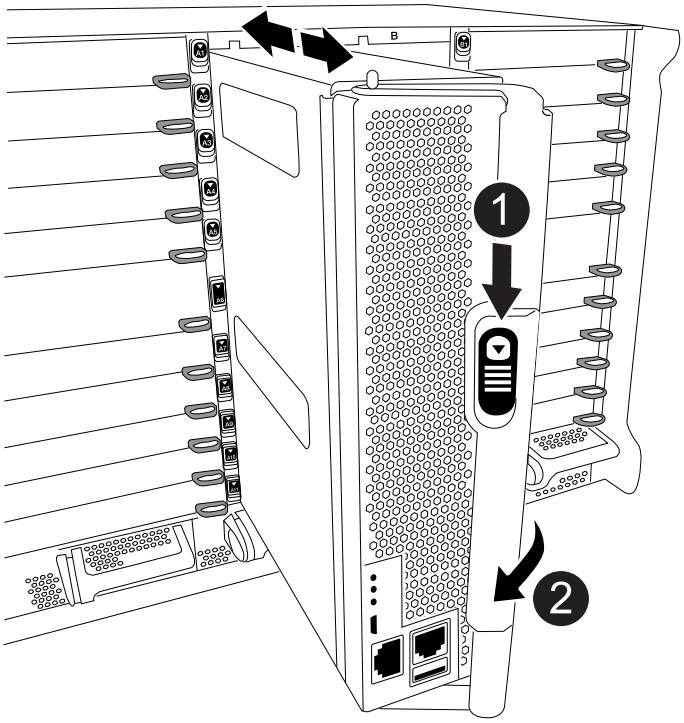
Step 2: Remove the controller

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. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
3. Slide the terra cotta button on the cam handle downward until it unlocks.

   Animation - Remove controller module

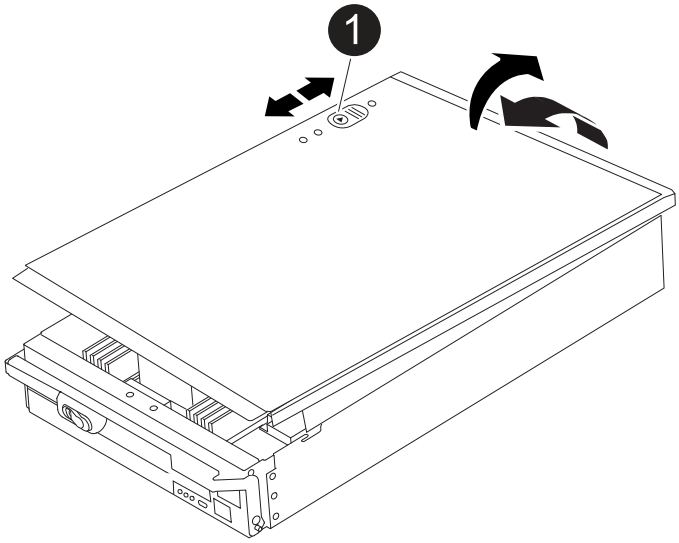


1	Cam handle release button
2	Cam handle

4. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

   Make sure that you support the bottom of the controller module as you slide it out of the chassis.

5. Place the controller module lid-side up on a stable, flat surface, press the blue button on the cover, slide the cover to the back of the controller module, and then swing the cover up and lift it off of the controller module.



1	Controller module cover locking button

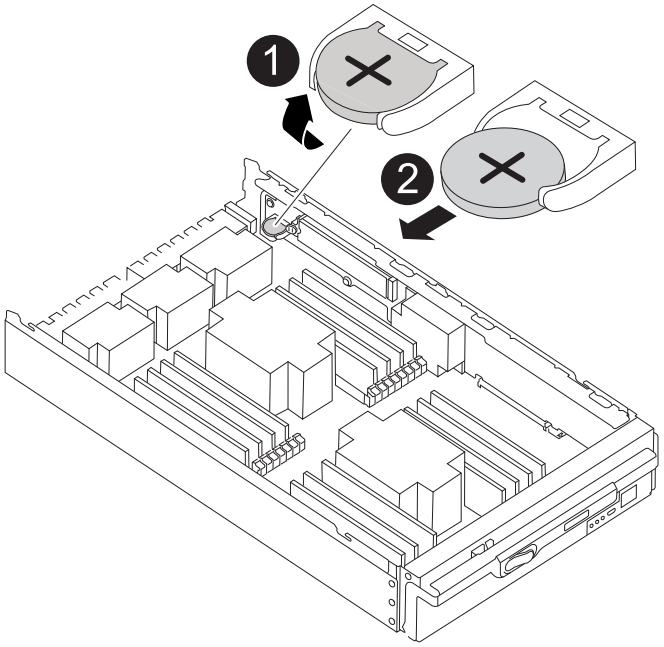
Step 3: Replace the RTC battery

To replace the RTC battery, you must locate the failed battery in the controller module, remove it from the holder, and then install the replacement battery in the holder.

Steps

1. If you are not already grounded, properly ground yourself.
2. Locate the RTC battery.

   Animation - Replace RTC battery



1	Rotate battery up
2	Slide battery out from housing

Steps

1. 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.

2. Remove the replacement battery from the antistatic shipping bag.
3. Locate the empty battery holder in the controller module.
4. Note the polarity of the RTC battery, and then insert it into the holder by tilting the battery at an angle and pushing down.
5. Visually inspect the battery to make sure that it is completely installed into the holder and that the polarity is correct.
6. Reinstall the controller module cover.

Step 4: Reinstall the controller module and set time/date

After you replace the RTC battery, you must reinstall the controller module. If the RTC battery has been left out of the controller module for more than 10 minutes, you may have to reset the time and date.

Steps

1. If you have not already done so, close the air duct or controller module cover.
2. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Do not completely insert the controller module in the chassis until instructed to do so.

3. Recable the system, as needed.

   If you removed the media converters (QSFPs or SFPs), remember to reinstall them if you are using fiber optic cables.

4. 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.

If your system stops at the boot menu, select the option for "Reboot node" and respond y when prompted,
then boot to LOADER by pressing `Ctrl-C`.

1. Reset the time and date on the controller:
   a. Check the date and time on the healthy node with the `show date` command.
   b. At the LOADER prompt on the target node, 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 node.
2. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components and let the node reboot.
3. Return the node to normal operation by giving back its storage: `storage failover giveback -ofnode impaired_node_name`
4. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto -giveback true`

Step 5: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the Part Return and Replacements page for further information.

Key specifications for FAS9500

The following are select specifications for the FAS9500. Visit NetApp Hardware Universe

(HWU) for a complete list of FAS9500 specifications. This page is reflective of a single high availability pair.

Key specifications for FAS9500

Platform Configuration: FAS9500 Single Chassis HA Pair

Max Raw Capacity: 15.0000 PB

Memory: 2048.0000 GB

Form Factor: 8U chassis with 2 HA controllers

ONTAP Version: b_startONTAP: 9.16.1P2b_end

PCIe Expansion Slots: 20

Minimum ONTAP Version: ONTAP 9.10.1P3

Scaleout Maximums

Type	HA Pairs	Raw Capacity	Max Memory
NAS	12	180.0 PB / 159.9 PiB	24576 GB
SAN	6	90.0 PB / 79.9 PiB	12288 GB
HA Pair		15.0 PB / 13.3 PiB	2048.0000

IO

Onboard IO

No onboard IO data.

Total IO

Protocol	Ports
Ethernet 100 Gbps	32
Ethernet 25 Gbps	64
Ethernet 10 Gbps	64
FC 32 Gbps	64
NVMe/FC 32 Gbps	64
	0
SAS 12 Gbps	64

Management Ports

Protocol	Ports

Ethernet 1 Gbps	2
RS-232 115 Kbps	6
USB 12 Mbps	2

Storage Networking Supported

CIFS;
FC;
iSCSI;
NFS v3;
NFS v4.0;
NFS v4.1;
NFS v4.2;
NVMe/FC ;
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: 8004 BTU/hr
- Worst-case Power: 9937 BTU/hr
- Weight: 220.5 lb 100.0 kg
- Height: 8U
- Width: 19” IEC rack-compliant (17.7” 45.0 cm)
- Depth: 28.8” (36.8” with cable management bracket)
- Operating Temp/Altitude/Humidity: 10°C to 35°C (50°F to 95°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): 7.4; Sound pressure (LpAm) (bystander positions): 65.0 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 for external shelves;

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.